Can Promotion Tournaments Produce Bad Managers? Evidence of the "Peter Principle"

Alan Benson, Danielle Li, and Kelly Shue*

October 31, 2015

Abstract

The best worker isn't always the best candidate for manager. In these cases, do firms promote the best potential manager or the best worker in their current job? Using data on the performance of sales workers from 214 firms, we find evidence consistent with the Peter Principle: firms prioritize current job performance when making promotion decisions, at the expense of other observable characteristics that better predict managerial quality. We estimate that the costs of managerial mismatch are substantial, suggesting that firms are either making inefficient promotion decisions or that the incentive benefits of emphasizing current performance must also be high.

*Respective affiliations are: University of Minnesota Carlson School of Management; Harvard Business School; and University of Chicago Booth School of Business and NBER. We thank Brian Phelan, Thomas Peeters, Lamar Pierce, and participants at the University of Minnesota Applied Microeconomics Workshop, Midwest Economics Association Meetings, NBER Summer Institute, and Wharton People Analytics Conference for helpful comments. We thank Menaka Hamplole and Leland Bybee for excellent research assistance. Correspondence: bensona@umn.edu.

Good workers do not always make good managers. When this is the case, who do firms promote: workers who excel in their current position, or those most likely to become excellent managers? If firms emphasize current performance, they may end up with worse managers. Yet if they prioritize managerial quality, workers may have weaker incentives to perform well in their current roles.

Using detailed microdata on sales workers, this paper provides the first large scale empirical evidence that firms systematically prioritize current performance in promotion decisions at the expense of optimizing managerial match quality. Our findings lend support for the "Peter Principle," which, in its extreme form, states that firms promote competent workers until they become incompetent managers (Peter and Hull, 1969). The data, provided by an anonymous firm that provides hosted sales performance management software, include standardized measures of sales transactions and pay at 214 US-based client firms across a range of industries. We use these data to track promotion decisions and performance measures from 2005 to 2011.

Because sales workers and their managers employ different skills, sales is a classic candidate for the Peter Principle (Fairburn and Malcomson 2001; Waldman 2003).[1] However, uinlike other classic settings for the Peter Principle, such as star engineers who are promoted to managers, sales workers and their managers have clear performance measures. For sales workers, we measure performance using the value of their credited sales, which the software uses to calculate commissions. For their managers, we measure performance using the change in their subordinates' performance after they fill the vacancy, analogous to value-added approaches applied in evaluations of managers and teachers.

Our analysis begins by showing that past sales performance is a very strong predictor of promotion. However among promoted workers, pre-promotion sales performance is in fact negatively correlated with managerial performance: when a managerial vacancy is filled by a work whose performance is at the 25th percentile instead of the 75th percentile, then after promotion, their subordinates' performance rises about 8% more. This is consistent with the Peter Principle. It is also consistent with the Beckerian insight that if firms are biased against poor sales performers

[1]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." 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."

who would nonetheless make good managers, then poor sales performers should outperform strong sales performers among observed promotions.

While these results are consistent with the Peter Principle, we face the empirical challenge that we only observe actual managerial performance among workers who are promoted. Ideally, we would observe the true latent managerial performance for all workers, regardless of whether they are actually promoted, and then examine whether firms promote the best candidates or are biased. However, firms do not promote at random, so the observed relationship between worker characteristics and managerial quality among promoted workers may not reflect the relationship for the worker population at large. For example, suppose that sales performance predicts managerial success and firms maximize match quality, so better sales performers tend to be promoted. Then poor sales performers who are nonetheless promoted in this scenario are likely to be ones with unobservably high managerial potential, and we may mistakenly conclude that the causal relationship between prior and the latent managerial performance is negative. Therefore, we want to use a method that distinguishes between this scenario and the alternative scenario that firms accept a lower acceptable threshold for managerial quality when evaluating star sales workers for promotion.

To do so, we develop a model of promotions based on the Heckman (1979) selection model with refinements by Chandra and Staiger (2011) and Abaluck and Agha (2015). In this model, firms form beliefs about a worker's managerial potential by observing the worker's characteristics, such as past sales performance and collaboration experience. Although this model allows us to measure whether firms promote the best potential managers, it does not impose any normative judgements regarding what firms should do. Firms may not promote the best potential managers because they are mistaken, they substitute match quality for tournament incentives, or because other criteria may be perceived as unfair.

Empirically, we identify whether firms promote to maximize managerial quality using a two step process. First, using the sample of all workers, we estimate each workers' propensity to be promoted as a function of their sales performance and other observable characteristics. Second, using the sample of promoted workers, we compare the performance of managers with the same propensity to be promoted. If firms overweight past sales performance in promotion decisions, then better

sales performance will be negatively correlated with subsequent managerial performance among two managers with an equal propensity to be promoted. The intuition behind this test is as follows: if firms are maximizing match quality, then it should only care about sales performance and other worker observables insofar as they affect managerial quality and should adjust a worker's propensity to be promoted accordingly. Controlling for promotion propensity, better salespeople should not make systematically worse managers. If this were the case, firms could improve managerial quality by placing less weight on sales performance in their promotion decisions.

We find that firm promotion decisions implicitly overweight the past sales performance of workers in their current positions. Holding the propensity to be promoted constant, a one standard deviation increase in pre-promotion sales is associated with a one-tenth standard deviation decline in managerial performance. Firms could improve managerial performance by placing more relative weight on worker sales collaboration experience. Overall, we estimate that firms could improve managerial performance by up to 50 percent if they adopted alternative promotion policies.

