Promotions and the Peter Principle*

Alan Benson
Danielle Li
Kelly Shue

July 29, 2019

The best worker is not always the best candidate for manager. In these cases, do firms promote the best potential manager or the best worker in her current job? Using microdata on the performance of sales workers at 131 firms, we find evidence consistent with the Peter Principle, which proposes that firms prioritize current job performance in promotion decisions at the expense of other observable characteristics that better predict managerial performance. We estimate that the costs of promoting workers with lower managerial potential are high, suggesting either that firms are making inefficient promotion decisions or that the benefits of promotion-based incentives are great enough to justify the costs of managerial mismatch. We find that firms manage the costs of the Peter Principle by placing less weight on sales performance in promotion decisions when managerial roles entail greater responsibility and when frontline workers are incentivized by strong pay for performance.

JEL Codes: M51, M52, J24, M12, J33, J31

Title page information requested in QJE Author Guidelines:

1. Publication correspondence and proofs can be sent to Kelly Shue, Yale School of Management, 165 Whitney Avenue, PO BOX 208200, New Haven, CT 06520-8200. kelly.shue@yale.edu.
2. Word count: 12,669 (main text)

*Respective contacts and affiliations are: bensona@umn.edu, University of Minnesota Carlson School of Management; danielle.li@mit.edu, MIT Sloan School of Management and NBER; and kelly.shue@yale.edu, Yale School of Management and NBER. We thank Peter Cappelli, Lisa Kahn, Steve Kaplan, Maximilian Kasy, Eddie Lazear, Inessa Liskovich (discussant), Vladimir Mukharlyamov (discussant), Paige Ouimet (discussant), Michaela Pagel (discussant), Amanda Pallais, Brian Phelan, Thomas Peeters, Lamar Pierce, Felipe Severino (discussant), Kathryn Shaw (discussant), Mike Waldman, and seminar participants at CEPR Milan, CUHK, FOM Conference, GSU CEAR, Gerzensee ESSFM, Harvard Business School, HKUST, Hong Kong U., ITAM, Melbourne FIRCG, MIT Economics, MIT Sloan, the MEA Meetings, Nanjing U., NBER (Corporate Finance, Personnel Economics, and Organizational Economics), Purdue, SOLE, Tsinghua PBC, U. of Chicago, U. of Hong Kong, U. of Miami, U. of Minnesota Carlson, Wharton People Analytics Conference, Wharton People & Organizations conference, and Yale Junior Finance Conference for helpful comments. We thank Menaka Hamplole, Leland Bybee, and Gen Li for excellent research assistance. This research was funded in part by the Initiative on Global Markets and the Fama Miller Center at the Chicago Booth School of Business and the International Center for Finance at Yale SOM.

I. Introduction

When management requires skills that are different from those required for lower-level work, the best workers may not make the best managers. In these cases, do firms promote someone who excels in her current position or someone who is likely to excel as a manager? If firms promote workers based on their current performance, they may end up with worse managers. Yet if firms promote workers based on traits that predict managerial performance, they may pass over higher-performing workers, thereby weakening incentives for workers to perform well in their current roles. Such promotion policies could also lead to perceptions of favoritism or unfairness, or the impression that effort in one's job goes unrewarded.

Using detailed microdata on sales workers in US firms, we provide the first large-scale empirical evidence of the Peter Principle, a hypothesis that firms prioritize current performance in promotion decisions at the expense of promoting the best potential managers (Peter and Hull 1969). In particular, we show that firms discriminate in favor of high-performing sales workers by promoting them ahead of lower performing sales workers with greater managerial potential. We then show that firms overweight sales in promotion decisions by constructing a counterfactual promotion policy that improves managerial quality by promoting fewer top salespeople.

These results suggest either that firms make mistakes in their promotion decisions or that the incentive benefits of promoting based on sales performance justify the costs of promoting workers with lower managerial potential. Consistent with the latter, we find that firms manage the costs of the Peter Principle by placing less emphasis on sales performance in settings in which salespeople are rewarded by strong pay for performance and where managerial roles entail greater responsibility.

The Peter Principle applies broadly to settings in which the skills required to succeed at one level in the organizational hierarchy may differ from the skills required in the next level, such as science, engineering, manufacturing, academia, or entrepreneurship (Baker, Jensen and Murphy 1988).[1] Among such settings, sales is particularly attractive from a research perspective. First, it

1. Kaplan, Klebanov and Sorensen (2012) and Kaplan and Sorensen (2016) show that execution, interpersonal, and general skills strongly predict executive performance, underscoring the possibility that promoting based on lower-level job skills rather than managerial skills can be costly. The Peter Principle may also be highly relevant for entrepreneurial firms, which must decide whether to retain founders in leadership roles (Ewens and Marx 2018;

is an economically important occupation, accounting for 9% of the US labor force.[2] Second, the sales setting offers a relatively clean and complete performance measure. Finally, it allows us to explore an interesting tension: sales is both widely cited as a canonical example of where the Peter Principle likely applies[3] and as a setting in which a simple economics model would predict that performance pay already incentivizes worker effort. Finding evidence of the Peter Principle in this setting suggests that the ability to observe and condition pay on performance cannot fully resolve the tension between providing incentives and promoting the most qualified managers.

Our analysis uses new transaction-level data that are well suited for the study of firms' promotion policies.[4] These data, provided by a firm that offers sales performance management software to client firms, include standardized measures of sales transactions and organizational hierarchy for a panel of 38,843 workers, 1,553 of whom are promoted into managerial positions during our sample period. Our data cover 131 different US-based client firms in a range of industries from 2005 to 2011, allowing us to study heterogeneity in how much firms prioritize current job performance as a function of firm organization or pay practices.

For sales workers, we use employment history and sales credit data to examine promotion as a function of sales performance (the dollar value of sales), sales collaboration (the number of colleagues with whom a worker shared credit on transactions), and other observable worker characteristics. For promoted managers, we evaluate managerial performance as their "manager value added" in shaping their subordinates' sales performance, that is, each manager's contribution to improving her subordinates' sales, controlling for subordinate and firm-year-month fixed effects as well as other potentially confounding factors (following the methods used in, for example, Abowd, Kramarz, and Margolis 1999; Adhvaryu et al. 2019; Bertrand and Schoar 2003; Lazear, Shaw and Stanton 2018).

In our setting, we define the Peter Principle as a promotion policy that 1) places positive weight

Hellmann and Puri 2002).

2. In 2018, the US labor force had 14.5 million workers in sales and sales-related occupations (Bureau of Labor Statistics 2018).

3. Deutsch (1986) points out that "American companies have always wrestled with ways to keep the Peter Principle at bay—to prevent competent salesmen, for example, from rising to become incompetent sales managers." Baker, Jensen and Murphy (1988) state that "in many cases, the best performer at one level in the hierarchy is not the best candidate for the job one level up—the best salesman is rarely the best manager."

4. We do not observe promotion offers. As such, when we refer to a "promotion policy," we refer to the combined impact of the firm's promotion offer and the worker's decision to accept the offer.

on worker sales performance and 2) places more weight on sales performance than a policy aimed solely at maximizing managerial quality. Our empirical analysis begins by testing the first part of this statement: we find a strong positive relation between past sales performance and promotion. To test the second part, we examine whether firms prioritize sales performance more than would be expected if they were simply trying to identify the best managers: that is, do they "discriminate" in favor of workers with strong sales performance by promoting them even if they have lower managerial potential?

We first show that pre-promotion sales performance is negatively correlated with post-promotion manager value added. The negative correlation is consistent with the Peter Principle: if promotion policies discriminate against workers with low sales, then low sales workers who are nevertheless promoted should be better managers. However, differences in *average* manager value added across promoted low and high sales workers need not be evidence that firms discriminate in favor of high sales workers—as long as firms equate expected manager value added on the *margin*. To test for discrimination, we conduct a Becker outcomes test to compare the managerial performance of marginally promoted high and low sales workers. Intuitively, if firms lower their standards for managerial potential to promote top sales workers, marginally promoted high sales workers will be worse managers than marginally promoted low sales workers (Becker 1957, 1993).

Using a variant of the models of Angrist, Imbens and Rubin (1996) and Abadie (2003), we identify marginally promoted workers by instrumenting for each worker's promotion using the average firm-level promotion rate in each month, leaving out the focal worker and her teammates. The compliers to this instrument can be thought of as marginal because they are promoted only if overall promotion rates are high, but would not have been promoted had average promotion rates been slightly lower. Our approach is analogous to that of Arnold, Dobbie and Yang (2018), who study discrimination in bail decisions.

We show that the instrument strongly predicts individual promotion. To satisfy the exclusion restriction, the instrument must also be uncorrelated with managerial potential. One may be concerned that high average promotion rates may reflect strong consumer demand or other time-varying firm shocks that affect the performance of all sales workers and may thus be

correlated with managerial quality. In our setting, however, we measure a manager's quality as her value added to subordinate sales, net of firm-year-month fixed effects. As a result, our measure of manager quality is, by construction, orthogonal to any firm-level time varying conditions that may also affect average promotion rates. Further, using a leave-out mean removes the direct impact that an individual's own promotion status or contribution to team performance can have on her value of the instrument.

We use this instrument to compare marginally promoted low and high sales workers (that is, the quality of low versus high sales workers who are compliers to our instrument for promotion). Across a variety of specifications, the managerial quality of marginally promoted workers is declining in their pre-promotion sales performance, providing evidence that firms apply lower standards when evaluating top sales performers for promotions. We then show that firms can improve managerial quality by promoting fewer top sales performers on the margin.

Our analysis also identifies another observable worker characteristic, sales collaboration experience, which is positively related to managerial performance but not consistently correlated with promotion. Sales collaboration may be a measure of a worker's experience working in teams or with more complex products that require coordination. We cannot pinpoint the exact channel through which collaboration predicts managerial performance, but these results suggest that firms wishing *only* to maximize managerial quality could potentially achieve better outcomes by placing less weight on sales and more on collaboration experience in promotion decisions.

We also provide evidence that our results are not driven by potential issues arising from mean reversion as in Lazear (2004), nonrandom assignment of managers to subordinates, or the unwillingness of some top sales workers to accept promotion offers. We further test whether managers with high pre-promotion sales contribute to the firm in other ways, such as by engaging directly in sales (which may substitute for subordinate sales) or by retaining more skilled subordinates. We find no evidence that these managers are better along these dimensions.

To assess the magnitude of the costs associated with firms' existing promotion policies, we compare the managerial performance of promoted workers with the predicted managerial performance of workers who would have been promoted under a counterfactual promotion policy

that maximizes expected managerial quality. We find that average managerial quality, as measured by value added to subordinate sales, is 30% higher under this counterfactual policy. These findings do not necessarily imply that firms are making mistakes. Rather, they suggest that the costs of not promoting the best potential managers may be high: firms value the incentive benefits of promoting based on demonstrated job performance enough to sacrifice managerial quality by up to 30%.

Lastly, we explore how firms trade off the benefits of using promotion-based incentives against the costs of managerial mismatch. We examine how promotion policies vary with managerial responsibility and the power of incentives. We expect to find less evidence of the Peter Principle in settings in which managerial quality is more important or where the firm offers strong non-promotion-based incentives for worker effort. We find that firms where managers supervise large teams place less weight on sales performance and more on collaboration experience when promoting workers. We also find that firms with stronger pay for performance put less weight on sales performance when making promotion decisions. However, we do not find that pay for performance can eliminate the costs associated with the Peter Principle. Indeed, relative to other occupations, sales is associated with high pay for performance; yet we continue to find evidence consistent with the Peter Principle, suggesting that the incentive power of promotions may be quite important in practice.

This paper is organized as follows. Section II presents our definition of the Peter Principle in the context of the related literature. Section III introduces our setting and data. Section IV provides baseline evidence consistent with the Peter Principle. Section V develops our empirical framework and provides our main results. Section VI discusses alternative explanations. Section VII explores the trade-offs associated with promoting based on current performance. Section VIII concludes. A supplementary model in which firms may optimally bear the costs of the Peter Principle, proofs pertaining to our empirical strategy, descriptive statistics, and additional results are available in the Online Appendix.

II. The Peter Principle and Related Literature

Peter and Hull (1969) first introduced the Peter Principle as a satirical commentary on the seemingly dysfunctional reasons why people are promoted. The book's introduction defines the Peter Principle as the idea that "in a hierarchy, every employee tends to rise to his level of incompetence," but the remainder of the book treats this principle as the outcome of organizations' tendency to promote workers who excel at their current jobs while downplaying or ignoring their aptitude for management. The idea that organizations promote based on current performance at the expense of maximizing the match quality between a worker's skills and the new position has come to define the Peter Principle in both the popular press and the academic literature that followed Peter and Hull's original work. For instance, Fairburn and Malcomson (2001:46), argue that "distortion [in assignments] takes the form of promoting employees who would not be promoted for assignment reasons alone, the Peter Principle effect." Similarly, Faria (2000:4) defines the Peter Principle as, "Some firms try to avoid rent-seeking workers... by imposing simple rules of promotions, based on... past performance. One shortcoming... is that people can be placed in important jobs for which they are ill qualified."

Both Peter and Hull (1969) and the economic literature argue that suboptimal matching to managerial positions may be the price that organizations pay to incentivize worker effort. Peter and Hull argue that promoting a productive worker "serves as a carrot-on-a-stick to many other employees," (Peter and Hull 1969: 25-26). Milgrom and Roberts (1992: 364) write, "Promotions serve two roles in an organization. First, they assign people to the roles where they can best contribute to the organization's performance. Second, promotions serve as incentives and rewards." Similarly, Baker, Jensen, and Murphy (1988: 599) argue that "promotions are a way to match individuals to the jobs they are best suited... A second role of promotions is to provide incentives for lower level employees who value the pay and prestige associated with a higher rank in the organization."[5]

Building on the previous literature, we define the Peter Principle for the purposes of this paper

5. The trade-off between incentives and matching has also been incorporated into the theoretical literature. For instance, models of internal careers yield the prediction that the incentive purpose of promotions may lead firms to promote insiders over more qualified outsiders (e.g., Ke, Li and Powell 2018; Malcomson 1984; Waldman 2003).

as follows: *Firms promote workers who excel in their current roles, at the expense of promoting those who would make the best managers.* Note that this definition does not imply that firms make mistakes. Rather, evidence of the Peter Principle implies that firms face a costly trade-off between promoting the best potential managers and incentivizing workers.

