HARVARD | BUSINESS | SCHOOL

The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales

Ian Larkin

Working Paper

13-073

February 20, 2013

The Cost of High-Powered Incentives: Employee Gaming in Enterprise Software Sales

Ian Larkin
Harvard Business School

February 20, 2013

Forthcoming in the *Journal of Labor Economics*
(tentatively scheduled for publication in Vol. 32, No. 2, 2014)

This paper investigates the pricing distortions that arise from the use of a common non-linear incentive scheme at a leading enterprise software vendor. The empirical results demonstrate that salespeople are adept at gaming the timing of deal closure to take advantage of the vendor's accelerating commission scheme. Specifically, salespeople agree to significantly lower pricing in quarters where they have a financial incentive to close a deal, resulting in mispricing that costs the vendor 6-8% of revenue. Robustness checks demonstrate that price discrimination by the vendor does not explain the identified effects.

I am grateful to Paul Oyer (the Editor), three anonymous referees, George Baker, Max Bazerman, Bob Gibbons, Bronwyn Hall, Larry Katz, David Mowery, Al Roth, Jason Snyder, Steve Tadelis, Catherine Wolfram, Oliver Williamson and numerous seminar participants. Desmond Ang and Matthew Chao provided excellent research assistance. All mistakes are my own. I thank the Ford Foundation and Harry S. Truman Foundation for providing research funding, and the enterprise software vendor for providing the dataset and for many useful discussions.
Address for correspondence: Baker Library 463, Soldiers Field Road, Boston MA 02613
email: ilarkin@hbs.edu

1 Introduction

It is well-known that employees "game" incentive systems by taking actions that increase their pay but hurt the objectives of their employer. Some gaming behavior, such as falsifying invoices or backdating stock option grants, is clearly illegal or unethical. However, employees often manipulate incentive systems in legal ways where the ethics of the manipulation are at worst questionable. One prominent example is the gaming of non-linear, period-based compensation systems, which give employees the incentive to alter the timing of tasks to take advantage of the non-linearity in pay. This "timing gaming" (Oyer, 1998) can affect business outcomes by changing revenue flow, pricing, or other key variables in ways beneficial to the employee but detrimental to the firm. While the prevalence of "timing gaming" in the face of non-linear, period-based incentives is well-known, little is known about the scope or cost of this gaming to employers. Existing empirical studies on "timing gaming" focus on macro-level company data (Oyer, 1998), illegal employee actions (Healy, 1985), or non-business settings (Asch, 1990).

In this paper, I use a proprietary database of deals for a leading enterprise software vendor ("the Vendor") to investigate the extent of "timing gaming" by the Vendor's salespeople, and to empirically estimate one significant cost arising from this gaming: pricing lower than required to make a sale in order to earn a commission on the timeframe most favorable to the salesperson. I observe full information about the Vendor's commission structure, as well as detailed information on the deal, customer and salesperson. Like many business-to-business (B2B) technology companies, the Vendor compensated salespeople based on an accelerating commission scale, meaning a salesperson's commission for the same deal could vary dramatically depending on the quarter in which the deal closes. The sales cycle in enterprise software is usually well over a year, and gives salespeople considerable control over the exact quarter in which a deal closes. Once a customer purchases licenses for enterprise software, it typically spends a year or more customizing and integrating the software into its other information technology systems. Therefore, as opposed to salespeople, customers often do not have strong preferences over the exact timing of its software purchases. Since salespeople have some control over the level of discount granted to a customer, they have the ability to give the customer a strong incentive – a lower price – to purchase in the quarter preferred by the salesperson.

Economists have long been puzzled by the use of non-linear incentives like the one used by the Vendor (Lazear and Oyer, 2012). Recent research in behavioral economics suggests that these schemes may be due to behavioral biases of salespeople. For example, the experimental study in Larkin and Leider (2012) shows that convex piece-rate systems attract and retain highly overconfident subjects, which may be beneficial in the sales function. A convex system may also have other sorting benefits, or lead to increased effort provision.