However, we caution that our results do not imply that firms use suboptimal promotion policies. Political or morale considerations may constrain firms' ability to promote on characteristics that are not the chief performance metric. Promotions may serve also an important motivational purpose to justify the significant costs of mismatch (Lazear and Rosen 1982, Prendergast 1998, DeVaro and Gurtner 2005). What our results do say is that costs of not promoting the best manager is high.

We also address two potential concerns with our analysis. First, Lazear (2004) argues that the Peter Principle could also be explained by mean reversion: high performing salespeople may indeed be the best predicted managers, but their performance would nonetheless be expected to decline following promotion. Our results, however, cannot be explained be explained by mean reversion because we measure managerial quality not as the manager's own sales performance (which may be mean-reverting), but as the ability of managers to improve the performance of their subordinates. Mean reversion is also inconsistent with a negative correlation between pre-promotion and post-promotion performance. Second, we show that our results are unlikely to be explained differences in the way promoted workers are matched to sales teams. Workers with strong previous sales performance are not matched to systematically better or worse teams of subordinates or to

teams with significantly different growth trends in sales. In addition, we observe similar patterns using direct manager sales as our measure of managerial performance instead of the change in subordinate sales.

This study offers the first empirical evidence of the magnitude of the Peter Principle using data on promotions across a large number of firms. Among workers in sales and technical occupations, the Peter Principle and the challenges it poses are well-known and provide the impetus for a range of policies, such as the use of dual career tracks for individual contributors and managers, and for the use of separate promotion and bonus criteria. Although theoretical work and reviews have hypothesized that tournament incentives may yield mismatch, scarce research has tested the Peter Principle directly. In this sense, our work is most closely related to Grabner and Moers (2013). They find that the bank that they study places less weight on on current job performance when a promotion would be to a job performing dissimilar tasks. However, this study features a single firm and does not attempt to estimate the cost of the Peter Principle. Finally, our findings are related to the large literature exploring the declining popularity of internal promotions and rising popularity of external CEOs or directors (e.g., Murphy and Zabojnik, 2004). The decision to hire an external manager must weigh the benefits of expanding the field of candidates to improve the quality of an eventual match against the costs of reducing tournament incentives for internal candidates.

1 Setting and Data

Data come from a firm that offers sales performance management (SPM) software over the cloud. Client firms input their employee records, organizational hierarchies, and sales transactions into the software calculates calculate pay and performance for each individual worker. Transaction inputs can be entered manually, but they are typically linked to customer relationship management (CRM) software. Pay outputs are typically linked directly to payroll software. The software is also used reporting and analysis. Salespeople and sales managers can log in to the SPM software's website to view their sales credits, progress toward quotas, commissions, and other data. It can also generate reports for use in auditing and compliance with Sarbanes-Oxley.

The data include 214 client firms, 48,209 sales workers, 734,278 worker-months, 6,308 sales

managers, and 20,929 manager-months. Client firms' most represented industries include manufacturing (62), information (56), and professional services (38). Because client firms pay a per capita license fee, client firms and their salespeople tend to be highly skilled and highly compensated, and should not be considered representative of a typical US sales worker. In 2011, sales occupations employed 13.6 million workers at a median wage of $24,840 (BLS 2011), about half of whom worked in retail sales. In contrast, sales workers in our data predominantly work in business-to-business sales and earn a median pay of $78,000, and the quartiles for variable pay are $0, $1,802, and $7,256 per month.

[Table 1]

Table 1 provides descriptive statistics for sample coverage. Because our sample firms are diverse and many are engaged in business to business transactions, worker performance varies widely. The median worker is credited on $122,632 USD worth of sales per month, with an interquartile range of $109 to $938,191. Reflecting the skew in the types of products that workers sell, the mean of this figure is $3.9 million. Panel 1 of Figure 1 plots the raw distribution of sales credits. Much of this variation may be explained by firm-level characteristics and seasonality. The second panel plots variation residual pay after accounting for firm by month fixed effects.

1.1 Measuring Sales Worker Performance

Sales workers are typically 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 sales industry refers to this process as the sales cycle. Furthermore, sales workers are typically assigned a market consisting of a territory, a set of products, or a type of client. This market determines the allocation of leads to salespeople. When a sale touches many people throughout the sales cycle, the allocation of credit to sales workers will depend on the organization's crediting rules. Splitting credit is common in business-to-business transactions, where an individual sale can involve multiple territories and products, and where transactions can be further split by the stage of the sales cycle (e.g. origination, execution, and renewals). Commission rates are then

specified at various intervals of the rolling sum of split credits. For example, commissions might be 0% on the first $1 million of sales, 2% on the next $1 million, and 3% thereafter. As such, these split credits are the fundamental unit of performance in the data; they are how compensable transactions are mapped onto performance, and also how the software determines commissions and bonuses.

Our primary measure of a sales worker's performance is the total dollar value of the worker's share of attributed credit on these transactions. Note that a successful sales worker (i.e. one with a high total sum of split credits) can generally achieve that status in two ways: by being the only person credited on a normal volume of transactions, or by receiving split credit on a larger volume of transactions. Indeed, much of the practitioner literature emphasizes the differences in performance management among these groups. "Lone wolves" might be recruited for their self-confidence, resilience, and autonomy, and are stereotypically marked by their reticence 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. For example, may successfully leverage others in the organization 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 typically entitle the salesperson to split credit.

Our data include 156 million credited transactions and how these credits are split among salespeople. We define a salesperson's performance to be their mean sum of split sales credits over the past twelve months. We define teamwork experience as the number of other workers with whom they shared split credit on any transaction in a given month. Other controls include a deal size variable measured as a Herfindahl index for that approaches one as all monthly credit comes from a single transaction, product variety measured by the count of distinct products credited to the salesperson, tenure in months, company fixed effects and month fixed effects.

