An Anatomy of Managerial Attention: Evidence From Retail*

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

July 24, 2025

Abstract

How does a manager's ability to manage attention–both their own scarce attention and that of their subordinates–affect productivity? We study this question using administrative data from a large retail firm operating over 200 stores and employing more than 20,000 workers across Colombia. Leveraging both the inherent complexity of store management and a policy of rotating middle managers across stores, we examine the role managers play in driving productivity. Specifically, we ask: Who are the managers that increase sales and what do they do differently? The average store carries 55,000 products and works with 800 suppliers, making it impossible for managers to oversee everything at once. We find that managers who successfully increase sales navigate this complexity through three strategies that less effective managers fail to implement. First, they reduce stockouts. Second, they decrease the size of the inventory—both in terms of overall value and the inventory-to-sales ratio. Third, they execute more effective pricing decisions. We also observe that the arrival of a high-performing manager leads to changes in store organization, reallocating personnel toward back-of-store roles. We complement these results with a survey capturing the managerial style and traits of managers at our partner firm, allowing us to better understand which types of managers drive superior performance.
Keywords: managerial attention, span of attention, retail, Colombia
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, Harvard Business School; Digital Reskilling Lab; jtamayo@hbs.edu

1 Introduction

A large body of research has established that managers account for a substantial share of variation in firm performance (Bloom and Van Reenen, 2007). In particular, managerial fixed effects have been shown to explain a significant portion of productivity differences across firms and establishments (Lazear et al., 2015; Metcalfe et al., 2023). Dessein and Santos (2021) further develop a model in which, under conditions of high complexity, differences in how managers allocate their attention amplify these fixed effects—suggesting that complexity can intensify the impact of individual managerial styles on performance.

Management involves a broad array of tasks, and managers must constantly make difficult decisions about how to allocate their scarce attention across competing priorities such as employee supervision, hiring, firing, promotions, pricing, product administration and general operations. However, managers face significant informational constraints and often do not know where their own marginal effort—or that of their subordinates—has the highest impact on outcomes. These span-of-control and cognitive hierarchy problems can lead to the misallocation of attention and effort (Adhvaryu et al., 2022; Bandiera et al., 2014).

Importantly, these decisions unfold in environments where attentional bandwidth is limited. It is well-established that decision-makers are bounded in their cognitive capacity and often rely on heuristics to manage complexity (Simon, 1973). In such environments, having more choices or more information does not necessarily lead to better decisions—in fact, large assortments can induce cognitive strain and reduce engagement (Iyengar and Lepper, 2000). This suggests that what sets effective managers apart may not only be what they know or what they do, but also how they allocate attention and simplify decision-making under constraints.

A growing literature has begun to unpack the specific types of skills that matter for effective management, moving beyond broad measures such as education or tenure. For instance, Weidmann et al. (2024) and Adhvaryu et al. (2023) highlight how interpersonal and cognitive skills are closely tied to managerial effectiveness and productivity outcomes. Another set of recent papers opens the black box of *what* managers actually do to influence performance. Rather than focusing solely on traits, these studies document specific actions and behaviors—such as performance monitoring (Adhvaryu et al., 2024) and rapport-building with subordinates (Adhvaryu et al., 2025)—that distinguish high-performing managers from their peers.

Building upon this literature, we examine how differences in managers' ability—and willingness—to navigate complexity translate into differences in firm performance. We focus on a setting in which complexity is not just present, but defining: large retail stores managing tens of thousands of products, hundreds of suppliers, and dozens of employees.

This complexity is especially pronounced in core operational areas such as supply chain coordination, inventory management, and pricing. In our setting, each store works with an average of 826 suppliers and carries 54,875 products, making it impossible for managers to monitor every item. Yet, they remain responsible for ensuring that products are consistently stocked, competitively priced, and aligned with firm-wide goals. These conditions create a natural testbed for studying how managers cope with complexity: who succeeds, what decisions they prioritize, and how attention is allocated when it cannot be spread evenly.

To answer this question, we use rich administrative data from a large Colombian retailer. Our dataset includes scanner-level sales data, detailed records on inventory, supplier characteristics, and supplier-store transactions. Crucially, we exploit the firm's policy of randomly rotating managers across stores between January 2017 and June 2020. Leveraging this quasi-experimental variation, we first estimate an Abowd-Kramarz-Margolis (AKM) model using store sales to separately identify manager and store fixed effects. We then classify managers as high type or low type, with high type managers defined as those whose estimated fixed effects fall above the median of the distribution.

Next, we combine these estimates with personnel records to characterize what high-type (good) 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 internal data to examine how these high type managers operate—focusing in particular on inventory management, pricing, and other internal metrics that reflect how they handle complexity on the ground.

We find that effective management of 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 critical areas. Specifically, we identify three mechanisms through which high-performing managers navigate this complexity:

First, good managers reduce stockouts by 16.7 percent, directly enhancing sales by ensuring product availability.[1] Second, they simultaneously decrease inventory size, with 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. Finally, they make more effective pricing decisions, as evidenced by increased revenues immediately following

[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.

a price change.

Upon a good manager's arrival, we also observe a significant restructuring, with more personnel being allocated to the non-customer-facing roles relating to pricing, inventory, and general supply chain management. Specifically, good managers reduce turnover among back-of-store 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, 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. For example, while not statistically significant, we observe that good managers are nearly three times as likely to list basic math skills as critical to their job compared to managers of lower-performing stores—suggesting that quantitative skills may facilitate better decision-making in complex environments.

The rest of the paper is organized as follows. Section 2 provides background on the institutional context, store operations, and the organizational structure of our partner firm. We also present evidence from a survey of managers as additional context. Section 3 describes the administrative data used in our analysis. Section 4 presents our empirical strategy, beginning with the estimation of manager fixed effects and continuing with our main event-study design and mechanism analysis. Section 6 offers robustness checks and addresses potential concerns around identification. Section 7 concludes with a discussion of the broader implications of our findings for managerial effectiveness and organizational design in complex environments. An online appendix provides additional results and robustness exercises.

2 Context

Our partner firm is the Colombian branch of one of the largest supermarket chains in Latin America. We leverage administrative data from the firm's 220 locations in Colombia from 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 central 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 oversee production 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 or Costco), the type of products sold varies, with some stores selling everything from non-perishable goods like flat-screen TVs and car batteries to perishable goods like sushi and donuts. 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 (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).

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 customer-facing 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.

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).