The existing literature has pointed to at least four reasons why firms may optimally choose to use promotion-based incentives (in addition to other forms of compensation), despite the potential downside of lowering managerial match quality. First, workers may value managerial titles associated with promotion because titles confer status and can be readily advertised on resumes (DellaVigna and Pope 2016; DeVaro and Waldman 2012; Waldman 1984*a*,*b*, 2003). Second, promotion-based incentives reduce the potential negative spillovers associated with wide horizontal pay inequality. Cullen and Perez-Truglia (2018) present empirical evidence that horizontal pay inequity can demotivate worker effort, while vertical pay inequality (as would be associated with promotion-based incentives) can motivate effort. Along the same lines, Larkin, Pierce and Gino (2012) argue that strong performance pay poses psychological costs that spill over into the rest of the organization. Third, firms may commit to promoting on objective performance measures to avoid perceptions of inconsistency, influence activities (Milgrom 1988), and favoritism (Fisman et al. 2017; Prendergast and Topel 1996) that could make cash compensation costly compared with promotions. Fairburn and Malcomson (2001) offer a specific theory for the Peter Principle by which firms require senior managers to promote productive workers because cash rewards are more susceptible to influence activities. Lastly, promotion policies based on verifiable performance metrics such as sales may discourage the manipulation of other, more fungible performance metrics such as credit sharing and collaboration experience (DeVaro and Gürtler 2015; Fisman and Wang 2017).[6]

A number of theoretical papers relate to the Peter Principle, but the empirical evidence is

6. DeVaro and Gürtler (2015) develop a model in which workers strategically allocate effort among multiple tasks to be assigned to their preferred jobs. In our setting, workers can potentially choose the allocation of effort to individual sales or more collaborative activities. Concerns regarding strategic gaming may explain why firms do not promote based on sales collaboration even though it predicts managerial performance: collaboration experience can be gamed by strategically sharing and trading credits, while the revenues associated with sales are relatively difficult to game and more directly aligned with firms' objectives. Indeed, if firms began to heavily weight collaboration experience in promotion decisions, workers could potentially add fake collaborators by sharing credits. Recent examples of the gaming of various sales evaluation metrics include Benson (2015), Larkin (2014), and Oyer (1998).

much more limited. This study offers the first empirical test of the Peter Principle using data on promotions across a large number of firms. Our paper is most closely related to Grabner and Moers (2013), which uses detailed promotions data from a single bank. They show that the bank places less weight on current job performance when a promotion would be to a job performing dissimilar tasks, illustrating how the bank attempts to mitigate the costs associated with the Peter Principle. This paper differs in that we use data from a large sample of firms, and our goal is to estimate the overall cost of the Peter Principle, thereby characterizing the importance of the trade-off that firms face when deciding on promotion policies.

Finally, our analysis is motivated by research showing that managerial quality is an important determinant of firm productivity (e.g., Bloom and Van Reenen 2007). A large related literature on corruption in leadership dating back to Weber (1947) attributes the existence of bad leaders to selection policies that are polluted by nepotism and cronyism. Our findings show that promotion policies that are more meritocratic or "fair" may still be problematic, as promoting based on merits in the current job—rather than on managerial potential—may still result in bad leaders.

III. Setting and Data

Our data come from a firm that offers sales performance management (SPM) software over the cloud. The firm's clients input their employee records, organizational hierarchies, and sales transactions into the software, which then calculates pay for each individual worker. Transaction inputs can be entered manually or linked to order management and customer relationship management (CRM) software. Pay outputs are typically linked directly to payroll software. The software also provides reporting and analysis. Sales workers and sales managers can view their sales credits, progress toward quotas, commissions, and other data. The software can also generate reports for use in auditing and compliance with Sarbanes-Oxley.

The data include 131 client firms and 38,843 sales workers, 1,553 of whom are promoted to managerial roles. The most-represented industries are information technology and services (57 firms), manufacturing (30 firms), and professional services (21 firms). Table I provides descriptive statistics. All firms have at least one complete fiscal year of data, and no firm constitutes more than

14% of employee observations. The Online Appendix provides further details, including industry coverage (Figure A1).

III.A. Overview of Sales Positions

Sales workers are typically assigned a market consisting of a territory, a set of products, or a type of client. Within their market, they are responsible for generating leads on potential new clients, making first contact, executing the initial sale, cross-selling other products, selling upgrades, and maintaining relationships.

The primary measure of a salesperson's performance is the total dollar value of the sales to which he contributes. Our data include 156 million sales transactions tied to individual workers. Table I describes the distribution of monthly sales. Because sales tend to be intermittent, we report rolling averages of sales credits in the previous 12 months. The quartiles for monthly worker sales are \$39,395, \$294,928, and \$1.68 million (in 2010 dollars). Reflecting the wide and skewed distribution of sales across markets in which workers operate, the mean of this figure is \$3.26 million.

Figure I also illustrates the skewness in the distribution of sales. The top left panel presents a histogram for the raw distribution of worker-level monthly sales (measured as 12 month rolling averages). The middle left panel plots the log of monthly sales, which follows a less skewed distribution. The bottom left panel, which reflects our main measure of sales performance, shows the residual distribution of monthly sales after controlling for firm-year-month fixed effects. In other words, we measure sales performance as the recent performance of a sales worker compared with others in her same firm at the same period in time. Even with these fixed effects, we still observe wide variation in sales across workers. The interquartile range of residual log sales is 2.44, meaning that, among rank-and-file sales workers in the same firm in the same year-month, a worker in the 75th percentile generates approximately $e^{2.44} = 11.5$ times as much revenue as one in the 25th percentile. Although this difference is stark, it's also consistent with the so-called 80-20 rule, a well-known heuristic in the sales industry that states that the top 20% of the sales force is responsible for 80% of sales.

In addition to total sales, we also observe collaboration experience, which we explore in

Section V.D. For complex products and services, a single transaction can involve salespeople across many sales functions, products, and territories. In our data, we observe all workers credited on a transaction and define a salesperson's collaboration experience as her average number of distinct collaborators per order over the past 12 months (or for her tenure if less than a year).

Table I presents summary statistics for collaboration, and Figure I presents histograms of the distribution of collaboration experience. Over 40% of workers worked alone in the past year, while the remainder vary greatly in their number of collaborators. This difference does not merely reflect differences in work organization across firms or over time. Panel D of Figure I shows that even within the same firm-year-month, there is substantial variation in the extent to which workers collaborate on sales (Online Appendix Figure A2 shows the distribution of team sizes within and across firms). The within firm-year-month interquartile range of sales collaborators is 0.71, signifying that the 75th percentile worker has $e^{0.71} = 2.03$ times as many collaborators as the 25th percentile worker.

This variation in collaboration highlights two archetypal sales workers described in the practitioner literature. "Lone wolves" are known for their self-confidence, resilience, and autonomy and are stereotypically marked by their reluctance to share leads, best practices, and client relationship responsibilities with others in the organization. The most effective team players, by contrast, enable those around them by forwarding leads, crafting sales that include many others' territories and products, forwarding established clients to account managers, and developing team members so they can be effective in these capacities. These lead generation and origination activities would also generally entitle that salesperson to a portion of the sales executed by others.[7]

The correlation between our sales and collaboration measures is 0.19. The moderate correlation shows that there is substantial variation across these measures.

Table I also provides summary statistics for worker compensation. Because our data provider's software is designed to track and distribute pay for sales performance, salary is an optional field and can be missing or measured with error. Based on these limited data, we believe that the median

7. We do not assume that collaboration experience is freely chosen by the worker. Indeed, some workers may be assigned to work alone or in teams. We instead focus on showing that collaboration experience, which is observable by the firm, positively predicts manager value added.

worker in our sample receives at most $89,000 in base pay per year, and more likely $50,000 to $60,000 per year in base pay, which is approximately half that of managers. Given that the software outputs commission data that are often linked to payroll, we're more confident in these measures, although they can still be missing. The median sales worker earns $3,842 per month in commission pay, slightly less than our estimates of workers' median base pay, and the 75th percentile sales worker earns more in commission pay than in base pay. Our sample is largely composed of sales workers who engage in big-ticket business-to-business sales whose pay is substantially greater than Bureau of Labor Statistics estimates for sales workers ($49,430 to $70,200 per year in the middle of our sample). However, the pay mix is similar to benchmark data for skilled sales jobs involving high degrees of autonomy.

Our analysis uses monthly sales as the measure of pre-promotion sales performance, which has the advantage of being highly standardized, and after controlling for firm-year-month fixed effects, has an easy interpretation. A limitation of our sales performance measure is that we do not observe the profit margins associated with sales transactions. Nevertheless, we believe that the relative levels of sales credits among workers in the same firm and time offer a reasonable approximation of relative sales performance. In theory, we could instead use worker compensation as a measure of sales performance. However, this approach would also have disadvantages. First, some firms in our sample use the software to track sales performance but not to record compensation, so we lack compensation data for the full sample. Second, compensation doesn't always correspond to recent performance; for example, in a given month, workers may receive commissions for origination or renewals for sales made in the distant past. Third, the base pay data can be unreliable since they are not required by the software and are not directly linked to payroll. Therefore, we prefer relative sales credits as our measure of sales performance.

Our data have the unique advantage of offering detailed organizational structure and worker productivity measures, but unfortunately we do not observe employee demographic characteristics such as age, gender, or education. We do observe worker tenure, which may affect both worker sales and promotion prospects. The tenure variable is censored by the date the firm began using the SPM software. Therefore, we control for tenure within the SPM system and its interaction with

whether tenure is potentially censored.

III.B. Overview of Managerial Positions

We observe the hierarchical structure linking sales managers to sales subordinates. For each person in the data, we observe the ID number of at most one direct superior within the hierarchy, as well as the ID numbers of any direct subordinates. Therefore, we define a worker as someone with zero subordinates and a manager as someone with at least one subordinate.

Managers typically have titles such as "territory manager," "sales director," "regional director," "regional manager," and "regional vice president." The last panel of Table I summarizes the characteristics of managers in our data. On average, each manager has five subordinates. Conversations with our data provider suggest that managers typically receive greater total compensation than their subordinates and have a pay mix that favors base pay rather than commission pay. Consistent with this, managers in our data have significantly higher reported salaries than workers on average and at each quartile of the pay distribution. In absolute terms, managers also have greater commissions than workers at each quartile of the commission pay distribution, though managers' overall pay mix is more weighted toward base pay. In addition, nonpecuniary rewards are also likely to favor managers, who typically enjoy greater prestige, opportunities for career progression inside and outside the firm, benefits, job security, pay security, and better work conditions than their subordinates.

Managers perform substantially different tasks. As summarized on O*NET, sales workers are primarily engaged in direct sales activities, whereas sales managers are responsible for building a high-performing sales team and earn commissions as a function of their team's performance (see Online Appendix Table A1). A survey of frontline sales managers by the Sales Management Association (2008) reports that sales managers spend the most time on performance management, followed by company administration, sales planning, selling and market development, and staff deployment. Performance management requires leadership, coaching, and training skills that may be imperfectly related to those used in direct sales activities. Administrative duties require general management knowledge so that the sales manager can interface with other functions, such as

marketing and operations. Sales planning requires data analysis skills so that managers can read market research, set quotas, assign territories, monitor performance, and prioritize sales activities. Sales managers also oversee the development of playbooks that compile best practices and outline the company's strategy for selling their products. Successfully executing these activities reflects in the performance of their teams. For example, if the manager misreads market research, sales workers could be misallocated to unproductive products or territories, quotas could be set at unattainably demotivating thresholds, or training could encourage salespeople to emphasize the wrong product features for their market.

Measuring Manager Quality

Because sales managers are ultimately responsible for improving the performance of their subordinates, we measure managerial performance as the impact of the manager on the sales of her subordinates. In general, any measure of managerial performance that relies on subordinate performance may be biased by the nonrandom assignment of managers to subordinates. For example, if a manager is assigned to high performing subordinates, the high sales numbers for these subordinates should not be attributed to the manager's skill.

To address these concerns, we follow Lazear, Shaw and Stanton (2018), Hoffman and Tadelis (2018), and a large literature on employer-employee and teacher-student matched data (e.g., Abowd et al. 2001) by estimating the manager's *value added* to their subordinates. We do so using a regression of the form

$$\text{Sales}_{imft} = a + \delta_i + \delta_m + \delta_{f\times t} + X_{it} + \varepsilon_{imft} \tag{1}$$

where the dependent variable is the log of one plus the sales performance of worker i under manager m at firm f and year-month t; the δ terms include fixed effects, and X_{it} includes seven bins for worker tenure, each interacted with an indicator for whether tenure is potentially censored.[8] The coefficients of interest are the manager fixed effects, δ_m, which is the average, time invariant

8. We estimate this regression using the Stata package `felsdvreg`. Rather than estimating $\delta_{f\times t}$ directly, we demean the outcome variable by firm-year-month prior to estimation to reduce computational demands.

component of a manager's quality or value added.

By including both manager and worker fixed effects, manager value added is identified from workers whom we observe under multiple managers. A manager's fixed effect represents the average change in sales performance across all workers who switch to or from that manager. As such, a manager with a high value added is one under whom workers perform above their individual mean across all the managers under whom they have worked. Whether a manager is assigned to strong or weak subordinates should not affect our measure of value added because a manager is credited only for *changes* in the performance of her subordinates. Further, firm-year-month fixed effects net out macroeconomic, industry-specific, and other firm-time specific conditions that may affect subordinate sales performance. Tenure effects net out returns to experience.