[Table 1]

[Figure 1]

Figures 1 and 2 show the distributions of both sales performance and teamwork experience. Both are skewed. Table 2 provides descriptive statistics for the key pre-promotion characteristics of sales workers. It also serves as a reminder that sales performance is highly skewed. Indeed, salespeople note that the distribution of sales credit typically follows a "80-20" rule, in that 80% of sales are made by the top 20% of workers. This rule of thumb is similar to the skew in our data. Log transforming sales credits roughly yields a normal distribution.

1.2 Measuring Sales Manager Performance and Promotions

For each worker and manager in the data, we observe their position in the organizational hierarchy as well as their parent position over time. In the software, these hierarchies are used to allocate administrative privileges, to aid in reporting, to allow managers to monitor their subordinates, to measure managers' performance, and so-on. We define a promotion to be a change in position from one that does not have subordinates and into one that does; we do not count changes in job titles alone as promotions. Typical job titles for these positions include territory manager, sales director, regional director, regional manager, and regional vice president. We further restrict the sample of promotions to those for which we observe at least twelve months of performance data both before and after promotion. This process yields 4,871 observed promotions.

Using these data, we construct a value-added measure of managerial performance similar to those used to study the value of bosses given changes in the performance of their subordinates (e.g. Lazear, Shaw, and Stanton 2015) or teachers given changes in their students' scores. Specifically, to measure manager performance, we use the change in their subordinates' performance after the promoted manager assumes the vacancy. We then take the average difference in performance across all of that manager's subordinates to arrive at the value-added measure of managerial quality.

To justify this measure, it is important to note that sales managers are primarily responsible for building an effective sales team. Front-line sales managers surveyed by the Sales Management Association (2008) reported spending only about 20% of their time in selling activities, which tend to be in concert with subordinates (such as training, negotiating pricing, or handling disputes). Sales managers report spending the remainder of their time performing administrative duties, managing

subordinates' performance, allocating staff and resources, and coordinating across functions (such as marketing or product development). Sales managers require staffing skills so that they can hire, fire, and train a salesforce. They require leadership skills so that they can coach and motivate subordinates to do jobs that can require substantial resilience and autonomy. They require data analysis skills so that they can read market research, set quotas, assign territories, monitor performance, and prioritize sales activities. They require business acumen so that they can prepare budgets and work with other functions, such as product development, marketing, and operations. Successfully executing these activities reflects in the performance of their team. For example, if the manager misreads market research, manpower 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 features for their market.

2 What predicts managerial quality?

Our first empirical exercise examines the correlation between post-promotion managerial quality and pre-promotion worker characteristics.

$$\text{Managerial Quality}_{ift} = a_1 X_{if} + a_2 Sales_{if} + a_3 Team_{if} + \delta_{f\times t} + u_{if} \tag{1}$$

We run Equation (1) on a manager-month level panel among managers who were promoted as of month t. The sample includes both managers who have been promoted, for whom we observe pre-promotion characteristics, as well as those who have always been managers during our sample period. The coefficients of interest are a_2 and a_3, which describe the relationship between pre-promotion sales performance and teamwork experience and post-promotion managerial performance. Equation (1) also includes a number of additional variables X_{ift}: a worker's log commissions for the past month, an indicator for zero commissions, log of monthly bonus pay, an indicator for no bonus pay, information about their quota attainment, dummies for the size of their division, and their months of tenure. Finally, we also include company by month fixed effects, $\delta_{f\times t}$. Because of the

fixed effects, we estimate this equation using OLS, although we report probit specifications with fewer fixed effects.

[Table 3]

The key finding in Table 3 is that sales performance is negatively correlated with managerial quality, among the selected set of promoted workers. Could be optimal – for example, firms are proritizing

3 What predicts promotion?

Next, we consider how firms make promotion decisions. We first document the determinants of promotion among our sample sales workers:

$$\text{Promote}_{ift} = a_1 X_{ift} + a_2 Sales_{ift} + a_3 Team_{ift} + \delta_{f\times t} + u_{ift} \tag{2}$$

We run Equation (2) on a worker-month level panel, for workers i at firm f who have not yet been promoted as of month t. The dependent variable, Promote_{ift} is an indicator for whether a worker is promoted in the next month. Sales_{ift} is the log of worker i's monthly sales credits, averaged over the past 12 months or for the worker's total tenure if tenure is less than 12 months. Team_{ift} is defined analogously for the log of the unique number of sales collaborators in a month who have shared credits with worker i.

[Table 4]

Table 4 reports the results of this regression. We find that firms are more likely to promote higher performing salespeople, and this result is robust across specifications. The most simple model estimates that the number of sales collaborators is also correlated with promotion, although this is not robust to the full model.

To give an intuition of the magnitude in the full model, a worker with sales one standard deviation below the mean sales has a monthly promotion hazard of 0.15 percent, versus 0.36 percent for those with one standard deviation above.

Our results in the previous section show that firms favor high performing salespeople when making promotion decisions even though, conditional on promotion, sales performance is negatively correlated with managerial quality. This suggests that firms are placing more weight on current worker performance, at the expense of hiring the best managers.

A concern with interpreting our results in this way, however, is that the correlation between sales performance and managerial quality among promoted workers may not reflect the correlation between sales performance and managerial potential in the population of workers at large. For example, it may be the case that firms promote high–performing sales workers because sales performance is in fact positively correlated with managerial performance among their full population of sales workers. Among the selected sample of promoted workers, however, it may be the case that poor sales performers have high managerial potential on some other dimension which justifies their promotion. In this example, firms are using sales performance to accurately forecast managerial quality even though the correlation between sales performance and managerial quality among promoted workers is low or possibly even negative.

We next discuss a model of firm promotions and illustrate how it can be used to assess whether firms are indeed promoting the best potential managers.