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.9 shows that Coordinators and

Leaders are far more common than Bosses or Directors. The responsibilities and actions of 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, we exclude the lowest level of managers (coordinators), re-run the AKM analysis, and replicate all the results presented in the main body of the paper in Appendix A.4.1.[2]

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 are far more rectangular 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, for any given product a manager must decide *when*, *how much*, and *from which suppliers* to order. 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. However, even within a store, demand is sensitive

[2]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 attenuated when we include coordinators. We interpret this attenuation as confirmation of our anecdotal evidence that these lower-level managers are less likely to be engaging in hiring and firing or other departmental restructuring decisions.

to everything from weather, time of year, and day of the month.

To appreciate the complexity of this task, imagine a toy version of this problem in which a manager's only duty is to worry about maximizing sales for a single product (i.e., this manager doesn't worry about personnel issues like training or hiring). In this example, the manager's only job is to ensure that a single good is appropriately stocked (i.e., she wants to avoid excess inventory while avoiding stockouts) and priced (i.e., regular sales should generate enough margin to make the product worth selling, but not so much so that consumers seek out the good via other channels) so as to maximize store sales.

To fix ideas, suppose that product is a soccer jersey for Colombia's national team. After looking through historical data to determine the right proportion of shirts to order in each size, our manager might predict that her store will experience increased demand before each game. If so, she may consider it a good use of time to look up the national team's schedule and place orders ahead of time to ensure her store has enough jerseys for each game, perhaps ordering a few more jerseys for games she considers more important than others. Once the jerseys arrive (not too far in advance so the store doesn't carry excess inventory, but also far enough in advance so that consumers can purchase the jerseys in time for the big game) she must consider how best to price the jerseys. The profit-maximizing price may be a function of both time and the team's success; It likely makes sense to charge a high price the day of the game, but it's unclear how to price the jerseys the day after the game: a loss may portend a price cut, while a win may require an increase.

In this example our manager is likely to do a good job maximizing sales for her one product. Of course, as the manager's attention becomes more divided, she will need to spend less time maximizing sales on soccer jerseys to ensure she maximizes sales on all the products she manages. Intuitively, there is a trade-off between extracting a higher margin from a lower number of products and extracting a lower margin from a higher number of products.

Thus, managers who are most likely to increase store sales are those who optimize along two dimensions: (1) how many goods to spend time thinking about (i.e., forecasting, stocking, and pricing), and (2) how much inventory to hold. 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 store-level sales. Conversely, spending too little attention on all products increases the chances that a store carries excess inventory or experiences excessive stockouts. As the number of products (and/or suppliers) increases, so too does the likelihood that attention constraints bind.

2.4.1 Survey of Managers

In 2023, Good Business Lab, in conjunction with Weidmann et al. (2024) surveyed the universe of managers at our partner firm (achieving a 77.6% response rate), capturing information on managers' demographics, job histories, psychometric profiles, personality traits, managerial practices, leadership styles, job descriptions and job satisfaction. Because this survey was conducted for an unrelated project it has two limitations:

First, the survey was conducted nearly three years after the end of our administrative data sample.[3] 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.[4] 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 results include all middle managers.Thus, even though our empirical results include all middle management, our survey evidence comes from only the top 20 percent of middle management[5].

While these limitations mean we must take care to not overanalyze the findings, this survey guided the questions we asked in the empirical analysis presented in section 4, making it a useful, albeit suggestive, guide to those findings. We present the findings from subset of relevant questions 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 as the dependent variable. Thus, a large and positive coefficient means there is a strong, positive correlation between store productivity and that outcome. Panel A presents managers' results on standardized tests of cognitive and emotional ability (e.g., the Ravens progressive matrices, a digit-span recall test, a test of simple arithmetic ability, and the reading the mind in the eyes test) along with self-reported measures of a manager's leadership style. Perhaps surprisingly, we observe that high-performing managers tend to receive low scores on these tests of cognitive and emotional ability. On the other hand, higher-performing managers are more likely to report that traits like friendliness and flexibility are critical skills in their

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

[4]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 more likely to have responded to the 2023 survey (p=0.29)).

[5]Specifically, managers used in the AKM includes approximately 3000 individuals, while the survey targeted only the 500 senior-most store managers.

Figure 1: Survey Results

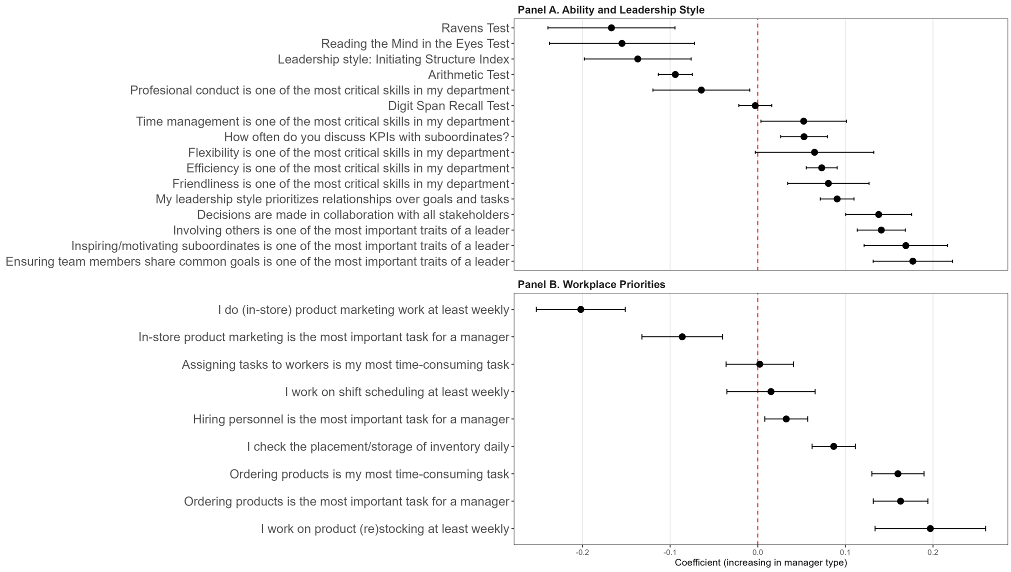

Notes: Each row of Figure 1 corresponds to a question (or index from a set of questions) asked to 183 managers at our partner firm. We run a separate regression for each row, using standardized manager fixed effects (on log productivity) from the AKM decomposition described in Section 4 as the explanatory variable, and survey responses as the dependent variable.

department and that involving, motivating, inspiring, and collaborating with subordinates are important traits in leaders. 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.

3 Data and Descriptive Statistics

We use administrative data from the supermarket chain's 220 stores that operated across all of Colombia 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.[6] 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

[6]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.

to excess demand.[7] 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[8], and because of potentially negative impacts on customer satisfaction. Prior to receiving a good manager, stores experience an average of 886 stockouts per month costing approximately $407K USD each 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.

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.

[7]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.

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

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.