Estimating managerial quality as the manager's value added has clear advantages, as described above. However, we also acknowledge that the measure is imperfect. First, our estimates of manager value added are likely to be noisy. In equation (1), the dependent variable is worker monthly sales, which varies widely. Classical measurement error in worker sales will add noise to our measures of manager value added, raising our model's standard errors and increasing our estimates of the *variance* of the manager fixed effects. Second, our estimates of manager value added may be systematically biased if managers are nonrandomly assigned to subordinates on the basis of time-varying worker or manager characteristics, or potential match quality. Systematic bias may pose a problem if it is correlated with managers' pre-promotion sales, a possibility we discuss in detail in Section VI.

Summary Statistics: Manager Quality

We observe 5,956 managers in our data, of whom we are able to estimate fixed effects for 4,887. This lowered number comes from the high bar required to identify manager fixed effects: we must observe that manager supervising multiple subordinates whose own fixed effects are known through their work under other managers. Our sample is also constrained to managers within groups of workers and managers who are connected through moves. For instance, a connected group might contain a manager, her new subordinates, the previous managers of those subordinates, and the

other subordinates of those managers. Fixed effects for managers within the same connected group are comparable relative to a group-specific normalization. For the average firm in our sample, 76.5% of workers are part of this largest connected group. To make these fixed effects more comparable across firms, we further demean them by firm specific averages. Because we estimate manager fixed effects with varying precision, we weight summary statistics and regressions involving these fixed effects by the inverse variance of our estimates. Finally, to estimate the relation between pre-promotion characteristics and post-promotion managerial performance, we must further restrict the sample to observed promotions. We have information on both manager value added and pre-promotion characteristics for 1,054 managers who are promoted during our sample period.

By construction, manager value added has a mean of zero. The 25th percentile of this distribution is -0.71, implying that, when assigned to a 25th percentile manager, a worker's output is $e^{-0.71} = 0.49$ of what it would have been under the median manager. Conversely, when assigned to a 75th percentile manager, a worker's output increases by a factor of $e^{0.85} = 2.34$. Note that this interquartile range may be large because it reflects real differences in managerial performance or because of noise in the estimation of manager fixed effects, which exaggerates the variance.[9]

IV. Baseline Evidence of the Peter Principle

IV.A. Are Better Sales Workers More Likely to be Promoted?

Our first empirical exercise examines how the sales performance of frontline sales workers predicts their promotion to management:

$$\text{Promote}_{ift} = a_1 \text{Sales}_{ift} + X_{ift} + \delta_{f \times t} + \varepsilon_{ift}. \tag{2}$$

We estimate an OLS model for equation (2) on a worker-year-month level panel for worker i at firm f who has not yet been promoted as of year-month t in which at least one worker at the firm

9. See the Online Appendix for details regarding the manager sample and summary statistics (Table A2); the distribution of manager value added (Figure A3), and robustness checks where observations are not weighted by the inverse variance of the fixed effect estimates (Table A3 and Figure A4).

is promoted. The dependent variable, Promote_{ift}, is an indicator for whether a worker is promoted in the next month. Sales_{ift} is the log of one plus worker i's monthly sales credits, averaged over the past 12 months or over the worker's total tenure if it spans fewer than 12 months. The other covariates X_{ift} include: the log of one plus worker i's average number of collaborators per order, again averaged over the past 12 months or over the total tenure if it spans fewer than 12 months; an indicator for having no collaborations (whom we label "lone wolves"); and fixed effects for seven bins of worker tenure, interacted with whether tenure may be censored in the data. Some specifications also control for the firm-wide average promotion rate in the current month, leaving out the focal worker and her colleagues, or firm-year-month fixed effects.

Equation (2) estimates the determinants of firm "promotion policies," which we use as an umbrella term for the ultimate outcome in terms of which workers transition into managerial positions. We caution that firm "promotion policies" refer to more than the firm's choice of which workers to receive promotion opportunities. It also depends on the terms of the promotion offer and whether workers accept. We present a detailed discussion of non-random selection into the sample of promoted workers in Section VI.

Panel A of Figure II and Table II report our results. We find that firms are significantly more likely to promote higher performing salespeople. Accounting for firm-year-month fixed effects, the estimate in column (2) of Table II implies that a doubling of a worker's relative sales performance corresponds to a 0.074 percentage point increase in a worker's probability of being promoted, or a 32% increase relative to the base rate.[10] We also note that a doubling of a worker's relative sales performance is not an unusual occurrence in our data given the wide dispersion in worker sales—it is equivalent to a worker moving from the 50th to the 67th percentile in terms of relative worker sales.

Because promotions can be considered a tournament, columns (3) and (4) explore the role of each worker's relative sales ranking on promotions. We rank workers by sales within each team (e.g., those who share a common manager) in a firm-year-month (rank 1 implies the top salesperson). We then take the average of the ranks over the past 12 months. Column (3) shows

10. This follows from ln(2)*0.107=0.074. The base monthly rate of promotion is 0.23%.

that, controlling for a worker's actual sales output, her rank still matters: a decrease in ranking is associated with a substantially reduced probability of promotion. Column (4) shows that this is driven primarily by whether a worker is the top-ranked person within a sales team. Controlling for sales relative to the firm and team, being top ranked in sales (as measured by a 12-month rolling average) increases a worker's probability of promotion by 0.659 percentage points, corresponding to an approximate tripling of the base rate probability of promotion. Results are robust to a probit model for promotions (see Online Appendix Table A4).

The estimates presented in column (1) of Table II also show that the leave-out firm-year-month average promotion rate is highly predictive of an individual worker's promotion probability. We will use this result in later analysis when we instrument for a worker's probability of promotion.

IV.B. Do Better Sales Workers Make Better Managers?

Next, we examine the relation between pre-promotion worker sales performance and post-promotion manager value added:

$$\text{Manager Value Added}_{if} = b_1 \text{Pre-Promotion Sales}_{if} + X_{if} + u_{if}. \tag{3}$$

We estimate equation (3) at the manager level because manager value added is defined as a time-invariant manager characteristic. Pre-Promotion Sales_{if} is the log of one plus manager i's monthly sales credits as a worker, averaged over the 12 months prior to i's promotion or over the total tenure if it spans fewer than 12 months. Here, Pre-Promotion Sales_{if} is demeaned by the average sales performance of all workers in the sample in the same firm-year-month to account for variation in market conditions prior to a manager's promotion. Thus, Pre-Promotion Sales_{if} represents each manager's pre-promotion sales performance relative to other workers in the firm during the same time period. In some specifications, we control for a manager's pre-promotion collaboration experience, also defined relative to other workers in the firm during the same time period, an indicator for whether a manager was a lone wolf prior to promotion, and fixed effects for a manager's tenure in the month prior to promotion.

Panel B of Figure II and Table III show that there is a significant *negative* relation between pre-promotion sales performance and subsequent managerial performance. Column 2 of Table III shows, for instance, that doubling a manager's pre-promotion sales corresponds to a 0.061 point decline in manager value added. Since manager value added represents the change in log subordinate sales, this implies that a manager with double the pre-promotion sales leads each subordinate's sales to decline by 6.1%. Given that a typical manager is in charge of five subordinates, our results also imply that a doubling of a manager's pre-promotion sales predicts that total team sales under the new manager will decline by almost one-third of one worker. This result is, if anything, slightly stronger for managers who are assigned to manage a different team than the one he or she was originally on (e.g., managers whose new subordinates were not their prior teammates), indicating that our results are unlikely to be driven by team-specific factors such as group-level mean reversion. See additional discussion in Section VI.[11]

It may seem counterintuitive that good sales workers make *worse* managers because both roles are likely to require social skills, but the business press offers some insights into why excellence in sales may translate negatively into managerial quality. Sevy (2016), in a *Forbes* article entitled "Why Great Sales People Make Terrible Sales Managers," argues that great sales workers are motivated by a desire for personal—rather than team—achievement: "Success in sales is about me while success in sales management is about my team. This is where the downside of a strong achievement drive makes itself known. If I'm driven to prove my personal ability, I find it hard (nearly impossible sometimes) to step back and let others take the spotlight."

11. We measure worker performance as deviations from the firm-year-month mean to control for time trends in firm level sales that are unrelated to individual worker effort or ability. This method introduces a small bias against our conclusions. Suppose that a worker with high sales performance is promoted. This worker's sales will no longer be included in the computation of the firm-year-month mean, which in turn increases the measured relative performance of all other workers in the firm, including the subordinates of the newly promoted manager. This causes an upward bias in the estimate of the manager value added for managers with high pre-promotion sales. The direction of the bias goes against our findings that workers with high pre-promotion sales are associated with lower manager value added. Further, this bias is likely to be small in magnitude because we observe an average of 815 workers per firm-year-month, so the promotion of a high sales worker is unlikely to substantially change the average over a large sample.

V. Testing the Peter Principle: Comparing Marginally Promoted Workers

The empirical results so far show that firms promote based on current job performance even though pre-promotion sales negatively predict managerial performance. This evidence is consistent with the idea that firms favor strong sales workers even though they do not make the best managers.

However, we face a missing data problem: we do not observe managerial quality for workers who are not promoted. Because promotions are not random, the unobserved managerial quality of low sales performers who were not promoted may be much worse than the quality of low sales performers who were promoted. As such, even if weaker sales workers appear to make better managers, as in Figure II, it is difficult to know whether this pattern would also hold among the set of workers who were not promoted.

To address this issue, we formalize our analysis in the context of a potential outcomes framework where all workers have managerial potential that is observed only if they are promoted (Angrist, Imbens and Rubin 1996; Holland 1986; Rubin 1974). We then formally state the Peter Principle as a prediction about the causal impact of alternative promotion policies on the distribution of observed manager quality.

We then derive a test for the Peter Principle that is analogous to a Becker outcomes test for discrimination: a firm that prioritizes sales performance at the expense of maximizing managerial quality will set lower standards for managerial potential when evaluating strong sales performers, implying that *marginally* promoted workers with strong sales performance will have lower managerial quality than *marginally* promoted workers with weaker sales performance.

V.A. Model Framework

Consider a group of sales workers. Workers can be promoted to managerial positions ($P = 1$) or not promoted ($P = 0$). Each worker has a potential outcome, M, which captures her "managerial potential"—that is, her managerial quality if she were to be promoted. As is common in missing

data models, M is observed only when $P = 1$.[12]

Let S indicate a worker's pre-promotion sales performance and a remaining vector X indicate collaboration experience, tenure, and all other variables observable to the econometrician. While we observe M only if a worker is promoted, we observe S and X for all workers and treat these as conditioning variables. Firms may also observe variables U that are unobserved by the econometrician. We assume that firms form rational expectations of managerial potential, given what they observe:

$$Q = E[M|S, X, U]. \tag{4}$$

If firms make promotion decisions only to maximize managerial quality, they should promote if $Q > \tau$, where τ is a threshold set so that the firm promotes the desired number of managers: $P = \mathbb{I}(Q > \tau)$. If firms also care about sales performance or other variables, they may allow the promotion threshold τ to vary with these variables: $P = \mathbb{I}(Q > \tau(S, X, Z, U))$. To facilitate our discussion of identification in Section V.B., we allow for an instrument, Z, that predicts promotion but is unrelated to managerial potential.

Given this setup, the Peter Principle can be formally stated as follows:

DEFINITION 1. (Peter Principle) Firms use promotion policies $P = \mathbb{I}(Q > \tau(S, X, Z, U))$ that prioritize sales performance at the expense of maximizing managerial quality. Equivalently, there exists an alternative promotion policy $\tilde{P} = \mathbb{I}(Q > \tilde{\tau}(S, X, Z, U))$ such that $E[M|\tilde{P} = 1] > E[M|P = 1]$, $\partial E[\tilde{P}|S, X]/\partial S < \partial E[P|S, X]/\partial S$, and $E[\tilde{P}] = E[P]$.

In the following section, we describe how we construct such an alternative policy $\tilde{P}$.

12. This set up follows the Rubin Causal Model (RCM) described in Holland (1986). In our model, promotion P corresponds to the treatment D in the standard RCM model; a worker's managerial quality if promoted, M, corresponds to the standard potential outcome conditional on treatment Y^1. However, the potential outcome Y^0 (a worker's managerial quality if she is not promoted) is undefined in our setting. As such, instead of estimating the causal impact of promotion on managerial quality (which is also undefined), we focus on estimating the causal impact of different promotion *policies* on observed managerial quality: $E[M|P = 1]$ versus $E[M|\tilde{P} = 1]$ for some other promotion policy $\tilde{P}$. We do this because the Peter Principle can be stated as a hypothesis that promotion policies that strongly favor current performance do not necessarily maximize the managerial quality of promoted workers, relative to other policies that may place less emphasis on current performance.

V.B. Empirical Strategy

Under the above framework, a test of the Peter Principle is equivalent to a Becker outcomes test for discrimination, where we compare the managerial quality of marginally promoted high and low sales workers. Intuitively, if marginally promoted low sales workers make better managers than marginally promoted high sales workers, then the firm could improve managerial quality by following an alternative policy $\tilde{P}$ that promotes fewer high sales workers (and more low sales workers) on the margin.

Identifying Marginally Promoted Workers

We identify the managerial quality of marginally promoted workers using an instrument for promotion. Instrument compliers—workers who would not have been promoted but for the instrument—can be thought of as a set of marginally promoted workers.