4 Are firms promoting the best potential managers?

4.1 Theory

Firms may value sales performance in promotion decisions for three reasons: because workers with high sales performance make better managers, because promoting workers with higher sales performance creates incentives for all workers to exert effort on sales, or because firms mistakenly overweight the value of sales performance in predicting managerial performance. For brevity, we respectively refer to these as the matching, tournament, and biased belief motivations. In this section, we develop an empirical framework that allows us to decompose observed promotion patterns into a portion that is driven by matching, and another portion that is driven either by either tournaments or biased beliefs. We are unable to distinguish between the tournament or

biased belief motivations, but by separating out the extent to which firms do not promote the best expected managers (based on their ex ante information set), we provide a lower bound for the incentive benefits that firms must expect in order to rationalize this behavior.

Our approach to modeling promotion is based on classic selection models developed by Heckman (1979) and extended by Chandra and Staiger (2010) and Abaluck and Agha (2015) to study bias in medical treatment and diagnoses, and where the econometrician only observes the effect of the treatment on the treated.

In our version of this model, firms make promotion decisions for workers based on their performance in their current job. Empirically we will focus on two specific measures of worker level performance: sales performance (Sales_i) and team work experience (Team_i). We also observe other worker characteristics such as their pay, firm affiliation, and firm division (X_i), described in more detail in Section 4.2.

Suppose that the true relationship between these characteristics and a worker's latent managerial ability M_i is given by:

$$M_i = \beta_1 \text{Sales}_i + \beta_2 \text{Team}_i + X_i\beta_3 + \varepsilon_i \tag{3}$$

where ε_i represent the contribution of factors observable to the firm but not to the econometrician. Firms, however, may systematically misweight the importance of a worker's sales and team performance, so that a firm's beliefs about managerial quality are given by:

$$M_i' = (\beta_1 + \beta_1^e)\text{Sales}_i + (\beta_2 + \beta_2^e)\text{Team}_i + X_i\beta_3 + \varepsilon_i \tag{4}$$

where the $'$ superscripts denote errors in the firm's beliefs. In our model, we assume that firms have correct overall beliefs about the importance of other characteristics X_i, as well as other unobserved characteristics ε_i. We discuss the reasons for this assumption in Section 4.3.

Firms promote workers based on their beliefs about managerial quality M_i', but may apply different standards to different types of workers. Intuitively, firms prioritize certain pre-promotion characteristics and penalize others by raising or lowering the threshold of M_i' at which they would be willing to promote the worker. We formalize this by saying that firms promote if $M_i' > \tau_i$, where

the threshold τ_i may depend on a worker's performance:

$$\tau_i = \tau_0 + \tau_1 \text{Sales}_i + \tau_2 \text{Team}_i \tag{5}$$

Equation (5) is intentionally flexible: it allows a firm's promotion threshold to reflect or differ from the managerial quality forecasts in Equation (4). For example, a firm may observe that two workers have the same managerial potential M_i'. However, one worker expresses that managerial potential through higher sales, while the other expresses it through other observable characteristics. In order to provide incentives, the firm may adopt a lower promotion threshold for salespeople with higher sales, setting $\tau_1 < 0$, relative to the baseline threshold τ_0. In this case, high performing salespeople get promoted even when they are not likely to become a successful manager and vice versa—a practice consistent with trading match quality for tournament incentives. If the firm wants to incentivize teamwork, it may instead adopt $\tau_2 < 0$.

In this case, Equations (4) and (5) imply a worker's probability of promotion is:

$$\begin{aligned}
\Pr(\text{Promote}_i) &= \Pr(M_i > \tau_i) \\
&= \Pr\left((\beta_1 + \beta_1^e)\text{Sales}_i + (\beta_2 + \beta_2^e)\text{Team}_i + X_i\beta_3 + \varepsilon_i > \tau_0 + \tau_1 \text{Sales}_i + \tau_2 \text{Team}_i\right) \\
&= \Pr(I_i > -\varepsilon_i) \\
&\qquad \text{where } I_i \equiv (\beta_1 + \beta_1^e - \tau_1)\text{Sales}_i + (\beta_2 + \beta_2^e - \tau_2)\text{Team}_i + X_i\beta_3 - \tau_0
\end{aligned} \tag{6}$$

In Equation (6), I_i denotes worker i's propensity to be promoted, based on the factors that are observable to both the firm and the econometrician. Intuitively, a worker with high Sales_i is more likely to be promoted for two possibilities. The first possibility is that better salespeople tend to make better managers β_1 is high. In this case, firms promote high-performing salespeople in order to maximize match quality. Alternatively, firms may apply a lower promotion threshold for good sales workers, or have positively biased expectations regarding the managerial potential of high performing salespeople: $\beta_1^e - \tau_1$ is high. This would be consistent with cases in which firms forego managerial match quality to prioritize the chief pre-promotion performance metric tournament, as would be consistent with tournament theory and political models (e.g. Lazear and Rosen 1981,

Prendergast 1998, DeVaro and Gürtner 2005). Whether for tournaments or bias, we define this lower promotion threshold for high performing salespeople to be the hallmark of the Peter Principle.