4 Empirical Analysis

We first establish the existence of large manager fixed effects on sales and productivity. We then observe that the arrival of a high-type (good) manager to an underperforming (bad) store does, in fact, increase store productivity and sales. Finally, we uncover the specific mechanisms managers used to achieve such increases and then consider how the organizational structure responds to the arrival of a good manager.

4.1 Measuring Value-Added of a Manager

We start with the, by now well-established, observation that individual managers may exert large effects on productivity (Lazear et al., 2015; Metcalfe et al., 2023) and classify managers according to their fixed effects on store performance. These fixed effects will be biased unless moves are uncorrelated with time-varying components of store productivity. Consequently, 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 quasirandom (see Section 2), permitting us to estimate the differences in sales attributable to individual managers. We use the AKM framework (Abowd et al., 1999) to estimate manager and store fixed effects on 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}:

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

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.

This assumption does not restrict the possibility that managers are assigned to stores based on permanent components of managerial ability θ_i or store components α_j. While our partner firm has told us that the assignment of managers to stores is random (specifically, managers are shuffled across stores periodically as part of their career development), we note that even if sorting on these (permanent) fixed effects does occur (e.g., managers prefer 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.11 shows the results of the AKM decomposition, suggesting that manager fixed effects explain more than 12 percent of log productivity.

4.1.1 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.13 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.13 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.14 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.14 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.8 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.

4.1.2 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 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 35 percent in Alvarez et al. (2018). Next, we present the bias correction of Andrews et al. (2008) in Table A.13. 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.13. Again, the results are very similar to our baseline model.

Taken together, these tests suggest that use of the AKM framework is appropriate in our setting.

4.2 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. [9]

[9] We believe this to be the most intuitive and interesting definition of treatment: 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 will make a difference. Even so, we consider the effects of several alternative definitions of treatment in Table A.12. 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.8 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

4.3 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 performance measure of 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).

4.4 Identification

Our identifying assumption is that the matching of managers to stores is exogenous (see Section 4.1). Our partner firm has assured us that the rotation of managers across stores is mechanical and assignments are made without considering manager traits or store needs. While we present the formal tests to help us rule out endogeneity in manager mobility above in Section 4.1.1, in this section we present the results of test for the existence of pre-trends. 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, suggesting that the rotation of managers across stores is indeed quasirandom.

falls to (a still significant) 0.089 log points, and falls further to a (marginally significant) 3.4% increase when we expand the definition to include the top 50% of differences. Treatment effects are negative when we consider the bottom 50% of differences. Finally, when we consider the arrival of 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 manager is unlikely to cause major changes.

Table 3: Check for Exogeneity of Manager Arrival

Variable	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
Stockout value (millions COP)	-17.280	-0.898	1353.435	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: Table 3 presents the coefficients from a regression of each variable of interest on the time to a manager's arrival. Clustering is done at the store level.

4.5 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. Table 4 shows that log sales and log productivity both increase by more than 10 percent after the arrival of a good manager. Figures 2 and A.9 show that these increases occur shortly after the good manager's arrival and increase over time.

Table 4: 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 4 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 2: Effect of Treatment on Log Sales

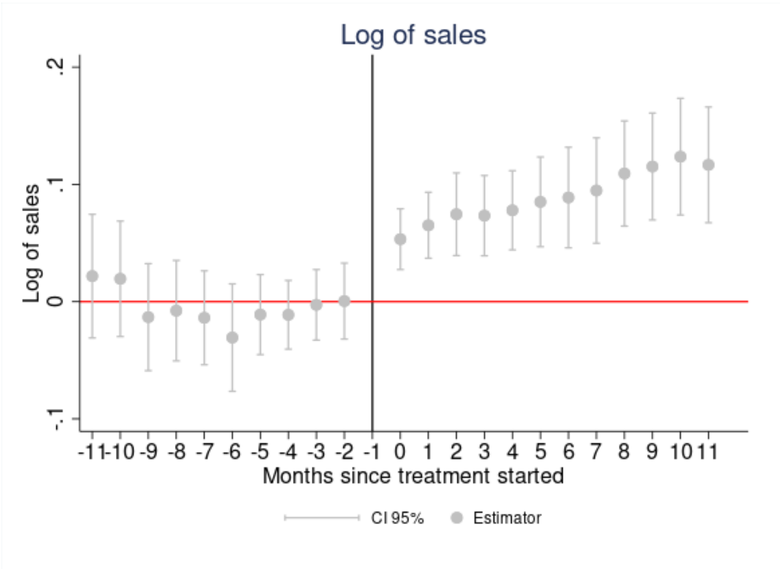

Notes: Figure 2 shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's arrival on log sales. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The vertical line represents the arrival of the good manager.

4.6 What do good managers do?

In this section we ask what specific mechanisms are associated with the large increase in sales and productivity. 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[10]) *price changes*.

4.6.1 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 5). Fewer stockouts means that the total value of stocked-out inventory decreases by $53K USD, or 16.1 percent (column 2). Panels (a) and (b) of Figure A.10 show that both effects begin shortly after the arrival of a good manager and increase over time.

Table 5: Effects on stockouts and on-hand inventory value

	(1) Stockouts	(2) Value of Stocked-out Inventory (USD)	(3) Inventory value (USD)	(4) Inventory value over sales
Treatment effect	-87.82***	-53,180***	-28,996***	-0.0290***
	(32.59)	(15,587)	(8,779)	(0.00579)
Observations	8,890	8,890	8,890	8,890
R-squared	0.886	0.884	0.986	0.808
Treated Stores	87	87	87	87
Stores	220	220	220	220
Mean of Dv	763.3	329,642	404,962	0.247
Relative Effect	-11.5%	-16.1%	-7.2%	-11.7%

Notes: Table 5 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%.

4.6.2 Action 2: Good managers make their inventory leaner

Column 3 of Table 5 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 3 shows that this drop occurs in the period after the manager's arrival and that inventory values further decrease post-arrival. Column 4 of Table 5 shows that

[10]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.