Before discussing our instrument and its validity, we first formalize our approach. For intuition, consider a binary instrument Z, where workers with values of $Z = 1$ are more likely to be promoted. We define $k(S, X)$ to be the average quality of instrument compliers—that is, workers who are promoted under Z=1, but not under Z=0:

$$k(S, X) \equiv E[M|S, X, P^{Z=1} > P^{Z=0}]. \tag{5}$$

We are interested in estimating and comparing $k(S, X)$ for workers with high and low sales performance. The Peter Principle implies that $k(S^{\text{High}}, X) < k(S^{\text{Low}}, X)$. The following proposition shows how we can estimate $k(S, X)$ from our data:

PROPOSITION 1. (Estimating equation) Consider the regression

$$M_i \times P_{it} = a_{0j} + a_{1j}P_{it} + \beta_j X_{it} + \varepsilon_{it} \tag{6}$$

for workers with sales S_{it} falling in sales bin S_j. Suppose that we have a valid binary instrument for promotion: Z such that $P_{it}^{Z_{it}=1} \geq P_{it}^{Z_{it}=0}$ and the exclusion restriction holds

$(M_i \perp \{Z_{it}, P_{it}^{Z_{it}=1}, P_{it}^{Z_{it}=0}\}|X_{it}, S_{it})$. Then it is the case that

$$\hat{a}_{1j}^{IV} = E[M_i|S_{it} \in S_j, P_{it}^{Z_{it}=1} > P_{it}^{Z_{it}=0}] \equiv k(S_{it} \in S_j, X_{it}) \tag{7}$$

That is, $\hat{a}_{1i}^{IV}$ is a consistent estimate of the average managerial quality of workers with $S_{it} \in S_j$, who are compliers to the promotion instrument Z.[13]

Proof. See Online Appendix Section C.

This result is analogous to Angrist, Imbens and Rubin (1996), who show that IV estimates identify local average treatment effects for instrument compliers. Following Abadie (2003), Proposition 1 takes this same framework and focuses on estimating local selection rather than treatment effects: we estimate $E[Y^1|D^{Z=1} > D^{Z=0}]$ rather than $E[Y^1 - Y^0|D^{Z=1} > D^{Z=0}]$.

In equation (6), $M_i \times P_{it}$ is the *observed* managerial quality of worker i if worker i is promoted at time t. If that worker is not promoted in that period—or if he or she is never promoted—the left side takes a value of 0. The coefficient of interest is on the dummy P_{it} for whether or not worker i is promoted at time t. This regression is structured so that the OLS coefficient a_{1j}^{OLS} estimates *average* managerial quality of promoted workers with pre-promotion sales performance falling in the $j-th$ bin. To identify the managerial quality of *marginal* promoted sales workers, we instrument P_{it} with Z_{it}. The IV estimate a_{1j}^{IV} is equivalent to $k(S, X)$ for $S_{it} \in \mathbf{S}_j$.

We estimate equation (6) separately for three groups of pre-promotion sales performance. If a_{1j}^{IV} is decreasing in j, then the managerial quality of the marginally promoted worker is lower for higher sales performers, indicating discrimination in their favor.

13. In practice, our instrument will be continuous. In the continuous case, equation (5) is replaced by $k(S, X, z) = \frac{\partial E[MP|S,X,Z=z]/\partial z}{\partial E[P|S,X,Z=z]/\partial z}$ and equation (7) is replaced by the corresponding LATE representation $k(S, X, z) = E[M|S, X, \lim_{z' \downarrow z} P^{z'} = 1, \lim_{z' \uparrow z} P^{z'} = 0]$. We focus on the binary case for intuition, because it can be interpreted as a close analogue to the Wald estimate. In general, we require that the probability of promotion is monotonic in the instrument, following Angrist, Imbens and Rubin (1996) and Heckman and Vytlacil (2005).

Instrument and Identifying Assumptions

The approach illustrated above requires a valid instrument Z for promotion. We use a jackknife IV approach where we instrument for an individual's promotion status P_{it} in equation (6) with the average promotion rate in their firm-month, leaving out worker i and her teammates. The estimated coefficient a_{1j}^{IV} identifies the manager value added of sales workers who were promoted *on the margin*—that is, those who were promoted only because the firm made many promotions that month overall, and who would not have been promoted in months with fewer promotions. By separately estimating equation (6) for high-, mid-, and low-performing sales workers, we can compare the managerial quality of marginally promoted workers from each of these groups.

This instrument must satisfy two identifying conditions. First, it must be positively and monotonically correlated with workers' individual probabilities of promotion (instrument relevance). Table II, discussed earlier, shows that there is indeed a strong positive relation between jackknife firm-year-month promotion rates and individual promotion.

Second, the promotion rate instrument must be orthogonal to her managerial potential M, conditional on observables (instrument exclusion). One may be concerned, in particular, that promotion rates reflect other firm level factors that may subsequently have a direct impact on how well managers perform after promotion. As an illustration, suppose that demand for a firm's products is particularly high in a given period and the firm responds by promoting more workers, who then take on managerial roles. If demand continues to increase, then these newly promoted managers will preside over strong subordinate sales growth: we would not want to attribute this trend to their managerial quality.

However, recall from equation (3) that we estimate M as a manager's value added to subordinate sales *controlling for* worker and firm-year-month fixed effects. The inclusion of these fixed effects creates a measure of managerial quality that is unrelated to aggregate firm-time patterns such as overall consumer demand or firm expansion plans that may be correlated with our promotion rate instrument. The instrument is not significantly correlated with manager value added, nor is manager value added correlated with other factors that may drive promotion opportunities (see Online Appendix Table A5).

Another potential concern is reverse causality: if a given worker is particularly strong, the firm may increase its promotion rate to promote him or her. Using a jackknife approach and leaving out a worker's own promotion status (and that of her teammates) severs the correlation between our instrument and an individual worker's quality.

Other scenarios may bias our estimates of the quality of marginally promoted workers: for example, workers promoted in high-promotion months may be assigned to different client portfolios than those promoted in other months. We lack the data to fully rule out these types of scenarios. However, we do note that potential bias in our measure of manager value added should not affect our findings unless the direction of the bias is also correlated with pre-promotion sales performance. Our analysis is primarily concerned with comparing the quality of marginally promoted workers from different bins of pre-promotion sales performance. As such, biases in the measured quality of marginally promoted workers will not affect our conclusions unless they apply *differentially* for high and low sales workers.

More generally, one may be concerned that marginally promoted low sales workers are somehow different from marginally promoted high sales workers in a way that makes them difficult to compare. For example, suppose that all marginal low sales workers were promoted in 2005, whereas all marginal high sales workers were promoted in 2010. If sales conditions were worse in 2010, this would not necessarily mean that firms were discriminating in favor of high sales workers. A similar concern would apply if marginal high sales workers were associated with one set of firms, whereas marginal low sales workers came from another. To increase the likelihood that marginally promoted high and low sales workers are drawn from comparable groups, we measure pre-promotion sales performance within a firm-year-month. This means that, by construction, we are comparing the quality of marginally promoted low and high sales workers coming from the same firm, at the same time.

Constructing a Counterfactual Promotion Policy

Finding differences in the managerial quality of marginally promoted workers with different sales records allows us to construct an explicit alternative promotion policy $\tilde{P}$ that improves expected

managerial quality among the promoted:

$$\tilde{P}(S,X) = \begin{cases} P^{Z=1} \text{ if } k(S,X) > \tilde{\tau}, \\ P^{Z=0} \text{ otherwise} \end{cases} \tag{8}$$

This promotion rule essentially increases the promotion rates of individuals from groups with high managerial quality on the margin and decreases the promotion of groups whose marginally promoted managers appear to be low quality. For simplicity, suppose we divide our instrument into above ($Z = 1$) and below ($Z = 0$) median promotion rates firm-year-months. The alternative rule $\tilde{P}$ takes a firm's existing promotion policy P and assigns individuals from high marginal quality groups (e.g., those with covariates S and X such that the expected quality of compliers given their covariates $k(S,X)$ is greater than some threshold) to the promotion status they would have under the existing policy P if they were faced with the high value of the instrument. If an individual is from a low marginal quality group, then $\tilde{P}$ assigns individuals to the promotion status they would have if $Z = 0$. The specific threshold for what is considered "high" marginal quality is given by $\tilde{\tau}$, which is set to keep the number of promoted workers constant, $E[P] = E[\tilde{P}]$.

If we find that the marginally promoted high sales worker has lower managerial quality, then $\tilde{P}$ essentially tells the firm to promote low sales workers as if they were planning to have a high promotion rate and to promote high sales workers as if they were planning to have a low promotion rate. Such a rule would, by construction, promote fewer high sales workers. To see that it would also improve expected managerial quality, consider how $\tilde{P}$ differs from P. If a low sales worker is promoted under P, she would also be promoted under $\tilde{P}$: these are the always takers. The low sales people who are promoted under $\tilde{P}$ and not P are, by construction, compliers: those who would not have been promoted had they faced a low average promotion rate, but who are promoted if they face a high promotion rate. Similarly, the high sales people who are no longer promoted under $\tilde{P}$ are also compliers: those who are promoted under the original policy P, but who are no longer promoted once they face low promotion rates under $\tilde{P}$. The return to promoting based on $\tilde{P}$ instead of P is then the quality of the low sales compliers minus the quality of the high sales compliers: this is exactly what is estimated by $a_{1L}^{IV} - a_{1H}^{IV}$. Recall that in Section V.A, we defined

the Peter Principle as the claim that there exists some alternative promotion policy that places less emphasis on current job performance while achieving a better managerial match. If $a_{1L}^{IV} - a_{1H}^{IV} > 0$, then $\tilde{P}$ is an explicit example of such a promotion policy.[14]

V.C. Main Results

Figure III reports our estimates of the manager value added of marginally promoted salespeople across terciles of prior sales performance, as specified by equation (6). Panel A presents a monotonically decreasing relationship between sales and managerial quality: the marginally promoted worker among those in the lowest tercile of sales performance has higher manager value added than those in the middle sales tercile, who in turn have higher manager value added than those in the high sales tercile. We can reject equality across the terciles with a p-value of 0.004.

Panel B of Figure III presents a particularly key contrast between the managerial quality of marginally promoted workers who were top ranked in sales within a team versus those who were not. In Table II, we showed that the top-ranked salespeople in a given team were almost three times as likely to be promoted as the average sales worker. Here, we find that top-ranked sales workers are actually worse managers on the margin (we can reject equality with a p-value of 0.09). Panels A and B of Table IV provide the regression analogues for this result. Together, these results suggest that firms discriminate in favor of high sales workers by applying a lower promotion threshold for expected managerial quality, leading marginally promoted high sales workers to be worse managers.

As a result, firms can improve expected managerial quality by placing less weight on sales in promotion decisions. In particular, we can construct an alternative promotion rule $\tilde{P}$ as described in equation (8). Recall that Z is our instrument for promotion, the jackknife average promotion rate for a firm in a given year-month. For simplicity, consider the following rule that uses a binary

14. We can construct $\tilde{P}$ even if the probability of being a complier varies across sales bins. For example, if low sales workers are promoted only when a firm has many vacancies, then they may be more likely to be instrument compliers, relative to high sales workers who are more frequently promoted. Such differences do not affect our ability to construct the alternative promotion policy $\tilde{P}$. The number of low and high sales compliers would affect the number of workers who switch promotion status between a firm's initial policy P and the proposed alternative $\tilde{P}$; however, regardless of how many (or few) compliers there are, the expected change in managerial quality resulting from following policy $\tilde{P}$ instead of P is still given by $a_{1L}^{IV} - a_{1H}^{IV}$, which we estimate to be positive.

version of this instrument (indicating above- and below- median average promotion rates):

$$\tilde{P}(S,X) = \begin{cases} P^{Z=1} \text{ if } S \in \mathbf{S}_1, \\ P^{Z=0} \text{ if } S \in \mathbf{S}_3, \\ P \text{ if } S \in \mathbf{S}_2 \end{cases} \tag{9}$$

The policy presented by equation (9) promotes low sales workers according to the firm's existing promotion rule, but as if these workers faced high average promotion rates ($Z = 1$). Conversely, it promotes high sales workers according to the existing promotion rule, but as if they faced low average promotion rates ($Z = 0$). It keeps the promotion policy constant for average sales workers. This policy has the effect of promoting more low sales workers and fewer high sales workers on the margin. Given the results in Figure III, $\tilde{P}$ would lead to an increase in managerial quality among promoted workers.[15]

The existence of a promotion rule that increases managerial quality by placing less weight on sales provides direct evidence in favor of the Peter Principle. The rule $\tilde{P}$ accomplishes this by changing promotion patterns slightly on the margin. In Section VII.A, we assess the potential increases that firms can achieve from a more drastic change in promotion policies that is aimed solely at maximizing managerial quality.

V.D. Other Sources of Information: Collaboration Experience

Next, we present results showing that workers' collaboration experience—as measured by the number of people they share sales credits with—is an observable characteristic that firms could weight more positively in promotion decisions if their only goal were to improve managerial match quality. We emphasize that our results relating to collaboration experience are meant to be suggestive rather than conclusive, as some of the estimated magnitudes are economically meaningful but not always statistically significant.

Panel A of Figure IV shows that, conditional on sales performance, tenure, and firm-time

15. In practice, one would have to make sure that the same number of workers were promoted, which could be achieved by implementing the continuous version of this rule, given by equation (8).

effects, workers with more collaboration experience appear less likely to be promoted. One possible explanation for this result is that assigning credit for collaborative work is difficult, making solo work de facto more rewarded, conditional on total output (Sarsons 2019). Panel B, however, shows that promoted workers with more collaboration experience appear to have better managerial performance. These patterns are also present in Table V, which shows that collaboration experience negatively predicts promotion but positively predicts manager value added within the sample of promoted managers.

We also find that workers who never collaborate with others—the so called lone wolves—fare particularly poorly when they are promoted into managerial roles. Stereotypically, lone wolves are known within the sales profession to be "the deeply self-confident, the rule-breaking cowboys of the sales force who do things their way or not at all" (Dixon and Adamson 2011). In Table V, we find that firms are significantly more likely to promote lone wolves, and yet lone wolves have lower manager value added. Collectively, these results suggest that firms underprioritize collaboration in promotion decisions, at least from the perspective of increasing the quality of managers. Our findings are also consistent with Oettl (2012), which documents the importance of helpfulness (as measured by acknowledgments) as an indicator of scientific productivity.