To identify the sign of $\beta_1^e - \tau_1$, we examine the managerial performance among promoted workers:

$$\begin{aligned}
E(M_i|\text{Promote}) =& E(M_i|I_i > -\varepsilon_i) \\
=& \beta_1 \text{Sales}_i + \beta_2 \text{Team}_i + X\beta_3 + E(\varepsilon|I_i > -\varepsilon_i) \qquad (7) \\
=& (\beta_1 + \beta_1^e - \tau_1)\text{Sales}_i + (\beta_2 + \beta_2^e - \tau_2)\text{Team}_i + X_i\beta_3 - \tau_0 \\
&+ (\tau_1 - \beta_1^e)\text{Sales}_i + (\tau_2 - \beta_2^e)\text{Team}_i + E(\varepsilon_i|I_i > -\varepsilon_i) \\
=& I_i + (\tau_1 - \beta_1^e)\text{Sales}_i + (\tau_2 - \beta_2^e)\text{Team}_i + \tau_0 + E(\varepsilon|I_i > -\varepsilon) \\
=& \tau_0 + (\tau_1 - \beta_1^e)\text{Sales}_i + (\tau_2 - \beta_2^e)\text{Team}_i + g(I_i) \\
& \text{where } g(I_i) \equiv I_i + E(\varepsilon|I_i > -\varepsilon) \qquad (8)
\end{aligned}$$

To better understand Equation (8), consider a simple regression of managerial quality on sales performance, among promoted workers:

$$M_i = a_0 + a_1 \text{Sales}_i + e_i \qquad (9)$$

Suppose that the estimated coefficient a_1 is negative, meaning that, among promoted workers, better sales people make worse managers. fact may suggest that firms are applying a lower promotion threshold for high performing workers rather than simply trying to promote the best managers ($\tau_2 < 0$). However, another explanation is that firms apply the same promotion threshold to all workers ($\tau_2 = 0$), but high performing workers are less likely to make good managers ($\beta_2 < 0$) so that they on average fall closer to that threshold.

The key observation in Equation (8) is that we can separately estimate the promotion thresholds τ by controlling flexibly for a worker's propensity to be promoted. Specifically, the function $g(I_i)$ captures all the information that a firm has about how worker characteristics relate to true managerial quality. As such, we do not need to estimate the β coefficients specifically because they only impact the quality of managers through the propensity to promote. Put another way,

Equation (8) says that a firm's beliefs about how worker characteristics impact managerial quality is already "priced" into its promotion decisions.

Once we control for $g(I)$, Equation (8) compares the managerial quality of two promoted workers who have the same likelihood of being promoted. If the firm applies the same promotion threshold to high and low performing salespeople, then after controlling for the propensity to promote, higher performing salespeople should not be systematically better or worse managers. Any excess correlation indicates that the firm is either overweighting or underweighting sales performance, relative to the benchmark of simply trying to maximize managerial performance. We use Equation (8) to recover estimates of promotion thresholds and biases τ and β^b, and in turn use these to back out estimates of the managerial quality maximizing weights β.

4.2 Estimation Strategy

We recover estimates of a worker's promotion propensity I_i by estimating the a worker's probability of promotion, as described earlier in Equation (2), and construct fitted values of a worker's propensity to be promoted, I_i. Next, we control flexibly for promotion propensity, I_i, in a regression of realized managerial quality on pre-promotion characteristics:

$$\text{Managerial Quality}_{ift} = b_0 + b_1 \text{Sales}_{ift} + b_2 \text{Team}_{ift} + I_i + \cdots + I_i^5 + \delta_{f\times t} + e_i \tag{10}$$

Equation (10) is the regression analogue of Equation (6) from our model. We estimate this at the manager-month level for the sample of promoted managers. $\text{Managerial Quality}_{ift}$ is our measure of managerial quality, the change in a subordinate worker's sales performance under the current manager i, from that subordinate's performance under his or her previous manager, averaged over all of manager i's subordinates. Sales_{ift} and Team_{ift} are defined as pre-promotion sales and sales collaborations, averaged over the 12 months prior to a manger's promotion.

The coefficients b_1 and b_2 are the primary coefficients on interest, and correspond to estimates of $\tau_1 - \beta_1^e$ and $\tau_2 - \beta_2^e$ in Equation (6). The null hypothesis is that firms maximize managerial match quality, implying we should estimate $b_1 = b_2 = 0$. In other words, controlling for the propensity to promote, higher performing salespeople (or workers with more teamwork experience) should not

be systematically better or worse managers. Any excess correlation that gets captured by b_1 is indicative of a firm either over or underweighting sales performance relative to the benchmark of simply trying to maximize managerial performance. If we find $b_1 < 0$, this indicates that firms could improve managerial match quality by placing less weight on sales performance in their promotion decisions. The same logic applies for b_2 and teamwork experience.

4.3 Identifying Assumptions and Interpretation

As with standard Heckman selection models, the estimation of Equation (10) relies either on functional form restrictions on $g(I)$, or on the use of exclusion restrictions. If we include the same set of control variables in the first stage promotion equation (2) as in the second stage managerial quality regression (10), then the estimated propensity I_i from the first stage will be a linear combination of the covariates in the second stage regression. In this case, the coefficients b_1 and b_2 will be identified only from the functional form imposed on $g(I_i)$, in our case, the quintic polynomials in I_i. To avoid, this, we exclude the variables X_{ift} from entering the second stage equation except through $g(I_i)$. Because these variables are excluded, we can think of them as instrumental variables that aid the estimation of I_i, but which do not impact managerial quality M_i except through the promotion process described by I_i. This reflects a standard IV identification in Heckman selection models.

Substantively, this exclusion restriction allows for firms to consider commissions, bonuses, quota attainments, tenure, and division size when making promotion decisions, and requires them to correctly assess their impact on managerial quality. That is, we allow firms to misweight sales performance and teamwork experience in predicting managerial quality (β_i^e need not be zero), but we assume that firms do not misweight these other variables. This restriction is equivalent to including X_{ift} in the second stage controlling for $g(I_i)$, but imposing that the coefficient on these variables be equal to zero.