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 3: Value of inventory

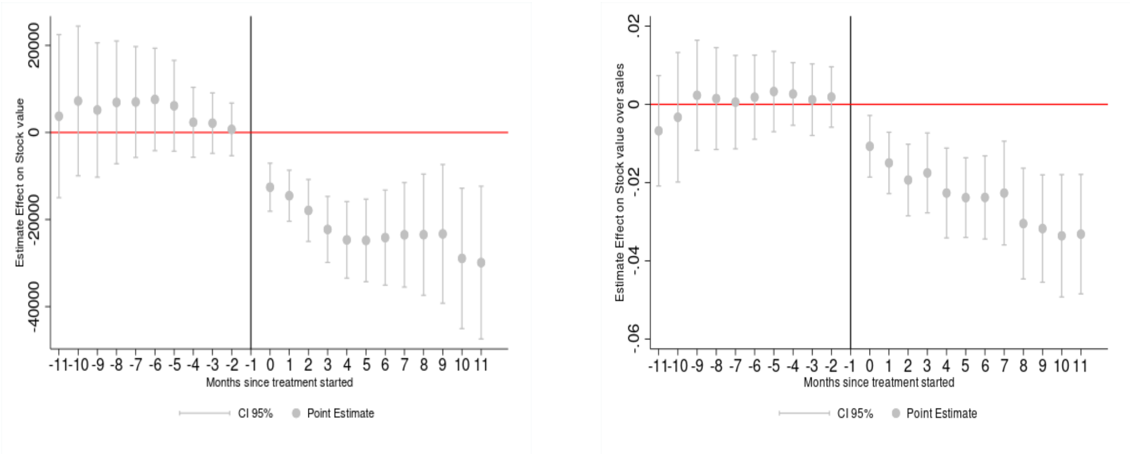

(a) Total value of inventory (USD)

(b) Value of inventory over sales

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's arrival on the total value of inventory and 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. The vertical line represents the arrival of a good manager.

4.6.3 Action 3: Good managers make better pricing decisions

Table 6 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.11 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 6: 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 6 shows the point estimates of log sales of product that did and did not undergo price changes, and the number of price changes in a month. 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%.

4.7 What do good managers give up?

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 4 shows that this reduction in trainings begins immediately following the arrival of a new manager.

Figure 4: Effect on number of trainings

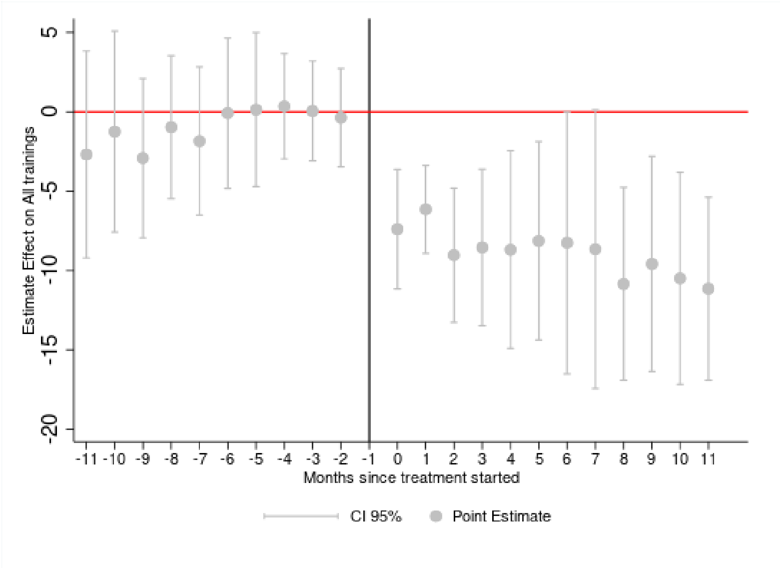

Because we do not directly observe scheduling, absenteeism is the closest proxy for how much attention a manager allocates to worker scheduling. We observe that the arrival of a good manager precedes an increase in overall absenteeism. Neither hiring nor turnover change following the arrival of a good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

4.8 Personnel Management and Changes to Store-level Organizational Structure

In this section we discuss shows 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.

Bad stores prior to arrival of a good manager 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 7 shows that the arrival of a good manager to a bad store causes backroom departments to gain an average of 1.8 workers. Panel (a) of Figure A.16a 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 8 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 including the lowest level of managers (coordinators) in the AKM decomposition of moving managers. When we compare Table A.14 to Table 7 the direction of our point estimates are unchanged, but the coefficients in the results that include coordinators tend to be smaller. Specifically, when we exclude coordinators (the junior-most level of managers who are likely not in charge of hiring or firing) we see that the treatment effect on the number of workers in backroom department shrinks from 1.8 to 1.3. Similarly, the effect on the number of managers in backroom departments falls from 0.51 to 0.36. Comparing Tables 8 and A.15 shows the same pattern: the magnitudes of both turnover and hiring in backroom departments shrink. This pattern is consistent with the interpretation that the arriving manager is the driving force behind these organizational changes.

Table 7: Workers and Managers (backroom/customer-facing)

VARIABLES	(1) Total workers (backroom)	(2) Total workers (customer-facing)	(3) Total managers (backroom)	(4) Total managers (customer-facing)	(5) Ratio managers to non-managers (backroom)	(6) Ratio managers to non-managers (customer-facing)
Treatment Effect	1.262**	-0.375	0.358**	0.175	0.0855	-0.0184**
	(0.539)	(1.486)	(0.146)	(0.112)	(0.0718)	(0.00926)
Observations	9,007	9,007	9,007	9,007	6,193	9,007
R-squared	0.990	0.980	0.972	0.964	0.660	0.602
Treated Stores	87	87	87	87	87	87
Stores	220	220	220	220	220	220
Mean of Dv	24.82	77.84	5.496	2.387	0.423	0.0732
Relative Effect	5.1%	-.5%	6.5%	7.3%	20.2%	-25.1%

Notes: Table 7 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 5: Workers and Managers (backroom/customer-facing)