These results are subject to several caveats. First, the positive relation between collaboration experience and manager value added remains sizable but becomes statistically insignificant if we adopt the framework in Section V to compare the quality of marginally promoted workers with low and high collaboration experience (see Online Appendix Table A6 and Figure A5). Second, high sales collaborators may be assigned to different types of management positions, although we find no evidence of such systematic assignment in the data (see Online Appendix Table A7). Finally, our measure of collaboration experience may capture more than teamwork skills. For example, workers assigned to team-based sales may acquire skills for dealing with complex products and get experience dealing with demanding clients. Alternatively, firms may select workers that are conscientious or detail oriented for team-based assignments. It could be these skills rather than teamwork skills *per se* that cause workers with high collaboration experience to be better managers. Our goal is this section is not to show a causal link between teamwork skills and management quality, but

rather to show that other observable variables may also predict manager value added. This kind of predictive power can help firms promote workers with high managerial potential (if that is the firm's only goal), even if the mechanism driving the correlation is unclear.

V.E. Could Managers Contribute in Other Ways?

Our results so far show that managers with high pre-promotion sales have lower value added from the perspective of increasing subordinate sales. It is possible, however, that managers with high pre-promotion sales contribute to firm value in other ways. In this section, we consider two key sets of potential alternative contributions: direct sales by managers and workforce management activities such as reducing turnover or recruiting new employees.

Manager Direct Sales

If managers are directly involved in making sales, then promoted managers with strong sales skills may contribute to the firm relatively more through direct sales. In addition, managers with high pre-promotion sales experience may take sales credits and opportunities away from their subordinates. If so, these managers may increase total team output, but at the expense of subordinate sales, making them appear as if they have lower value added according to our measure of managerial quality.

We begin by noting that we do not believe that managers engage in direct sales. As discussed in Section III.B, managers focus on directing sales activities, leaving actual sales to their subordinates. In fact, discussions with industry experts indicate that managers do not engage in direct sales because doing so would create conflicts of interest given that they are also charged with assigning subordinates to accounts and overseeing bonuses and performance reviews.

Nonetheless, managers do receive sales credits in our data. Managers' sales credits are highly correlated with the sum of monthly credits across their subordinates, indicating that managers are awarded credits as a function of the sales of their subordinates. Online Appendix Figure A6 shows the relation between a manager's monthly sales credits and the total monthly sales credits of her subordinates. The relationship is highly linear with a correlation of 0.8. In one-quarter of

observations, a manager's monthly sales credits are equal to the sum of her subordinates' credits.

However, our data do not specify the exact reason why managers are credited on sales, so we cannot completely rule out the possibility that their credits partly reflect direct participation in sales. We show that managers with high pre-promotion sales performance continue to make worse managers even if we assume (unrealistically) that the manager sales credits we observe in our data reflect direct manager sales that may have substituted for subordinate sales.

First, we show that managers with higher pre-promotion sales do not contribute a greater share of sales credits to their teams' total sales credits. If this were the case, we would expect managers with higher pre-promotion sales to be assigned a greater proportion of their team's total sales credits. Online Appendix Table A8 shows no substantial relationship between pre-promotion sales performance and the share of credits attributed to a manager, defined in several ways. Across all specifications, we find no relationship between "credit hogging" and pre-promotion sales performance. Our estimates are precise zeros. For instance, we estimate that doubling pre-promotion sales is associated with a 0.4 percentage point increase in the share of a team's credits that are attributed to the manager.

Our next approach addresses the concern that managers with high pre-promotion sales may take more sales opportunities away from their subordinates, which would cause these managers to have lower value added according to our measure of managerial quality. We treat manager credits as if they reflect direct real sales by the manager and *reattribute* these sales to the manager's subordinates so that a subordinate's performance is measured as the sum of her individual sales plus her allotment of the manager's sales credits. We then reestimate manager value added as a manager's contribution to increasing subordinates' sales plus each subordinate's share of reallocated managerial credits. If a high sales manager raises team production through direct sales, then he or she would have positive manager value added under this metric because direct manager sales are attributed back to subordinate sales. We find the same negative relation between pre-promotion sales performance and the value added of the marginally promoted manager under both allocation rules (see Online Appendix Table A9 and Figure A7).

Finally, in Online Appendix Table A10, we regress manager sales credits in each month on the

manager's pre-promotion sales performance. Managers with greater pre-promotion sales do indeed have more manager sales credits in the data. However, this positive correlation is driven by the fact that managers with greater pre-promotion sales are, on average, assigned to larger teams (see Table VI column (1)) and assigned to subordinates with better sales performance *ex ante* (see Online Appendix Table A11 Columns (1)-(3)). After controlling for total subordinate sales in each month, the manager's pre-promotion sales performance is no longer predictive of manager sales credits. The coefficient on the manager's pre-promotion sales becomes negative, shrinks toward zero, and is insignificantly different from zero. Altogether, these results imply that, even if manager credits in our data partially reflect manager direct sales, managers with higher pre-promotion sales do not contribute more to the firm through direct sales.

Workforce Management

In addition to supporting the sales efforts of their subordinates, managers may also contribute by reducing costly worker turnover or by recruiting new sales employees to expand the operations of the firm. In Table VI and Online Appendix Table A12, we show that managers with high pre-promotion sales do not appear to be associated with better performance along these dimensions.

Table VI begins by showing that high sales performers are more likely to be assigned to larger teams upon promotion: column (1) indicates that doubling pre-promotion sales is correlated with 0.124 more subordinates. Given this, the remaining columns report the correlation between a manager's pre-promotion sales and various measures of workforce management. First, to assess whether managers are able to grow their teams, column (2) considers the net change in a manager's team size each month, averaged over that manager's tenure. Column (3) examines a manager's ability to bring on new workers by looking at the fraction of the team that is new. Similarly, column (4) examines whether managers reduce turnover by looking at the fraction of a team that exits. Finally, column (5) examines a manager's ability to retain good sales workers while letting go of poor performers by examining the sales percentile of workers who exit. In all cases, we find an economically small and insignificant relationship between a manager's pre-promotion sales and performance on these dimensions. Online Appendix Table A12 finds similar results when we

consider the workforce management performance of the *marginally* promoted manager.

Overall, we don't find that pre-promotion sales performance is associated with better managerial performance in terms of team expansion, reducing turnover, or selective retention. If anything, the fact that high sales performers are assigned to manage larger teams suggests that firms incur greater losses by giving poor managers greater responsibility. However, we acknowledge that we cannot rule out the possibility that managers with high pre-promotion sales contribute in other unobserved ways. For example, these managers may effectively allocate financial resources, choose new product lines, or be good matches for higher-level managerial positions.

VI. Potential Alternative Explanations

The results above are consistent with the Peter Principle, which we define as promotion policies that favor higher–performing workers at the expense of promoting the best potential managers. In this section, we explore whether alternative explanations or biases could explain our findings such that firms in our sample actually are promoting the best potential managers.

VI.A. Individual Mean Reversion

Lazear (2004) cautions that mean reversion can generate patterns that, on the surface, look like the Peter Principle. Lazear argues that the best worker may, *correctly*, be the best candidate for management but may display a decline in performance after promotion because of mean reversion. In Lazear's model, there exists within-person mean reversion over time, but the cross-sectional correlation between pre-promotion sales and post-promotion manager value added remains positive, so the best potential managers are still the highest-performing sales workers. However, we find in our data that the cross-sectional correlation between pre-promotion sales and post-promotion manager value is negative, implying that firms could improve manager quality by promoting different workers who are less strong at sales.

VI.B. Group-Level Mean Reversion

A remaining concern is that there may instead be group-level regression to the mean. In some cases, workers are promoted to replace their former managers. If a worker is promoted as a result of a transitory shock to his or her team's sales performance, then other members of the sales team (who are now her subordinates) may subsequently experience a decline in their sales, reducing our estimates of the new manager's value added. We address this concern by restricting our analysis to the sample of managers who are promoted to manage subordinates who were not part of their original sales group. Within this sample, we continue to find a negative relationship between manager value added and the manager's pre-promotion sales (see Online Appendix Figure A8 and Table A13).

This set of results is also inconsistent with alternative explanations involving spillover effects. For example, when a high sales worker is promoted to manage her previous teammates, the sales worker's existing pipeline could roll over to the sales worker's former teammates, giving a temporary boost to the newly promoted manager's estimated value added. The fact that our results look very similar when we focus on the performance of workers who are promoted to manage *different* teams suggests that our findings are not driven by spillover effects.

VI.C. Nonrandom Assignment

Lastly, Online Appendix Table A11 explores the potential threat posed by the nonrandom assignment of managers to subordinates. We find that better salespeople who are promoted tend to be assigned better subordinates: a doubling of a manager's pre-promotion sales is correlated with an approximately 20% increase in the prior sales of the subordinates to whom she is assigned. In general, a correlation between the pre-promotion sales of newly promoted managers and the *level* of performance of their assigned subordinates should not affect our results because we estimate manager value added from *changes* in subordinate performance under the new manager. However, we remain concerned that our estimates of manager value added will be biased by a downward trend (for reasons unrelated to the true manager quality) in the performance of subordinates assigned to managers with strong pre-promotion sales. To check this, we examine subordinates' performance

within 7-9, 4-6, and 1-3 month windows prior to the manager's arrival and find no evidence of pretrends. We also consider the possibility that managers assigned to subordinates with high prior sales will appear to have lower value added because these subordinates have less scope for improvement. However, we find no significant or substantial relationship between subordinates' prior sales and our estimates of manager value added.

VII. The Trade-off Between Manager Match Quality and Promotion-Based Incentives

In this section, we consider the mismatch costs of the Peter Principle and whether firms appear aware of them.

VII.A. What are the Performance Losses from Mismatch?

First, we estimate the cost of managerial mismatch. We do so by comparing the managerial quality of the observed promotions against a counterfactual policy in which firms promote the best potential managers. This analysis sets aside tournament incentives and other potential benefits of firms' promotion rules to focus instead on the costs of managerial mismatch. Our estimates may be interpreted as the match quality that firms forgo to use promotions for other purposes.

To begin, recall that equation (4) states that firms form beliefs about a worker's managerial potential based on what it observes: $Q = E[M|S, X, U]$. Firms then promote workers whose Q exceeds some threshold that may vary for high sales workers: $P = \mathbb{I}(Q > \tau(S, X, U, Z))$.

In Section V, we tested the Peter Principle using aggregate promotion rates to identify lower managerial quality among marginally promoted high sales workers. This difference in managerial quality among marginally promoted workers allowed us to construct a counterfactual promotion policy, $\tilde{P}$ given by equation (9), that improves managerial quality by promoting fewer top sales workers on the margin.

This analysis did not require assumptions about the functional form of Q or P. However, if we impose additional assumptions, we can identify the average managerial quality resulting from

a wider range of promotion policies, including the managerial-quality-maximizing policy based on observables. Consider the following selection correction model, where we specify functional forms for Q and P:

$$Q = a_0 + a_1 S + a_2 X + e \tag{10}$$

and the firm's promotion policy is given by

$$P = \mathbb{I}(b_0 + b_1 S + b_2 X + b_3 Z + v > \tau) \tag{11}$$

where e and v represent jointly normally distributed errors. This model is a special case of the baseline model we considered in Section V.[16]

We estimate equations (10) and (11) using the standard Heckman selection procedure. We instrument for selection into the observed sample, (i.e., promotion to management), using the jackknife firm-year-month average promotion rates previously discussed. Our results are reported in Online Appendix Table A14. As in both the baseline OLS illustrated in Panel B of Figure II and the non parametric IV illustrated in Figure III, we continue to find that better salespeople make worse managers. The advantage of this parametric approach is that it allows us to recover selection–corrected estimates of a_0, a_1, and a_2 from equation (10), which we use to form estimates of managerial potential $\hat{M}^1$ for all workers.

To calculate the cost of mismatch, we examine how predicted managerial performance differs among three groups of workers: (1) actually promoted salespeople, (2) nonpromoted salespeople among the promoted worker's peers, and (3) the top predicted manager among a promoted sales worker's peers. Peers are defined as other salespeople in a team managed by the same manager in the same time period. We interpret case (3) to be the performance-maximizing promotion decision

16. We also assume that there are no additional variables U that are observed by the firm but not by the econometrician. Under this assumption, our analysis estimates the distribution of expected managerial quality under the promotion rule that maximizes managerial potential M given observables S and X. If this second assumption is not met—the firm observes additional information U—then the firm can construct a promotion rule that does an even better job of maximizing managerial quality. If so, the improvement in managerial quality that we estimate should be thought of as a lower bound for what the firm can achieve if it made use of all its information.

under the restriction that mobility and other frictions prevent the firm from promoting among the entire organization, so that firms must promote among the peers of promoted workers. If we relax this restriction, then the estimated costs of mismatch will further increase.

Figure V shows the distributions of predicted manager value added in the three groups of workers. The mean predicted improvement in subordinate sales performance is scaled to zero for the sample of promoted workers. The mean change in subordinate sales performance for the sample of nonpromoted workers is 0.12, implying that firms' current promotion policies do slightly worse than promoting at random. This is expected, as firms' current promotion policies strongly favor sales performance, and sales performance negatively predicts manager value added (even when we allow for nonlinear relationships in the counterfactual simulation).

The mean in the sample of best predicted managers is 0.28, implying that subordinate performance could improve by approximately 30% if firms pursued an alternative promotion policy of promoting the best predicted manager within a sales team.[17] Our estimate is not meant to suggest that firms would actually achieve 30% gains in sales if they switched to a promotion policy in which they promoted the best potential managers. This counterfactual estimate ignores potentially large productivity declines that could result from lost incentive and other morale effects that may occur if firms switched away from the current promotion policy of rewarding high sales with promotions. Thus, the 30% should be viewed as a lower bound for how large the incentive benefits of promoting the top sales workers would have to be to rationalize the current set of promotion policies.

VII.B. Which Firms Place Less Weight on Sales Performance in Promotions?