In practice, firms may also mis-weight a worker's wages when predicting managerial quality. In this case, other worker characteristics such as pay would separately enter Equation (10), but we would mistakenly omit these variables. To the extent that pay and sales performance are correlated,

our estimate of b_1 may reflect the firm's misweighting of pay in promotion decisions, rather than sales. We cannot simply include all X_{ift} variables because then $g(I_i)$ would only be identified from its functional form.

We view violations of this exclusion restriction as changing the interpretation of our results, rather than the validity of the exercise. Under the null that firms are, on average, accurately promoting the best managers, all variables, including our sales and team experience measures would be excluded from the second stage. Finding either $b_1 \neq 0$ or $b_2 \neq 0$ indicates that firms are not maximizing managerial match quality, and that they instead prioritize something correlated with sales performance or team experience. In our counterfactual simulations we can still examine the magnitude of the match quality loss associated with this promotion behavior.

In addition to this exclusion restriction, our identification strategy also requires that $E(\varepsilon_i|I_i > -\varepsilon_i)$ can be expressed as a function of I_i only. This is equivalent to assuming that the distribution of unobserved worker characteristics that determine promotion does not differ by sales performance or teamwork experience. For example, suppose firms groom high performers for management positions, and in doing so, they receive more precise information about ε_i (that is observable to the firm but unobserved to us). In this case, the variation in ε_i will be greater for high performers, and the truncated mean $E(\varepsilon_i|I_i > -\varepsilon_i)$ would be larger for any given truncation point, resulting in a higher expected managerial quality for good salespeople as firms screened upon their new information. Therefore, the example of superior screening among high sales performers would bias against the test of the Peter Principle. Alternatively, suppose that firms put less effort into determining managerial quality for better salespeople; in this case, $E(\varepsilon_i|I_i > -\varepsilon_i)$ may be smaller, leading to a lower expected managerial quality, given the same truncation point. This would appear as if firms were favoring good salespeople even if they were not (except through putting less effort into determining managerial quality).

4.4 Results

We begin by estimating Equation (2), which gives the probability of promotion based on observable characteristics.

[Table 3]

Table 3 shows that firms are more likely to promote higher performing salespeople, and this result is robust across specifications. The most simple model estimates that the number of sales collaborators is also correlated with promotion, although this is not robust to the full model.

To give an intuition of the magnitude in the full model, a worker with sales one standard deviation below the mean sales has a monthly promotion hazard of 0.15 percent, versus 0.36 percent for those with one standard deviation above.

Next, we examine the predictors of managerial performance by estimating Equation (10) without the $g(I)$ terms. This relates pre-promotion characteristics to post-promotion managerial performance, but does not yet address selection.

[Table 4]

Table 4 shows that, among observed promotions, there is a *negative* correlation between the pre-promotion sales performance and the subsequent change in performance among that manager's subordinates. This is also consistent with Chandra and Staiger's (2015) intuition as it applies to sales promotion: if weaker salespeople are being "discriminated" against in promotion decisions, then observed promotions among non-top salespeople should be more successful than observed promotions among top salespeople since they had to cross a higher promotion threshold.

To give an intuition of the magnitude, we cut the sample of observed promotions into two sub-samples based on pre-promotion sales performance. Subordinates working under salespeople who were in the top half of salespeople pre-promotion improved their performance by a median of 2.4 percent, versus 10.1 percent for those working for a manager in the bottom half. Given that the median manager has five credited subordinates, this is slightly less than half an average subordianate.

Table 5 presents the main test, presented in Equation (10).

[Table 5]

Table 5 shows that, after controlling for promotion propensity, sales performance declines after the firm promotes a higher performing salesperson. In contrast, controlling for promotion

propensity, sales performance improves after hiring a salesperson with a greater number of collaborators. These results are robust to controling for company and month effects.

To give an intuition of the magnitude in the full model, a one standard deviation increase in pre-promotion sales is associated with a 0.1 standard deviation decline in the mean subordinate's performance.

On net, results imply that firms are weighting pre-promotion sales performance moreso than would seem to be justified by its ability to predict future managerial performance. While firms tend to promote these top salespeople, the performance of a manager's subordinates rises most sharply when the firm departs from top salespeople and instead promotes a worker who was credited on greater sales prior to promotion. These results are consistent with managers lowering the promotion threshold for high-performing salespeople or expressing biased beliefs regarding the relationship between pre-promotion sales and post-performance managerial performance.

The theoretical and empirical setup of this paper places special emphasis on pre-sales performance. This is important because it is the measure that firms presumably wish salespeople to prioritize: it generates revenue and is used to calculate commsisions and bonuses. The other variables, such as team crediting and the controls, are used for intermediary calculations. Although we confirm that sales performance is a strong predictor of both promotions and (negatively) managerial performance, there may be other reasons why our estimates of b_1 are negative. First, managers may lower the thresholds or be biased toward other factors that are correlated with pre-sales performance. In that case, firms aren't prioritizing pre-promotion sales performance *per se*, but rather a correlated statistic.

Second, mean reversion could be a concern. Our correlations between pre-promotion and post-promotion performance are negative, so that the mean reversion hypothesis raised by Lazear (2004) is not a concern. However, firms may promote better-performing salespeople into higher-performing teams, and then these higher-performing teams may revert to their means. To examine this, we reproduce the main results, breaking subordinates' performance into quarters. Coefficient estimates remain stationary; we find no evidence that short-term performance blips shortly before promotions, which would put these teams at risk for mean reversion.

4.5 What are the producitivity lossses from mismatching?

How much do suboptimal promotions cost firms? To analyze this, we set aside the potential value of tournament incentives and the constraints placed on firms by imperfect monitoring and politics to focus instead on foregone match quality alone. Our estimates may be interpreted as match quality that firms forego to achieve tournament incentives or avoid constraints.