a) Workers in backroom

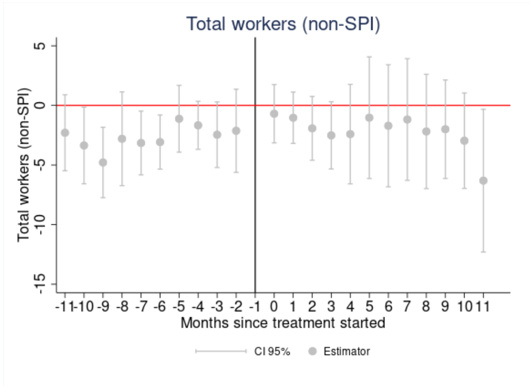

b) Workers in customer-facing

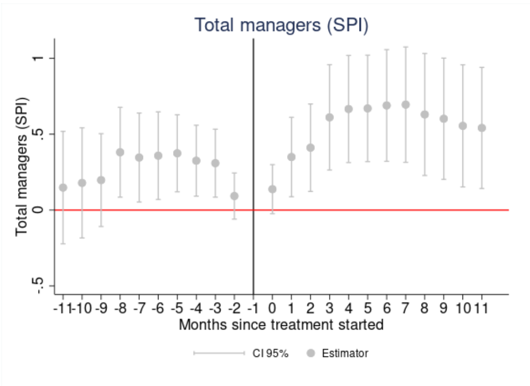

c) Managers in backroom

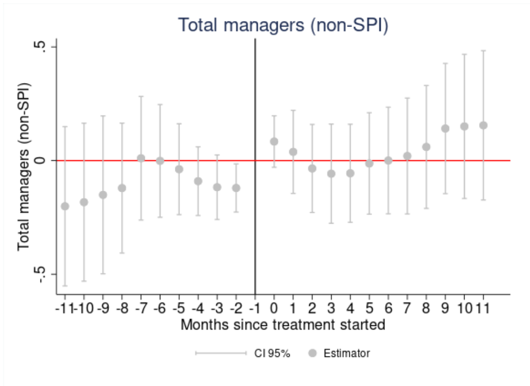

d) Managers in customer-facing

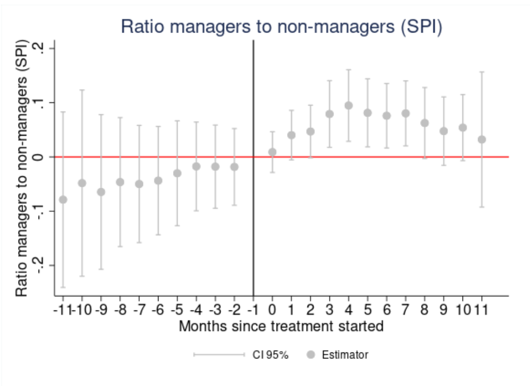

e) Mngrs/Non-Mngrs in backroom

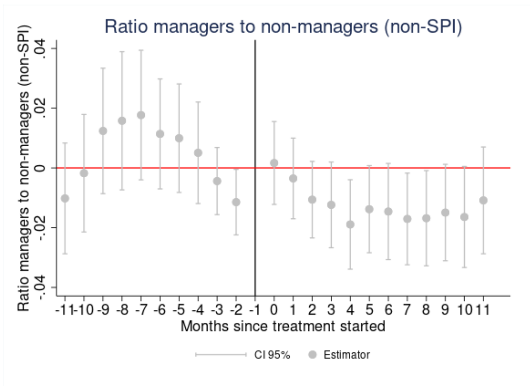

f) Mngrs/Non-Mngrs

Table 8: 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 8 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%.

5 Alternative Mechanisms

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 attention more effectively. We have shown that the arrival event causes three inventory- and pricing-related changes (specifically, (1) stockouts fall, (2) inventory becomes leaner, and (3) price changes result in higher sales)) along with a general restructuring of the organizational structure.

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 3 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

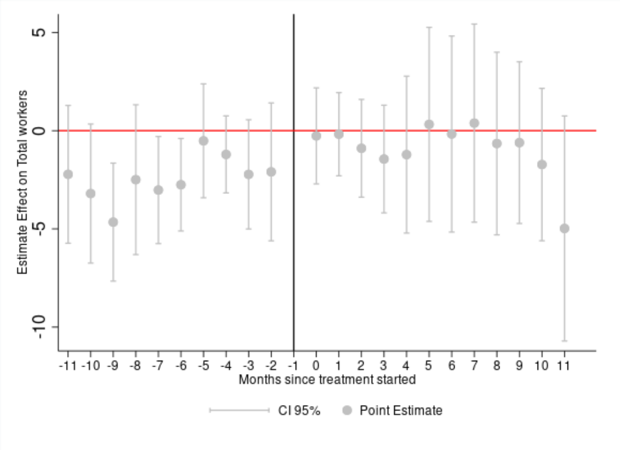

5.2 Demand for Additional Manager(s)

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 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.12 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).

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.

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 7: Heterogeneous effects

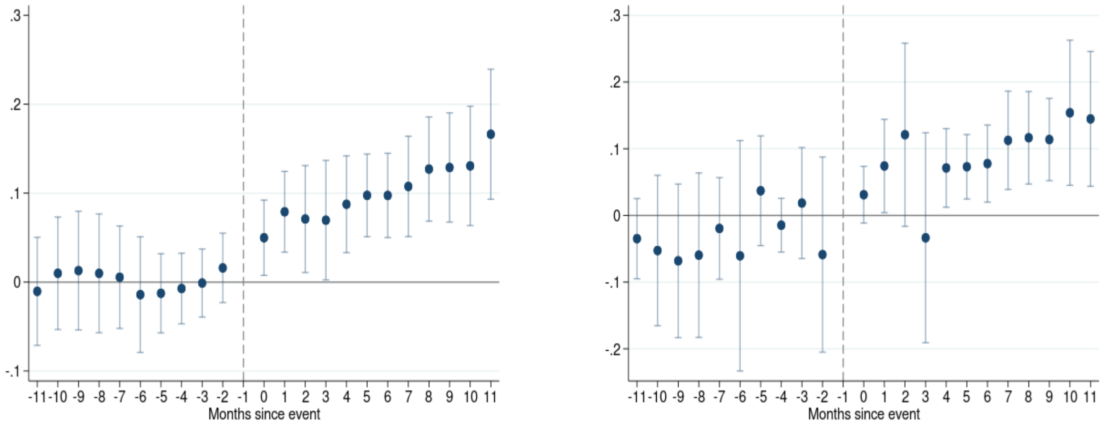

(a) Log sale Sun-Abraham

(b) Log sale Callaway-Sant'Anna

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., Howell, P., Nyshadham, A., and Tamayo, J. (2025). Rapport in Organizations: Evidence from Fast Food. *Working Paper*.

Adhvaryu, A., Kala, N., and Nyshadham, A. (2022). Management and Shocks to Worker Productivity. *Journal of Political Economy*, 130(1):1–47.

Adhvaryu, A., Nyshadham, A., and Tamayo, J. (2023). Managerial Quality and Productivity Dynamics. *Review of Economic Studies*, 90(4):1569–1607.

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

Alvarez, J., Benguria, F., Engbom, N., and Moser, C. (2018). Firms and the Decline in Earnings Inequality in Brazil. *American Economic Journal: Macroeconomics*, 10(1):149–189.

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.

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*.

Bandiera, O., Prat, A., Sadun, R., and Wulf, J. (2014). Span of Control and Span of Attention.

Bloom, N. and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The quarterly journal of Economics*, 122(4):1351–1408.

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

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.

Dessein, W. and Santos, T. (2021). Managerial Style and Attention. *American Economic Journal: Microeconomics*, 13(3):372–403.

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.

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

Iyengar, S. S. and Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology*, 79(6):995–1006.

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.

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.

Simon, H. A. (1973). Applying Information Technology to Organization Design. *Public Administration Review*, 33(3):268–278.

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., Deming, D. J., and Bhalotra, S. R. (2024). How Do You Find a Good Manager?