If firms are aware of the trade-off between maximizing managerial match quality and providing incentives for workers, we would expect firms to behave differently depending on the specific costs and benefits they face. For example, firms in which managers have greater responsibility may place

17. We estimate a 30% gain in sales if firms switched to a promotion policy in which they promoted the best potential managers. This number would be even higher if we adjusted for the fact that promoting lower sales workers results in fewer forgone sales, given that managers are no longer engaged in direct frontline sales.

more weight on picking the best managers and may be more willing to promote workers who are weaker in terms of sales performance. Similarly, firms that have chosen to use alternative ways of incentivizing worker effort may prioritize managerial match quality more in their promotion decisions.[18]

Table VII, columns (2) and (3) consider the first possibility by examining how promotion policies differ across firms by the supervisory responsibilities assigned to managers. We use the number of subordinates associated with each manager as a proxy for managerial responsibility and then take averages to obtain a firm-level measure. We then augment equation (2), which examines a worker's probability of promotion, by interacting our measures of worker sales and collaboration experience with the log of average team sizes each firm-year. Our estimation also controls for the direct effects of all variables.

We find that firms with larger subordinate teams tend to place less weight on sales performance in promotion decisions. A doubling of the average team size reduces the predictive relationship between sales performance and promotion by almost 30%. By contrast, firms with larger team sizes place relatively more weight on collaboration experience. These findings suggest that when the costs of managerial mismatch are particularly high, firms are more willing to sacrifice the incentive benefits of performance-based promotion tournaments to promote better managers.

Next, we consider how promotion policies vary with the use of incentive pay. We construct a firm-level measure of pay for performance as the ratio of commissions and bonuses to base salary, averaged across all workers in the firm within each calendar year. Before proceeding, recall that we observe base salaries for only a subsample of firms, leading to an approximate 25% decline in sample size for this analysis, and that salary data can be missing or measured with error. Nevertheless, we believe we are able to construct a reasonable, if noisy, proxy for the strength of pay-for-performance incentives across firms in our data.

We regress whether a worker is promoted on the interaction between our measure of pay for performance and worker sales and collaboration experience. In Table VII, columns (4) and (5), we

18. Promotion policies may also differ within a firm, depending on which manager conducts the performance evaluation and controls the promotion decision (see, e.g., Frederiksen, Lange and Kriechel 2017). Unfortunately, we lack the data to explore within-firm variation in promotion policies and therefore focus on across-firm variation.

find that firms with relatively strong pay for performance tend to implement promotion policies that are less sensitive to worker sales performance and more sensitive to collaboration experience. This is consistent with the idea that pay for performance incentives can partially offset the need to provide promotion based-incentives, as discussed in Ekinci, Kauhanen and Waldman (2018). However, pay for performance may be an expensive substitute for promotion-based incentives, especially if workers value the security, stature (e.g. DellaVigna and Pope 2016; Larkin 2011), or external signaling abilities associated with promotions (DeVaro and Waldman 2012; Waldman 1984*a*). The sales positions we study already have strong pay for performance relative to many other occupations. The fact that we still observe evidence of the Peter Principle in the sales setting suggests that it may be difficult to fully substitute for promotion-based incentives.

VII.C. Do Promotions Discourage Passed-Over Workers?

In Table VIII, we test whether workers appear to be discouraged when a worse-performing teammate is promoted. We examine the subsequent retention and sales for a promoted worker's teammates, depending on whether these teammates had higher or lower sales than the promoted worker at the time of the promotion event. We regress retention and changes in sales for worker i on a dummy for whether i had greater sales than the promoted teammate at the time of the promotion event. To account for possible within-person mean reversion, we also control for worker i's sales ranking within the team and level of sales at the time of the promotion. In other words, we compare workers with the same sales and team ranking, but with different ranks *relative* to the promoted worker.

We find that workers are much more likely to leave the firm if a teammate with worse sales performance is promoted. Twelve months after the promotion event, higher-performing sales workers who were passed over for the promotion are 23 percentage points less likely to remain with the firm. Sales workers also exhibit a greater decline in relative sales in the three months immediately following the promotion of a teammate with a poorer sales record. This decline, however, is short-lived conditional on the worker remaining with the firm and is insignificant by month 6 after the promotion event. These results suggest that promoting a lower-performing sales

worker may reduce the morale or incentives of higher-performing teammates, particularly in terms of retention. This may be an important reason why firms choose to promote high-performing workers, even if they do not make the best managers.

Overall, the analysis in this section shows that firms seem to be aware of the trade-off between maximizing managerial quality and providing promotion-based incentives. The Online Appendix presents a simple stylized model to illustrate that, under reasonable assumptions, it may be efficient for some firms to offer strong promotion-based incentives. In other words, the Peter Principle imposes costs that firms may optimally choose to bear.

VIII. Conclusion

We use detailed microdata on the performance and promotions of sales workers at a large number of firms to provide the first large-scale test of the Peter Principle, the notion that firms prioritize current performance when making promotion decisions, at the expense of choosing those best suited for the post-promotion role. Consistent with this hypothesis, we find that firms are substantially more likely to promote top salespeople, even when these workers make worse managers both on average and on the margin. This behavior results in firms promoting workers who decrease subordinate performance by 30%, relative to a promotion policy that optimizes match quality.

We caution against interpreting these results as evidence that firms have mistaken beliefs or behave inefficiently. Firms may heavily weight current job performance in promotion decisions to encourage workers to exert effort in their current job roles and to maintain norms of fairness. In addition, the availability of relatively clear measures of worker productivity among frontline sales workers may lead organizations to emphasize these characteristics rather than other, more subjective or fungible employee characteristics in promotion decisions.

Indeed, our results suggest that firms are aware of these benefits and appear to actively manage the trade-off between providing incentives and promoting the best potential managers: firms place less emphasis on current job performance in promotions where managerial roles entail greater responsibility and where current performance is rewarded by relatively strong pay for performance. Overall, our results imply that managerial match quality, tournament incentives,

and other objectives of job promotions are not perfectly aligned. The trade-off between incentives and match quality is likely to be an important consideration for any firm or institution in which the skills required to succeed at one level in the organizational hierarchy differ from the skills necessary to succeed at a higher level.

University of Minnesota
Massachusetts Institute of Technology
and National Bureau of Economic Research
Yale University
and National Bureau of Economic Research

IX. References

Abadie, Alberto, "Semiparametric instrumental variable estimation of treatment response models," Journal of Econometrics, 113 (2003), 231–263.

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," Econometrica, 67 (1999), 251–333.

Abowd, John M., Francis Kramarz, David N. Margolis, and Kenneth R. Troske, "The Relative Importance of Employer and Employee Effects on Compensation: A Comparison of France and the United States," Journal of the Japanese and International Economies, 15 (2001), 419–436.

Adhvaryu, Achyuta, Vittorio Bassi, Anant Nyshadham, and Jorge Tamayo, "No Line Left Behind: Assortative Matching Inside the Firm," Working Paper (2019).

Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin, "Identification of Causal Effects Using Instrumental Variables," Journal of the American Statistical Association, 91 (1996), 444–455.

Arnold, David, Will Dobbie, and Crystal S Yang, "Racial Bias in Bail Decisions," The Quarterly Journal of Economics, 133 (2018), 1885-1932.

Baker, George P., Michael C. Jensen, and Kevin J. Murphy, "Compensation and Incentives: Practice vs. Theory," The Journal of Finance, 43 (1988), 593–616.

Becker, Gary *The Economics of Discrimination* (Chicago: University of Chicago Press, 1957).

Becker, Gary S, "Nobel Lecture: The Economic Way of Looking at Behavior," Journal of Political Economy, 101 (1993), 385–409.

Benson, Alan, "Do Agents Game Their Agents' Behavior? Evidence from Sales Managers," Journal of Labor Economics, 33 (2015), 863–890.

Bertrand, Marianne, and Antoinette Schoar, "Managing with Style: The Effect of Managers on Firm Policies," Quarterly Journal of Economics, 118 (2003) 1169–1208.

Bloom, Nicholas, and John Van Reenen, "Measuring and Explaining Management Practices across Firms and Countries," The Quarterly Journal of Economics, 122 (2007), 1351–1408.

Bureau of Labor Statistics, *2018 Occupational Employment Statistics: Sales and Related Occupations*, *Bureau of Labor Statistics*, https://www.bls.gov/ooh/sales/home.htm.

Cullen, Zoë, and Ricardo Perez-Truglia, "How Much Does Your Boss Make? The Incentive Effects of Inequality," Working paper, 2018.

DellaVigna, Stefano, and Devin Pope, "What Motivates Effort? Evidence and Expert Forecasts," The Review of Economic Studies, 85 (2016), 1029-1069.

Deutsch, Claudia, "Holding on to Technical Talent," *New York Times* (1986).

DeVaro, Jed, and Michael Waldman, "The Signaling Role of Promotions: Further Theory and Empirical Evidence," Journal of Labor Economics, 30 (2012), 91–147.

DeVaro, Jed, and Oliver Gürtler, "Strategic Shirking in Promotion Tournaments," Journal of Law, Economics, and Organization, 32 (2015), 620-651.

Dixon, Matthew, and Brent Adamson, *The Challenger Sale: Taking Control of the Customer Conversation* (London: Penguin, 2011).

Ekinci, Emre, Antti Kauhanen, and Michael Waldman, "Bonuses and Promotion Tournaments: Theory and Evidence," The Economic Journal, (2018).

Ewens, Michael, and Matt Marx, "Founder Replacement and Startup Performance," The Review of Financial Studies, 31 (2018), 1532-1565.

Fairburn, James A., and James M. Malcomson, "Performance, Promotion, and the Peter Principle," The Review of Economic Studies, 68 (2001), 45–66.

Faria, João Ricardo, "An Economic Analysis of the Peter and Dilbert Principles," UTS Working Paper No. 101.

Fisman, Raymond, and Yongxiang Wang, "The Distortionary Effects of Incentives in Government: Evidence from China's Death Ceiling Program," American Economic Journal: Applied Economics, 9 (2017) 202–218.

Fisman, Raymond, Jing Shi, Yongxiang Wang, and Rong Xu, "Social Ties and Favoritism in Chinese Science," Journal of Political Economy, 126 (2017), 1134-1171.

Frederiksen, Anders, Fabian Lange, and Ben Kriechel, "Subjective performance evaluations and employee careers," Journal of Economic Behavior and Organization, 134 (2017), 408–429.

Grabner, Isabella, and Frank Moers, "Managers' Choices of Performance Measures in Promotion Decisions: An Analysis of Alternative Job Assignments," Journal of Accounting Research, 51 (2013), 1187–1220.

Heckman, James J., and Edward Vytlacil, "Structural Equations, Treatment Effects, and Econometric Policy Evaluation," Econometrica, 73 (2005), 669–738.

Hellmann, Thomas, and Manju Puri, "Venture Capital and the Professionalization of Start-up Firms: Empirical Evidence," The Journal of Finance, 57 (2002), 169–197.

Hoffman, Mitchell, and Steven Tadelis, "People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis," Working paper (2018).

Holland, Paul W, "Statistics and Causal Inference," Journal of the American Statistical Association, 81 (1986), 945–960.

Kaplan, Steven N., and Morten Sorensen, "Are CEOs Different? Characteristics of Top Managers," Working paper (2016).

Kaplan, Steven N., Mark M. Klebanov, and Morten Sorensen, "Which CEO Characteristics and Abilities Matter?" The Journal of Finance, 67 (2012), 973–1007.

Ke, Rongzhu, Jin Li, and Michael Powell, "Managing careers in organizations," Journal of Labor Economics, 36 (2018), 197–252.

Larkin, Ian, "Paying $30,000 for a Gold Star: An Empirical Investigation into the Value of Peer Recognition to Software Salespeople," Working paper, 2011.

Larkin, Ian, "The Cost of High-powered Incentives: Employee Gaming in Enterprise Software Sales," Journal of Labor Economics, 32 (2014), 199–227.

Larkin, Ian, Lamar Pierce, and Francesca Gino, "The Psychological Costs of Pay-for-performance: Implications for the Strategic Compensation of Employees," Strategic Management Journal, 33 (2012), 1194–1214.

Lazear, Edward P, "The Peter Principle: A Theory of Decline," The Journal of Political Economy, 112 (2004), S141–S163.

Lazear, Edward P., Kathryn L. Shaw, and Christopher T. Stanton, "Who Gets Hired? The Importance of Competition among Applicants," Journal of Labor Economics, 36 (2018), S133-181.

Malcomson, James M, "Work incentives, hierarchy, and internal labor markets," Journal of Political Economy, 92 (1984), 486–507.

Mas, Alexandre, and Enrico Moretti, "Peers at Work," American Economic Review, 99(2009), 112–45.

Milgrom, Paul R, "Employment contracts, influence activities, and efficient organization design," Journal of Political Economy, 96(1988), 42–60.

Miglrom, Paul R, and John Roberts, *Economics, Organization and Management* (Englewood Cliffs, NJ: Prentice Hall).

Oettl, Alexander, "Reconceptualizing Stars: Scientist Helpfulness and Peer Performance," Management Science, 58 (2012), 1122–1140.

Oyer, Paul, "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," The Quarterly Journal of Economics, 113 (1998), 149–185.

Peter, Laurence J., and Raymond Hull, *The Peter Principle* (New York: William Morrow & Co. 1969).

Prendergast, Canice, and Robert H. Topel, "Favoritism in Organizations," The Journal of Political Economy, 104 (1996), 958–978.

Rubin, Donald, "Estimating causal effects of treatments in randomized and nonrandomized studies" Journal of Educational Psychology, 66 (1974), 688–701.

Sales Management Association, "How Do Sales Managers Spend Their Time?" https://www.slideshare.net/TheSalesMgtAssoc/how-do-sales-managers-spend-their-time-524673.

Sarsons, Heather, "Gender Differences in Recognition for Group Work," Working paper, 2019.

Sevy, Bruce, "Why Great Salespeople Make Terrible Sales Managers." *Forbes*, 2016.

Waldman, Michael, "Job assignments, signalling, and efficiency," The RAND Journal of Economics, 15 (1984a), 255–267.

Waldman, Michael, "Worker allocation, hierarchies and the wage distribution," The Review of Economic Studies, 51 (1984b), 95–109.