Specifically, we examine how predicted managerial quality differs among three categories of workers: (1) actual promoted salespeople, (2) actual nonpromoted salespeople, and (3) the top predicted manager among subordinates under the same manager of observed promoted workers. We interpret the last case to be the optimal promotion under the restriction that mobility and other frictions prevent the frim from promoting among the entire organization, and rather, firms must promote among the peers (under the same manager) of the actual observed promotions. We respectively denote these as $\hat{M}_i^1$, $\hat{M}_i^2$, and $\hat{M}_i^3$.

[Figure 1]

Figure 1 shows the distribution of $\hat{M}_i^1$, $\hat{M}_i^2$, and $\hat{M}_i^3$. To provide an intuition of the magnitudes, the median predicted improvement in subordinate salesperson performance is 20% in both the promoted and nonpromoted salespeople. In contrast, the median predicted improvement under the optimal matching rule is 40%. To give a further intuition, given that the mean number of subordinates in the sample is 4.6, we estimate that observed promotions underperform optimal promotions by the equivalent of about one additional salesperson.

5 Conclusions

Sales are an archetypal setting for the Peter Principle. This reputation owes to two perceptions: frontline and managerial sales duties sharply differ, and organizations tend to promote top salespeople despite this disconnect. However, as Lazear notes, the Peter Principle may be an unnecessary theory; the tendency for firms to promote great salespeople who go on to become merely-good sales managers may be because organizations are making their best guesses for who will make a good manager. While theoretically unnecessary, this paper suggests that the Peter Principle

is nonetheless supported by evidence that, in making promotion decisions, organizations weigh pre-promotion sales performance more than would be the case if they were trying to maximize match quality. As a result, the performance of a managerial position's subordinates decline after that slot is filled by someone who as a strong salesperson, but typically worked alone, prior to promotion. This is consistent with both the stereotype that sales is a classic scenario for the Peter Principle.

The salespeople we examine in these data reflect a large, highly skilled, high earning segment of the US labor force. Given that firms already directly incentivize salespeople through commissions and bonuses, it is a puzzle why such firms would be willing to incur the costs of mismatch. This sets salespeople apart from other high-paying occupations, such as scientists and engineers, in which pay is largely distributed by a salary that depends on one's rank within either managerial or technical career tracks.

We caution against interpreting these results as evidence that firms are behaving inefficiently. Rather, firms may use total sales credits in promotion decisions because, consistent with tournament theory, this is what it wants its frontline salespeople to prioritize. As such, firms may be substituting between the use of promotions as an incentive and their use in matching. Alternatively, the availability of relatively clear measures of the aptitude of frontline salespeoples' aptitude for their own job may lead organizations to emphasize these in promotion decisions. Regardless, these results lend evidence that firms do not promote entirely to maximize match quality.

6 Bibliography- INCOMPLETE

Abaluck, J., Agha, L., 2013. Negative Tests and the Efficiency of Medical Care: Investigating the Determinants of Imaging Overuse.

Baker, G., Gibbs, M., Holmstrom, B., 1994. The wage policy of a firm. The Quarterly Journal of Economics, 921-955.

Baker, G.P., Jensen, M.C., Murphy, K.J., 1988. Compensation and incentives: Practice vs. theory. The journal of Finance 43, 593-616.

Bureau of Labor Statistics. 2011. Occupational Employment Statistics. Available online at http://www.bls.gov/oes/2011/may/oes410000.htm.

Chandra, A., Staiger, D.O., 2010. Identifying provider prejudice in healthcare. National Bureau of Economic Research.

Connelly, B.L., Tihanyi, L., Crook, T.R., Gangloff, K.A., 2014. Tournament Theory Thirty Years of Contests and Competitions. Journal of Management 40, 1647.

DeVaro, J., 2006. Internal promotion competitions in firms. The Rand Journal of Economics 37, 521542.

Fairburn, J.A., Malcomson, J.M., 2001. Performance, promotion, and the Peter Principle. The Review of Economic Studies 68, 45-66.

Lazear, E.P., 2004. The Peter Principle: A Theory of Decline. Journal of Political Economy 112, S141S163. doi:10.1086/379943

Lazear, E.P., Rosen, S., 1981. Rank-Order Tournaments as Optimum Incentive Contracts. Journal of Political Economy.

Medoff, J.L., Abraham, K.G., 1980. Experience, performance, and earnings. National Bureau of Economic Research Cambridge, Mass., USA.

Peter, L.J., Hull, R., Frey, L., 1969. The peter principle. W. Morrow.

Sales Management Association. 2008. A Survey of How Sales Managers Use Their Tim. Available online at http://www.slideshare.net/TheSalesMgtAssoc/how-do-sales-managers-spend-their-time-524673

Waldman, M., 2003. Ex ante versus ex post optimal promotion rules: The case of internal promotion. Economic Inquiry 41, 27-41.

Table 1. Descriptive statistics

Sample coverage	
Number of firms	214
Number of workers	48,209
Number of managers	6,308
Number of promotions	1,565
Number of worker-months	734,278
Number of manager-months	20,929
Years covered	2005-2011
Pr(Promotion)	
Overall	0.0325
Monthly hazard	0.0021

Table 2. Pre-promotion characteristics and post-promotion characteristics

	Mean	25th	50th	75th
Pre-promotion characteristics				
Monthly credits	$3,942,247	$109	$122,632	$938,191
Monthly commissions	$12,367	$0	$1,802	$7,256
Monthly bonuses	$687	$0	$0	$0
Active collaborators	19.7	1	1	9
Post-promotion characteristics				
Monthly credits	$11,100,000	$64,297	$956,912	$5,214,855
Monthly commissions	$16,595	$265	$3,923	$13,653
Monthly bonuses	$1,296	$0	$0	$0
Number of subordinates	4.6	2	4	6
Mean of subordinates' credits	$2,217,852	$55,203	$277,335	$1,039,032
Mean change in subordinates credits	$1,424,365	-$169,466	$2	$323,061