A Appendix

A.1 Additional Statistics

Table A.9: 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.9 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.8: Distribution of differences between arriving and existing management

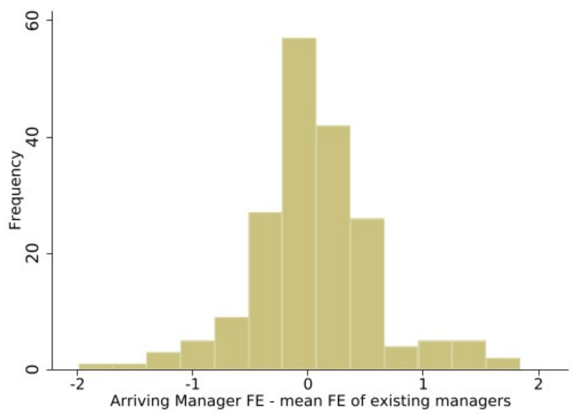

Notes: This figure shows the distribution of the differences between the fixed effect of the arriving manger and the mean of the fixed effects of all managers working at the store prior to the arrival of the new manager.

Table A.10: 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.10 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

Figure A.9: Effect of Treatment on Log Productivity

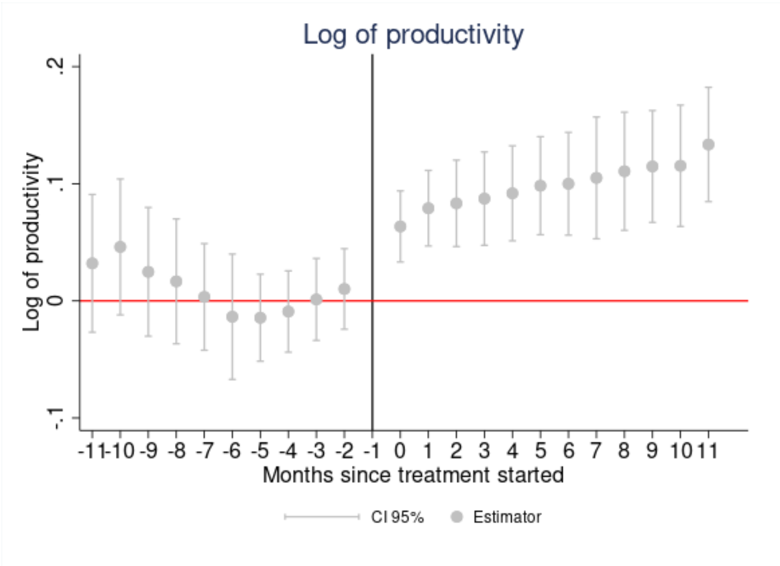

Notes: Figure A.9 shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's arrival on log productivity. The sample consists of store-monthly panel data for the 220 stores during our analysis period of January 2017 to June 2020. The vertical line represents the arrival of a good manager.

A.2.1 Additional Evidence on Manager Actions

Figure A.10: Number and Value of Stockouts

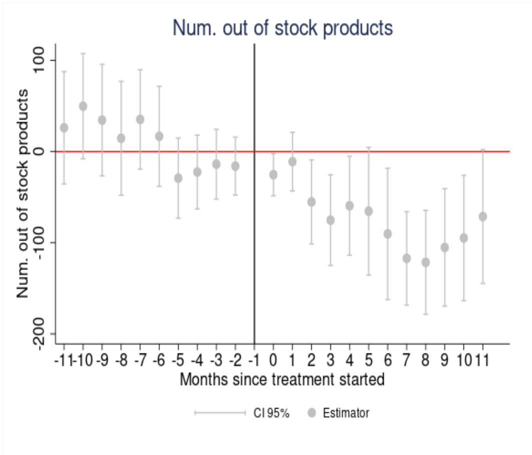

Stockouts

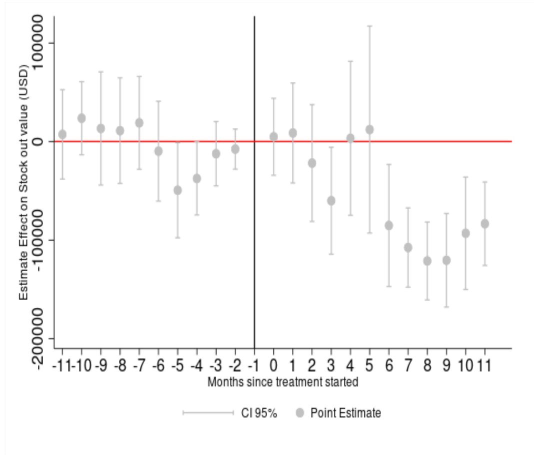

Value of stocked-out inventory (in thousands USD)

Figure A.11: Price Changes

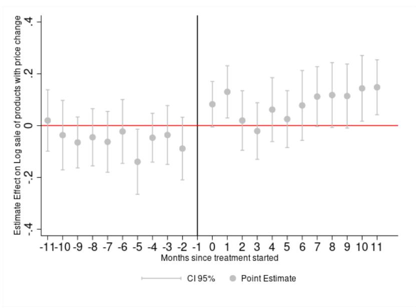

(a) Log sales of products *with* a price change

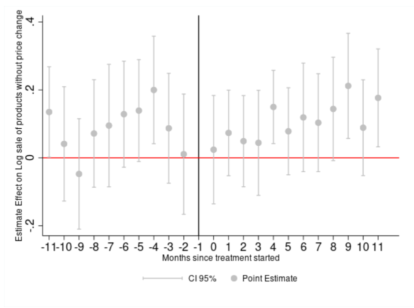

(b) Log sales of products *without* a price change

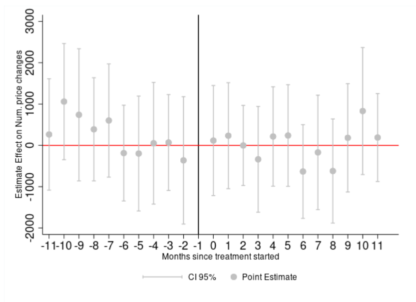

(c) Number of price changes

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's arrival on the log sales of products with and without price changes, as well as 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 vertical line represents the arrival of a good manager.