Waldman, Michael, "Ex Ante Versus Ex Post Optimal Promotion Rules: The Case of Internal Promotion," Economic Inquiry, 41 (2003), 27–41.

Weber, Max, *The Theory of Social and Economic Organization* (New York: AM Henderson and Talcott Parsons, 1947).

TABLE I:
DESCRIPTIVE STATISTICS

Sample coverage		Probability of promotion	
Number firms	131	Within sample	0.0400
Number workers	38,843	Monthly hazard	0.0023
Number workers promoted to management	1,553		
Years covered	2005-2011		

Summary statistics	*Mean*	*25th*	*50th*	*75th*
Worker characteristics				
Monthly sales*	$3,206,029	$35,715	$286,427	$1,641,797
Number of collaborators*	6.0	1	1.9	4.8
Monthly commissions*	$14,615	$925	$3,814	$ 10,458
Salary	$7,217	$4,426	$7,117	$9,380
Manager characteristics				
Number of subordinates	5.4	2	4	8
Monthly commissions*	$15,458	$2,562	$7,047	$17,052
Change in monthly commissions	$1,121	-$1,444	$713	$6,119
Salary	$11,994	$8,501	$11,207	$13,862

NOTES: * denotes 12 month moving average. To compute these summary statistics, sales, commissions, and salary are deflated to January 2010 dollars using the Consumer Price Index for All Urban Consumers (CPI-U). Worker summary statistics are calculated using observations at the worker-month level. The exceptions are salary summary statistics which are calculated using observations at the worker level, because we observe salary as a snapshot as of the start of each worker's tenure within our sample. Salary data are also not available for all workers in our sample; we observe salary for 21,243 workers. Manager summary statistics cover 5,956 managers and include managers who were not promoted internally within our sample period (see Online Appendix Table A2 for details concerning manager sample coverage). Manager summary statistics are calculated using observations at the manager-month level. The exceptions are salary summary statistics which are calculated using observations at the manager level, because we observe salary as a snapshot as of the start of each manager's tenure within our sample. Manager salary data are also not available for all managers in our sample; we observe salary for 3,070 managers. Change in monthly commissions represents changes in pay after promotion, estimated as the average of monthly commissions in the 12 months after promotion minus the average of monthly commissions in the 12 months before promotion, and is estimated from the subsample of 1,553 managers for whom we observe pre-promotion data.

TABLE II:
PROBABILITY OF PROMOTION BY SALES PERFORMANCE

	Worker is promoted			
	(1)	(2)	(3)	(4)
Log(sales)	0.0941***	0.107***	0.0948***	0.0866***
	(0.00860)	(0.00873)	(0.00906)	(0.00901)
Jackknife firm-month promotion rate	28.49***			
	(2.879)			
Team sales rank			-0.0271***	-0.00373
			(0.00396)	(0.00411)
Top sales rank				0.659***
				(0.0605)
Pre-promotion controls	Yes	Yes	Yes	Yes
Firm-month FE	No	Yes	Yes	Yes
R-squared	0.013	0.051	0.051	0.052
Observations	205,390	206,255	206,255	206,255

NOTES: This table presents the regression described in equation (2). We use data at the worker-month level for workers that have not yet been promoted. The dependent variable is an indicator for whether a worker is promoted in the next month, multiplied by 100 so that estimates represent a percentage point increase in the probability of being promoted. Log sales is the log of one plus worker i's monthly sales credits, averaged over the past 12 months or for the worker's total tenure if tenure is fewer than 12 months. It is demeaned within firm-year-month in Column 1 and the other columns control for firm-year-month fixed effects. Team sales rank is the rank of the worker among others who share the same manager, based on sales performance averaged over the past 12 months. Top sales rank is an indicator for whether a worker is top ranked in sales among the sales team. Pre-promotion characteristics include controls for a worker's collaboration experience (log of one plus the average number of other collaborators worker i has per order, again averaged over the past 12 months or for the worker's total tenure if tenure is fewer than 12 months, as well as an indicator for having no such collaborations), seven bins of a worker's tenure, interacted with an indicator for whether tenure may be censored. Jackknife firm-year-month promotion rate is the fraction of workers promoted within worker i's firm in the same month, excluding worker i and worker i's teammates. Standard errors are clustered by worker. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

TABLE III:
MANAGER VALUE ADDED BY SALES PERFORMANCE

	Manager value added among all promotions		Manager value added among salespeople promoted to different team	
	(1)	(2)	(3)	(4)
Pre-promotion log(sales)	-0.0914*	-0.0878*	-0.108**	-0.105**
	(0.0469)	(0.0452)	(0.0523)	(0.0500)
Pre-promotion controls	No	Yes	Yes	Yes
R-squared	0.017	0.038	0.022	0.044
Observations	1,039	1,039	792	792

NOTES: This table presents the regression described in equation (3). We use data at the manager level. The sample is restricted to promoted managers for whom we can observe pre-promotion characteristics and for whom we can estimate manager value added fixed effects using movements of subordinates across managers. The dependent variable is manager value added, estimated as the change in subordinate performance associated with each manager (see equation (1)). Log sales is the log of one plus manager i's monthly sales credits as a worker, averaged over the 12 months prior to i's promotion (or for i's total pre-promotion tenure, if fewer than 12 months), and demeaned within firm-year-month. Even-numbered columns include controls for the manager's pre-promotion collaboration, and tenure in the month prior to promotion, as described in Table II. Columns 3 and 4 further restrict the sample to managers who are assigned to subordinates, none of whom were their previous teammates. Observations are weighted by the inverse variance of the manager value added measures. Standard errors are adjusted for heteroskedasticity. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

TABLE IV:
Manager Value Added of Marginally Promoted Workers by Ranking

	Manager value added by pre-promotion sales tercile			Manager value added by top ranking	
	Top tercile (1)	Middle tercile (2)	Bottom tercile (3)	Ranked #1 (4)	Not #1 (5)
Promoted	-0.503***	-0.0212	0.615*	-0.726**	-0.165
	(0.121)	(0.252)	(0.357)	(0.296)	(0.134)
Pre-promotion controls	Yes	Yes	Yes	Yes	Yes
Observations	68,220	68,291	68,369	43,020	161,860
P-value, test of joint equality			0.0035		0.0921

Notes: This table reports the estimated managerial quality of the marginally promoted worker for each specified group, as described by equation (6) in Section V.B. In Panel A, Column 1 presents estimates from a separate regression on workers with bottom tercile pre-promotion sales for their firm-year-month. Columns 2 and 3 do this for middle and top tercile workers, respectively. Panel B estimates the quality of marginally promoted managers who were top-ranked in sales within their teams versus those who were not top-ranked. Pre-promotion characteristics include controls for a worker's collaboration experience (log of one plus the average number of other collaborators worker i has per order, again averaged over the past 12 months or for the worker's total tenure if tenure is fewer than 12 months, as well as an indicator for having no such collaborations), seven bins of a worker's tenure, interacted with an indicator for whether tenure may be censored. Observations are weighted by the inverse variance of the manager value added measures. Standard errors are clustered by worker. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

TABLE V:
COLLABORATION EXPERIENCE

	Worker is promoted			Manager value added		
	(1)	(2)	(3)	(4)	(5)	(6)
Log(collaborators)	-0.185***		-0.102**			
	(0.0402)		(0.0500)			
Lone wolf		0.274***	0.192***			
		(0.0458)	(0.0570)			
Pre-promotion log(collaborators)				0.392*		0.323
				(0.216)		(0.234)
Pre-promotion lone wolf					-0.345*	-0.198
					(0.193)	(0.208)
Pre-promotion controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm-month FE	Yes	Yes	Yes			
R-squared	0.051	0.051	0.051	0.016	0.014	0.017
Observations	206,255	206,255	206,255	1,039	1,039	1,039

NOTES: This table examines how promotions and manager value added are related to a worker's collaboration experience. Collaboration experience is measured as the log of one plus the average number of other collaborators worker i has per order, averaged over the past 12 months or for the worker's total tenure if tenure is fewer than 12 months. Lone wolf is an indicator for having no such collaborations. Columns 1-3 use data at the worker-month level, and examine how the probability of promotion (in percent) relates to collaboration experience. Standard errors are clustered by worker. Columns 4-6 use data at the manager level and examine how the manager's value added relates to the manager's pre-promotion collaboration experience. Log(collaborators) is measured relative to the firm-year-month mean in the month preceding the promotion event. Pre-promotion controls include the worker's sales performance and tenure in the month preceding the promotion event (as defined in Table III). ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

TABLE VI:
Sales Performance and Workforce Management

	Initial teamsize (1)	Net change (2)	Fraction joining (3)	Fraction leaving (4)	Percentile of leavers (5)
Pre-promotion log(sales)	0.180***	-0.00236	-0.000815	-0.000821	-0.00504
	(0.0405)	(0.00554)	(0.000646)	(0.000688)	(0.00328)
Pre-promotion controls	Yes	Yes	Yes	Yes	Yes
R-squared	0.203	0.071	0.119	0.104	0.119
Observations	1,019	1,019	1,019	1,019	665

Notes: This table examines the relation between manager pre-promotion sales performance and manager performance on non-sales metrics. Observations are at the manager level. Pre-promotion sales and collaborators are as defined in Table III. Initial teamsize is the initial number of subordinates assigned to a manager. Net change is the overall change in teamsize each month. Fraction joining is the fraction of new team members joining each month. Fraction leaving is the fraction of team members exiting each month. Percentile of leavers is the sales percentile of departing workers in each month. All of these variables are measured relative to the firm-year-month mean in the full sample and averaged over a manager's tenure. Pre-promotion characteristics include controls for a worker's collaboration experience (log of one plus the average number of other collaborators worker i has per order, again averaged over the past 12 months or for the worker's total tenure if tenure is fewer than 12 months, as well as an indicator for having no such collaborations), seven bins of a worker's tenure, interacted with an indicator for whether tenure may be censored. Standard errors are adjusted for heteroskedasticity. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

TABLE VII:
HETEROGENEITY IN PROMOTION POLICIES ACROSS FIRMS

	Worker is promoted				
	(1)	(2)	(3)	(4)	(5)
Log(sales)	0.0925***	0.677***	0.706***	0.222***	0.280***
	(0.00765)	(0.0897)	(0.100)	(0.0740)	(0.0800)
Log(sales) * log(mean team size)		-0.270***	-0.277***		
		(0.0399)	(0.0448)		
Log(sales) * share variable pay				-0.368**	-0.471**
				(0.173)	(0.186)
Log(collaborators)			-0.272		-0.588**
			(0.346)		(0.249)
Log(collaborators) * log(mean team size)			0.0376		
			(0.159)		
Log(collaborators) * share variable pay					0.943
					(0.617)
Pre-promotion controls	Yes	Yes	Yes	Yes	Yes
Firm-month FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.051	0.051	0.051	0.044	0.044
Observations	206,255	205,838	205,838	154,338	154,338

NOTES: This table examines how promotion policies vary with the mean size of teams and the strength of pay for performance across firms. Log mean team size is the log of the average number of subordinates assigned to each manager within each firm-year. The share of variable pay represents the share of variable pay (commissions and bonus) as a share of total pay (which also includes salary), averaged across all workers within each firm-year. Observations are at the worker-month level. All other variables and sample restrictions are as described in Table II. Pre-promotion characteristics include seven bins of a worker's tenure, interacted with an indicator for whether tenure may be censored. The sample size declines in columns (4) and (5) because we of incomplete compensation data within our sample. Standard errors are clustered by worker. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

TABLE VIII:
Disincentives of Promoting Lower Performing Workers

	Worker is promoted			
	(1)	(2)	(3)	(4)
Log(sales)	0.0941***	0.107***	0.0948***	0.0866***
	(0.00860)	(0.00873)	(0.00906)	(0.00901)
Jackknife firm-month promotion rate	28.49***			
	(2.879)			
Team sales rank			-0.0271***	-0.00373
			(0.00396)	(0.00411)
Top sales rank				0.659***
				(0.0605)
Pre-promotion controls	Yes	Yes	Yes	Yes
Firm-month FE	No	Yes	Yes	Yes
R-squared	0.013	0.051	0.051	0.052
Observations	205,390	206,255	206,255	206,255

Notes: This table examines the subsequent retention and change in performance of workers who were teammates of the promoted person at the time of the promotion event. Observations are at the worker by promotion event level. Retention is a dummy variable for whether the worker remains in the sample (as either a worker or manager) in the 3, 6, or 12 months after the teammate was promoted. Relative sales in each month is the worker's level of sales, demeaned by the firm-year-month average. Change in sales equals the average of worker's relative sales in the 3, 6, or 12 months after the promotion event minus the worker's 12-month average of previous relative sales at the time of the promotion event. Outsold promoted is a dummy variable for whether the worker had a higher 12-month average of previous relative sales than the promoted teammate. Own percentile rank is the worker's sales percentile (measured from 0 to 1) within the team at the time of the promotion event (higher is better), measured using average relative sales over the previous 12 months. Own relative sales is the worker's average sales relative to the firm mean over the past 12 months at the time of the promotion event. Standard errors are clustered by worker. ***, **, * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