Table 3. OLS regression of managerial performance

	(1)	(2)
	Change in Log(Subordinate Credits) | Promotion	
Manager's Pre-promotion Characteristics		
Log(Sales Credits)	-0.0529*** *(0.00827)*	-0.0487** *(0.0237)*
Log(# Sales Collaborators)	0.0552*** *(0.0197)*	0.111** *(0.0567)*
# Manager-Months	20,929	20,929
R^2	0.004	0.330
Company X Month FEs		X

Table 4. Linear probability model of whether worker is promoted in a given month

	(1)	(2)	(3)	(4)
	Worker is Promoted			
Log(Sales Credits)	0.0127***	0.0122***	0.0237***	0.0222***
	(0.00138)	*(0.00155)*	*(0.00219)*	*(0.00232)*
Log(# Sales Collaborators)	0.0157***	0.0116***	0.0153***	0.000116
	(0.00373)	*(0.00391)*	*(0.00592)*	*(0.00594)*
Tenure	-0.00507***		-0.000706	
	(0.000501)		*(0.000734)*	
# Worker-Months	734,278	698,965	734,278	698,965
R^2	0.000	0.001	0.056	0.057
Additional Controls		X		X
Company X Month FEs			X	X

Table 5. OLS regression of manager performance, with promotion propensities

	(1)	(2)	(3)	(4)
	Change in Log(Subordinate Credits) | Promotion			
Manager's Pre-promotion Characteristics				
τ_1: Log(Sales Credits)	-0.0384*** *(0.00835)*	-0.0405*** *(0.00836)*	-0.0431** *(0.0216)*	-0.0509** *(0.0227)*
τ_2: Log(# Sales Collaborators)		0.0704*** *(0.0214)*		0.117** *(0.0581)*
Manager's Promotion Propensity				
I	43.88 (49.06)	53.42 (49.76)	176.7*** (68.18)	184.6*** (68.63)
I^2, I^3, I^4, I^5 included but not reported				
# Manager-Months	20,929	20,929	20,929	20,929
R^2	0.014	0.014	0.339	0.341
Company X Month FEs			X	X

FIGURE 1—Distribution of Sales Credits

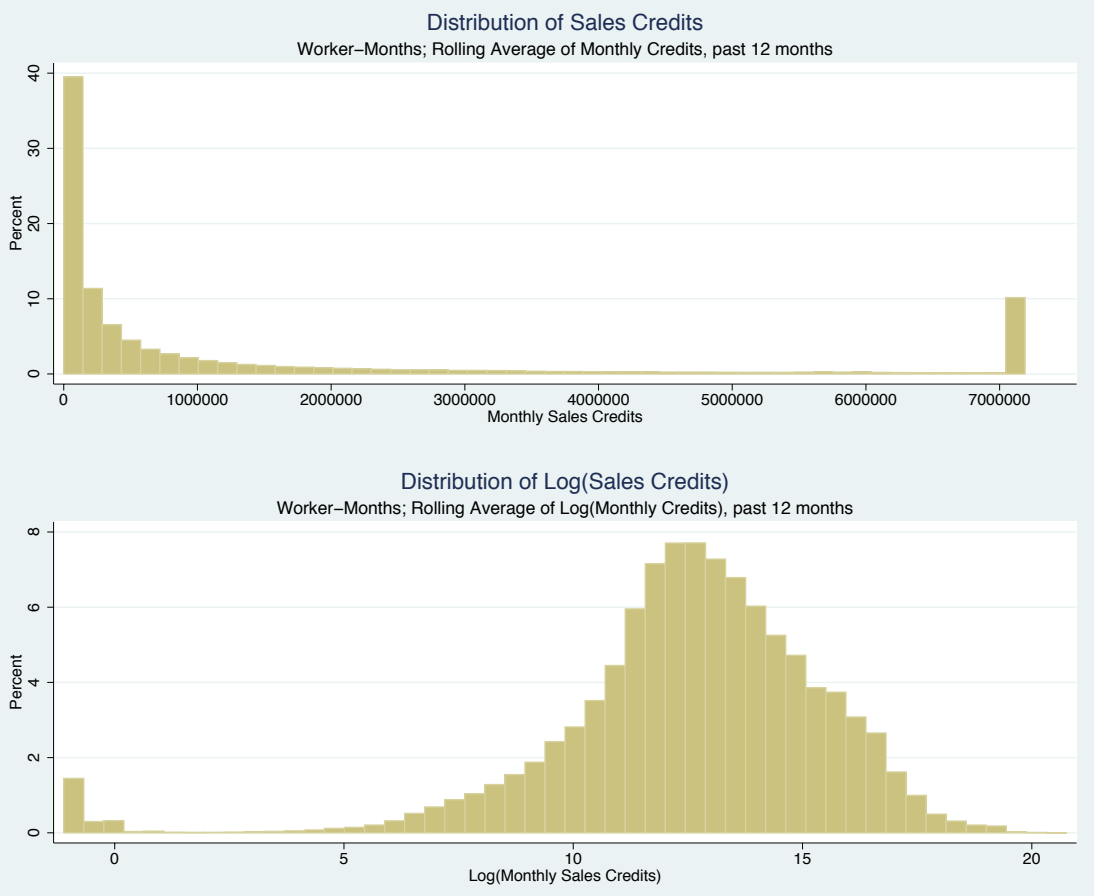

FIGURE 2— Kernel density of the predicted post-promotion change in subordinates' sales