A.2.2 Additional Evidence on Organizational Changes

Figure A.12: Turnover, hiring, and moves in backroom and customer-facing departments)

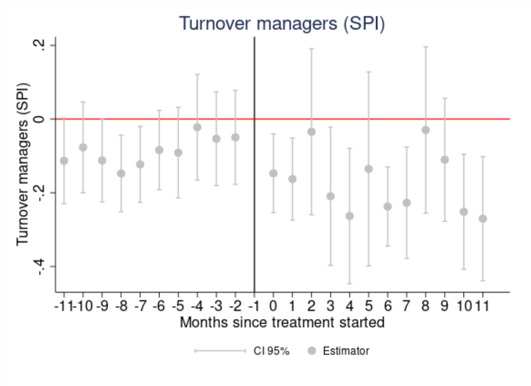

a) Turnover in backroom

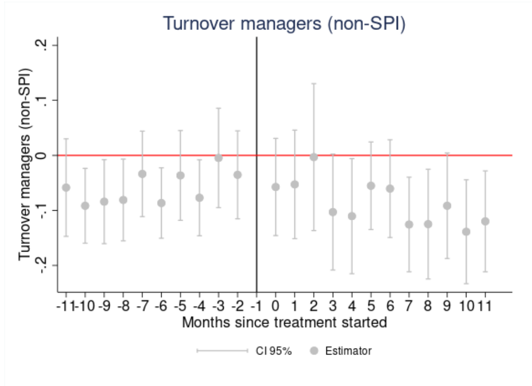

b) Turnover in customer-facing

c) Hiring in backroom

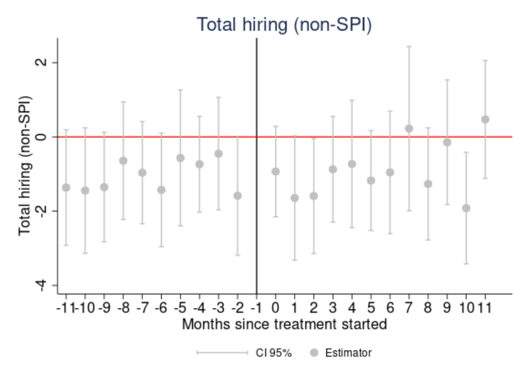

d) Hiring in customer-facing

A.3 Additional Evidence from the AKM Decomposition

Table A.11: 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.11 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.13: Event Study Around Movers - Log Productivity

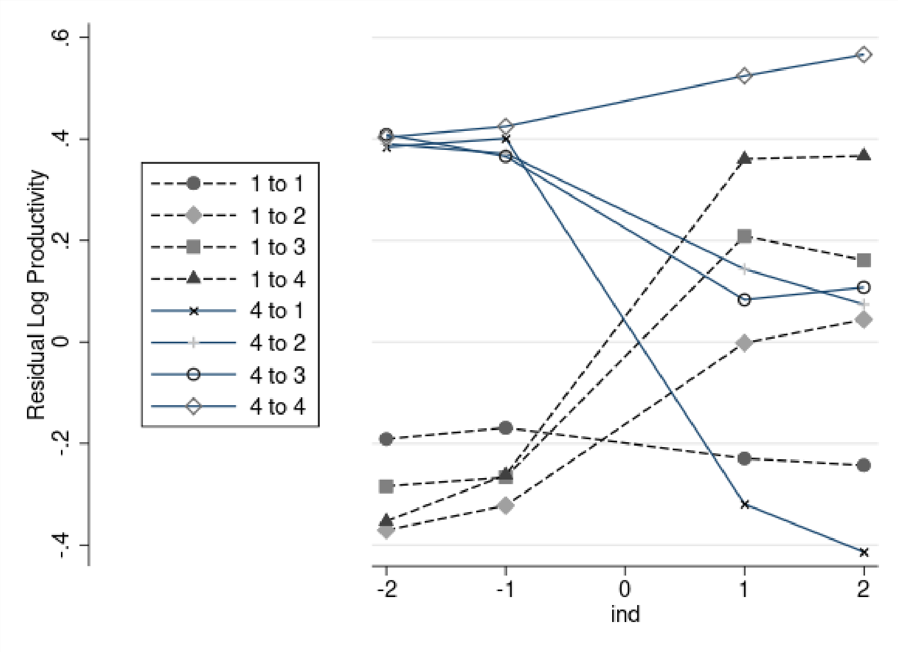

Notes: Figure A.13 ranks movers in terms of (i) quartiles of average log productivity in their initial store and (ii) quartiles of the average log productivity in the store where they moved to. The average log productivity is computed over the entire sample period, and quartiles are calculated for each 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.14: Symmetry Test - Log Productivity

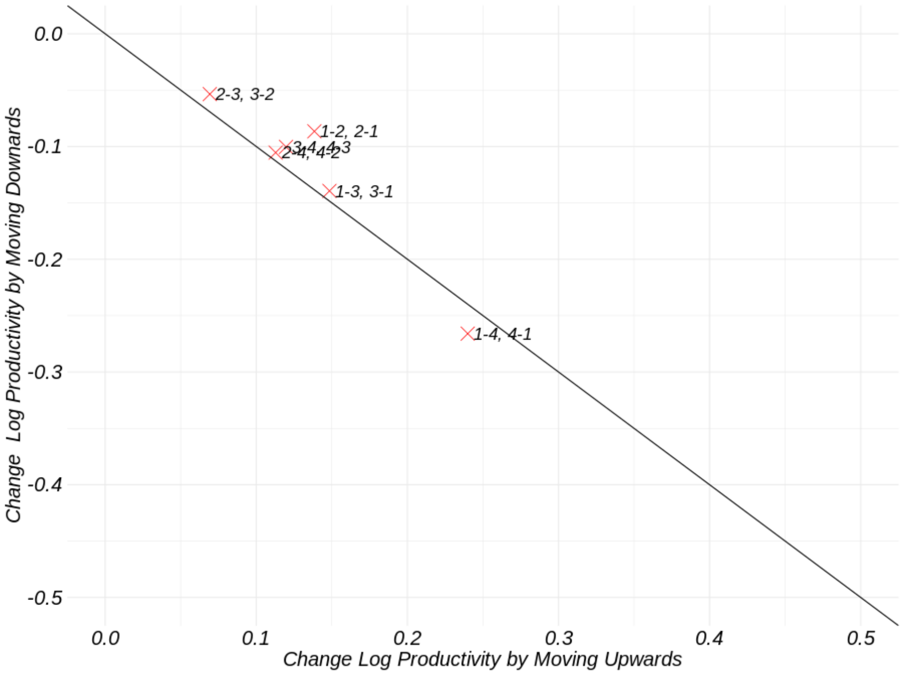

Notes: Figure A.14 ranks movers in terms of (i) quartiles of average log productivity in their initial store and (ii) quartiles of the average log productivity in the store where they moved to. Average log productivity is computed over the entire sample period, and quartiles are calculated for each store. 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.12: 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.12 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.13: AKM based on different approaches

	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.13 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 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.15: Effects of a good manager on log productivity: Sun-Abraham and Callaway-Sant'Anna specifications