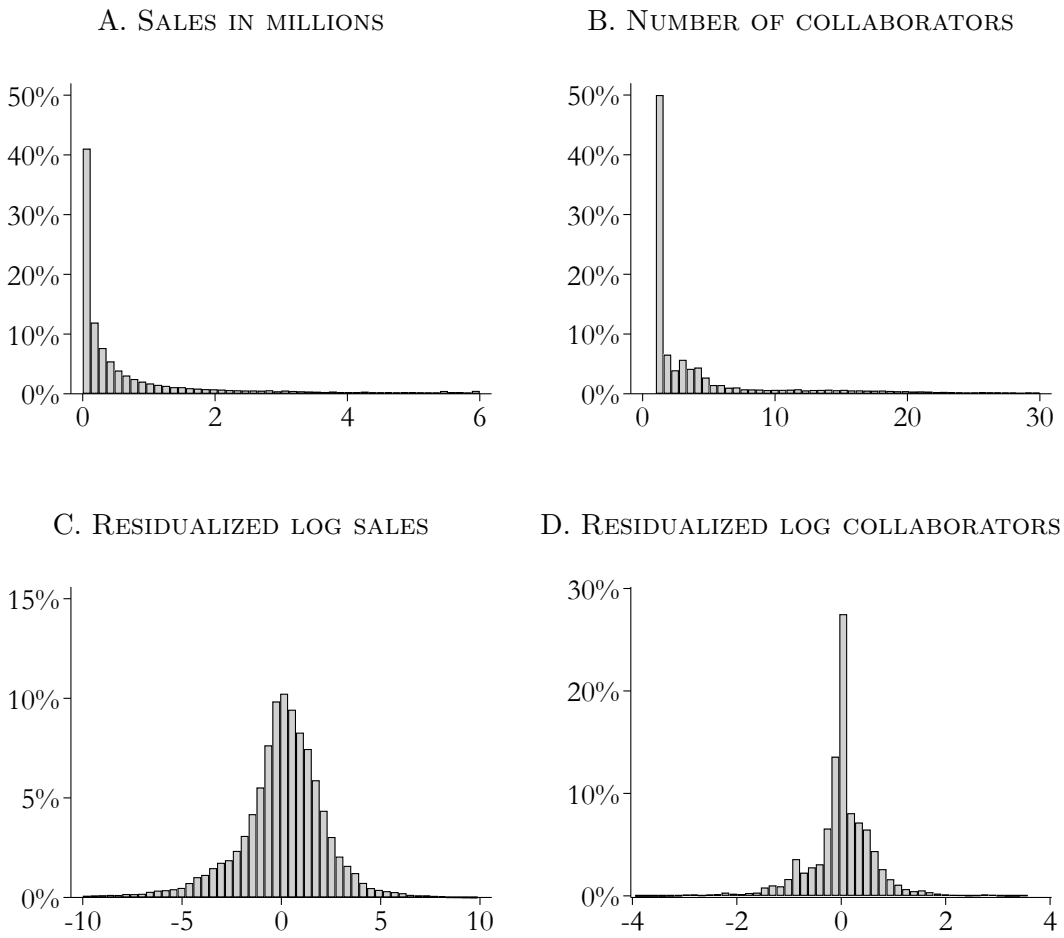

Figure I:
Distribution of Sales and Number of Collaborators

Panels A and C present the 12-month moving average of sales at the worker-month level, excluding those with zero sales over the past 12 months. Panels B and D do the same for the number of collaborators (including oneself). Panels A and B show the untransformed distributions. Panels C and D show the residuals after the log-transformed variables are regressed on firm-year-month fixed effects. Sales are deflated to January 2010 dollars using the Consumer Price Index for All Urban Consumers (CPI-U).

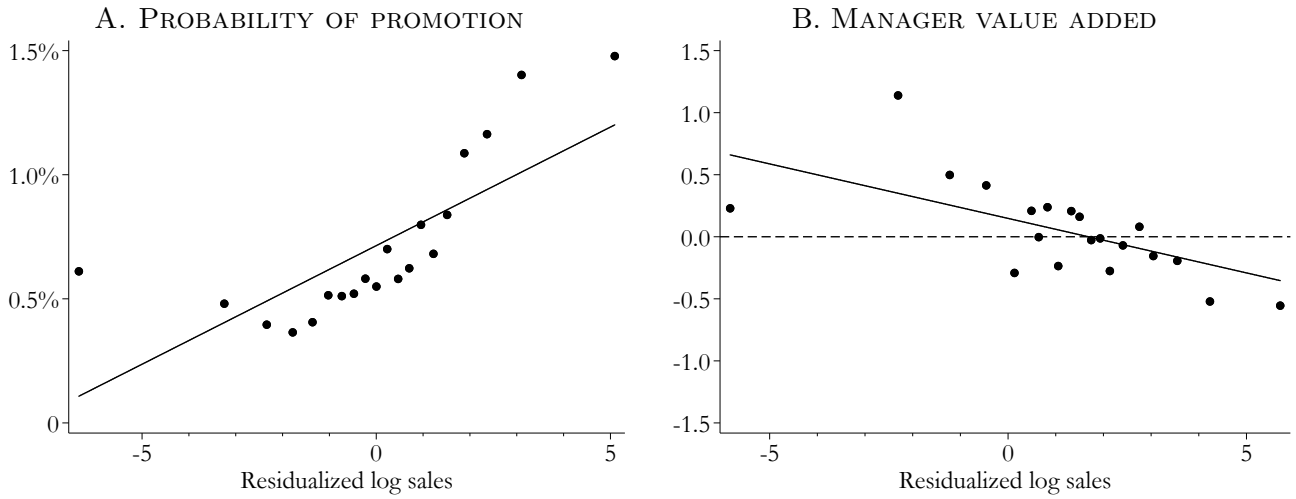

Figure II:
Correlates of Worker Sales Performance

Panel A shows a binned scatterplot relating worker sales and the monthly probability of promotion. Residualized log sales is the residual from a regression of the 12-month moving average of log pre-promotion sales on the following controls: the 12-month moving average of log pre-promotion number of collaborators, an indicator for having no collaborators, fixed effects for tenure bins, and firm by year-month fixed effects. Panel B plots the relation between the same residual pre-promotion sales performance variable and manager value added, weighted by the inverse variance of the estimated manager value added effect. These data are at the manager-level and include only promoted managers.

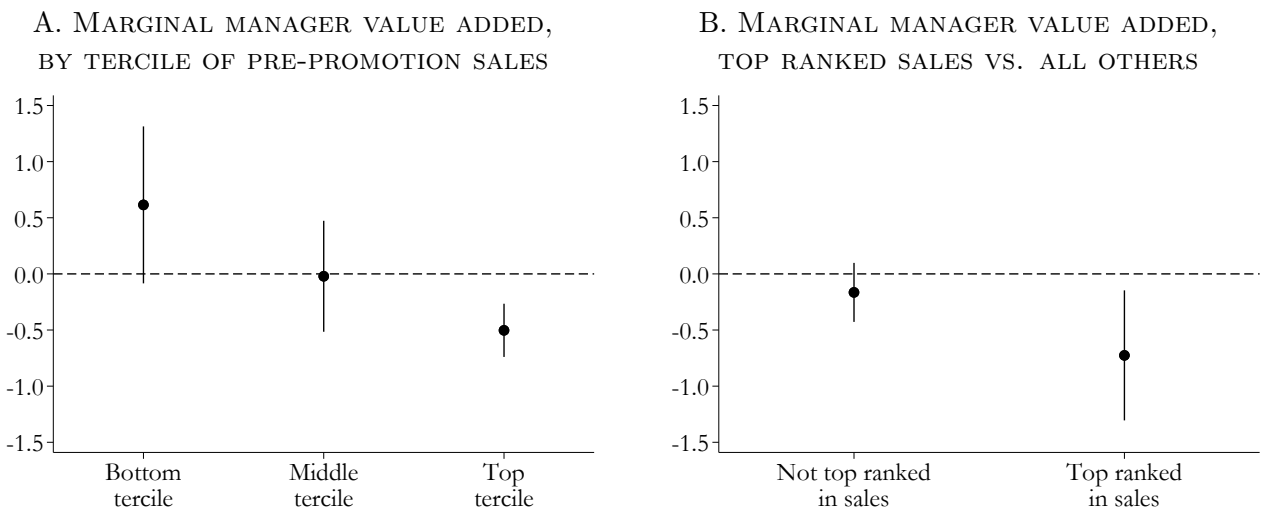

Figure III:
Manager Value Added for Marginally Promoted Workers, by Sales Performance

These figures plot the estimates from Table IV. Panel A plots the coefficient a_1 from equation (6) for each of three terciles of a worker's sales performance, instrumenting a worker's promotion status with the jackknife average promotion rate in each firm-year-month, weighted by the inverse variance of the estimated manager value added fixed effect. The coefficient can be interpreted as the manager value added of the marginally promoted manager, among workers with sales performance in each of three terciles. See Section V.B. for more discussion. Panel B plots the analogous graph of marginal managerial quality of promoted managers who were not the top-ranked sales person in their team versus the marginal quality of promoted managers who were. Bars represent 95% confidence intervals.

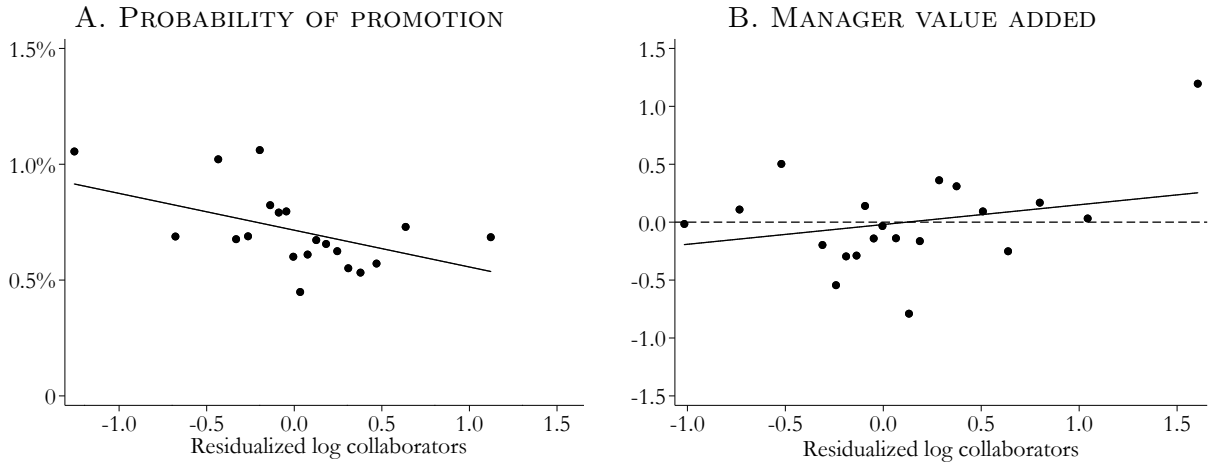

FIGURE IV:
COLLABORATION EXPERIENCE

These figures present analogues of Figures II and III by a worker's pre-promotion collaboration experience, rather than by pre-promotion sales performance. Panel A plots the monthly probability of promotion by collaboration experience, holding constant sales performance, fixed effects for tenure bins, and firm-year-month fixed effects. Panel B plots the relationship for manager value added and collaboration experience.

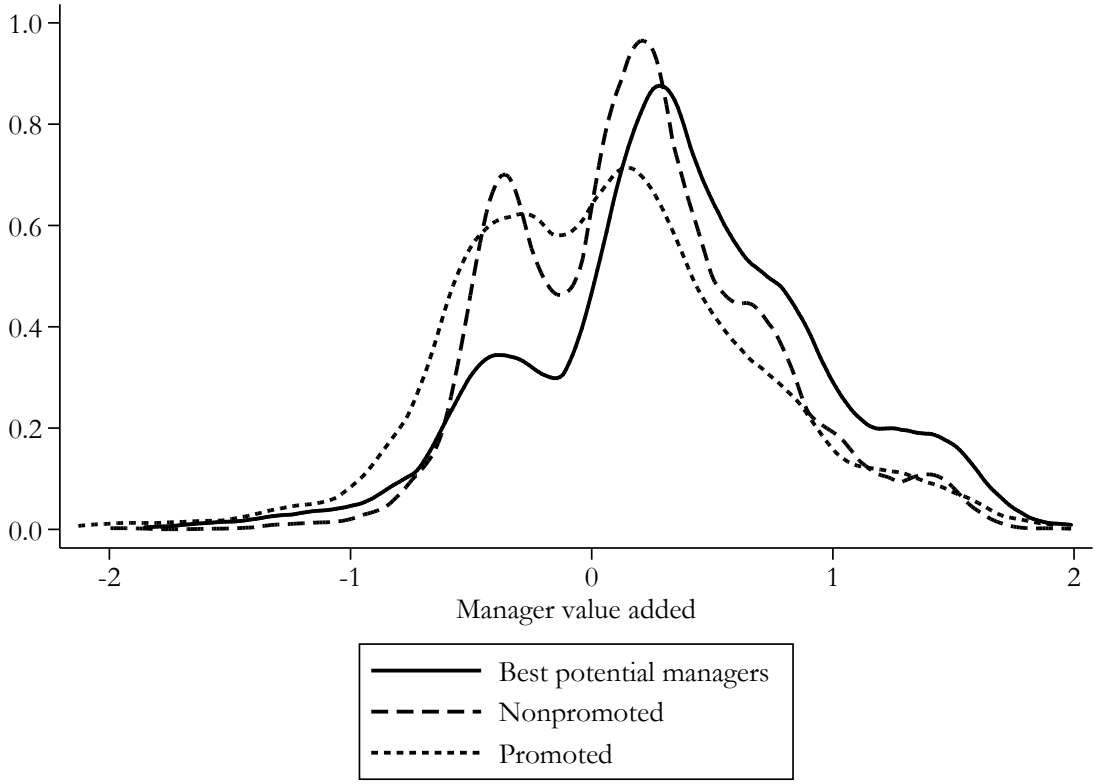

FIGURE V:
ACTUAL VS. OPTIMAL PROMOTION POLICIES

The figure shows the distributions (kernel densities) of predicted manager value added for the samples of workers that are promoted, workers that are not promoted, and workers that would make the best potential managers. All manager value added measures are predicted fitted values of manager value added based on estimating equations (10) and (11), with the modification that pre-promotion sales is measured as an equally spaced three-part spline to allow for potential non-linearities, and we instrument for selection into the promoted sample using the jackknife firm-year-month average promotion rates. To determine the best potential managers sample, we select the worker with the highest predicted manager value added within the same team and month when a worker is actually promoted. The nonpromoted sample consists of other, nonpromoted, workers in the same team and month when a worker is actually promoted.