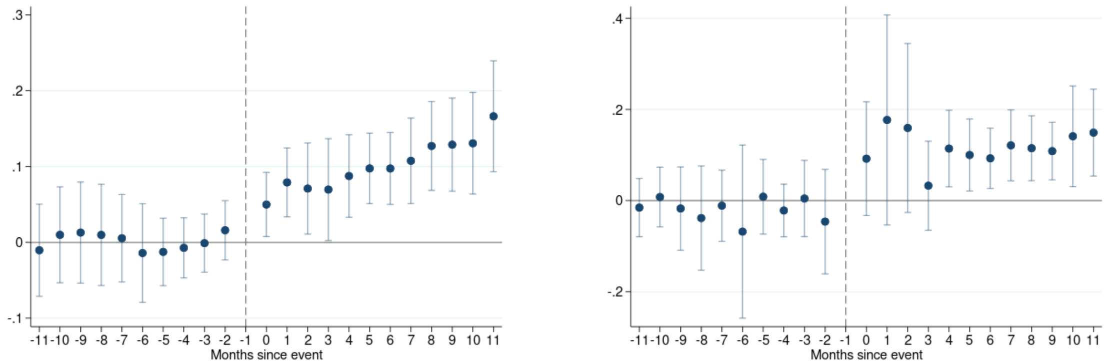

Sun-Abraham Log productivity

Callaway-Sant'Anna Log productivity

A.4.1 Robustness Results When Including 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 include 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.14: Workers and managers (backroom/customer-facing), excluding coordinators as movers

	(1) Total workers (SPI)	(2) Total workers (non-SPI)	(3) Total managers (SPI)	(4) Total managers (non-SPI)	(5) Ratio managers to non-managers (SPI)	(6) Ratio managers to non-managers (non-SPI)
Treatment Effect	1.778***	-1.351	0.507***	0.127	0.0401	-0.0307***
	(0.622)	(1.583)	(0.166)	(0.138)	(0.0346)	(0.0104)
Observations	8,163	8,163	8,163	8,163	5,461	8,163
R-squared	0.990	0.980	0.973	0.965	0.643	0.599
Treated Stores	87	87	87	87	87	87
Stores	220	220	220	220	220	220
Mean of Dv	24.32	76.25	5.549	2.846	0.351	0.0909
Effect on mean	0.0731	-0.0177	0.0914	0.0446	0.114	-0.338

Notes: Table A.14 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 A.16: Workers and Managers (backroom/customer-facing) Excluding Coordinators as Movers

a) Workers in backroom

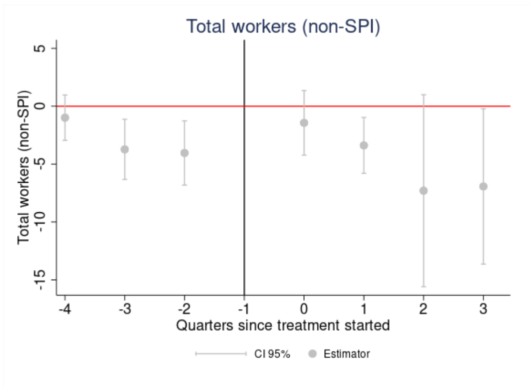

b) Workers in customer-facing

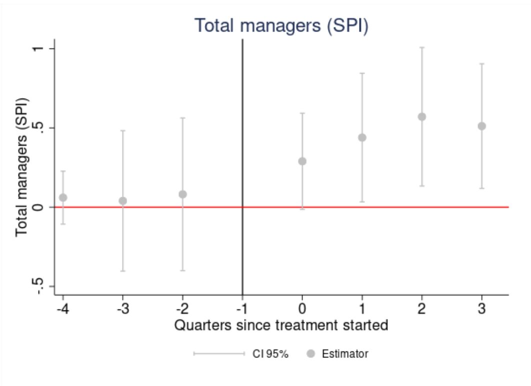

c) Managers in backroom

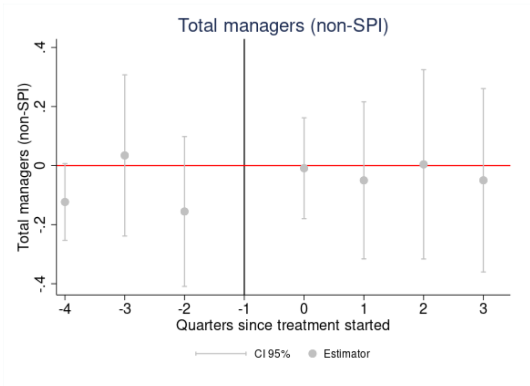

d) Managers in customer-facing

Notes: This figure shows the event-study coefficients and 95 percent confidence intervals from estimating a good manager's arrival on the number of total workers and/or managers in backroom and customer-facing departments. Workers counts all workers at the store, including managers. 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 arrival of a good manager.

Table A.15: Turnover, hiring, and moves in backroom and customer-facing departments

VARIABLES	(1) Turnover managers (SPI)	(2) Turnover managers (non-SPI)	(3) Total hiring (SPI)	(4) Total hiring (non-SPI)	(5) Moves managers to SPI	(6) Moves managers to non-SPI
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%.

Figure A.17: Turnover, hiring, and moves in backroom and customer-facing departments - Excluding Coordinators as Movers

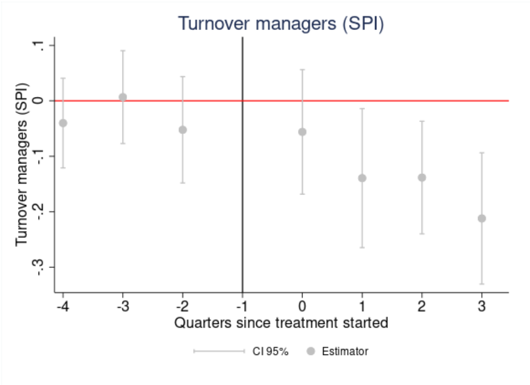

a) Turnover in backroom

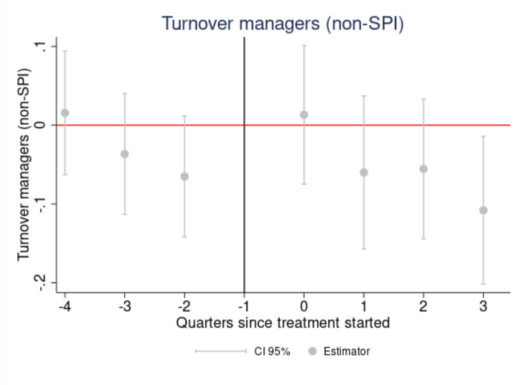

b) Turnover in customer-facing

c) Hiring in backroom

d) Hiring in customer-facing

e) Moves to backroom

f) Moves to customer-facing