The paper's empirical results suggest that gaming is widespread and costly to the Vendor, specifically because salespeople with strong financial incentives to close a deal in a given quarter appear to grant discounts that are larger than necessary to win the deal in order to guarantee that the deal closes on the salesperson's preferred timeline. I first demonstrate that salesperson compensation concerns appear to influence the timing of the large majority of deals in the database, leading to a considerable bunching of deals within quarters to take advantage of the accelerating commissions. I next demonstrate that salespeople are much more likely to agree to lower prices when they have the financial incentive to close a deal. This appears to lead to both "pulling" behavior, where a salesperson closes a deal very late in a quarter when waiting until the next quarter would significantly reduce his[1] commission, and "pushing" behavior, where a salesperson closes a deal very early in a quarter when closing it the quarter before would have reduced his commission.

Since there were no exogenous changes to the incentive system off of which to identify causality, I cannot definitively prove that "timing gaming" causes this large pricing effect. The strongest empirical concern is around endogeneity, which may specifically arise if customer utility for the product, an unobserved variable, is negatively correlated with a salesperson's financial incentive to close a deal. A correlation between salesperson incentives and lower prices could mean that customers with low utility only choose to purchase when salespeople have the incentive to cut price, and that the incentive system therefore acts as an implicit price discrimination mechanism.

I carry out a series of analyses that suggest that price discrimination does not appear to explain a significant portion of the estimated negative correlation between salesperson compensation concerns and pricing. For example, variance in historical price paid for a product, which is a proxy for heterogeneity in customer valuation, is not correlated with apparent gaming behavior; salespeople are equally driven by financial concerns to cut price even for products where there has been less historical price variation. Similar tests using alternative proxies for customer valuation suggest low value customers are not more likely to purchase when a salesperson has the

[1] Over 90% of salespeople in this study are male.

incentive to cut price. The preponderance of the evidence therefore suggests salesperson compensation drives the negative correlation between salesperson timing preferences and pricing.

The regression estimates suggest the Vendor foregoes 6-8% of revenue through lower-than-needed pricing. That the paper identifies one large cost of the incentive system does not mean the system is suboptimal, given the potential benefits of the systems discussed above.

This paper makes several contributions to the literature. The shifting of sales across periods has been demonstrated previously, but it has not been demonstrated that these shifts are costly to the firm. This paper is the first to identify on a micro-level and quantify the "gaming costs" of a non-linear incentive scheme in terms of foregone revenue for the firm. Indeed, while many theoretical and empirical studies assume that higher employee wages are the primary cost of performance-based incentive systems (e.g., Holmstrom, 1979), this study finds an entirely different but conceivably larger cost factor. Furthermore, the "pushing" behavior empirically documented in the paper, where salespeople artificially delay the closing of a deal, has not been previously demonstrated in the literature, and is less well-known than the traditional "pulling" or deadline effect. In sum, the paper furthers our knowledge of non-linear incentive schemes, around which scholars have made "remarkably little progress" in understanding (Prendergast, 1999).

The study also contributes to the marketing literature, which has examined the use of commissions and other output-based incentives for salesforce motivation, finding a strong agency-based rationale for these practices (e.g., Basu et al, 1985; Lal and Srinivasan, 1993; Shaw et. al, 2000). Finally, a related stream of research examines non-linear compensation structures, particularly related to deadlines, such as Asch (1990), Leventis (1997), Chevalier and Ellison (1997), and Healy (1985).

The paper is laid out as follows. In the next section, I briefly introduce incentive and sales dynamics in enterprise software and other B2B technology industries. Section three overviews the data used in the study. Section four presents the empirical results, including tests of the alternative explanation around price discrimination. The paper concludes with a discussion of the results, a review of the limitations of the study, and a discussion of potential avenues for further research.

2 Sales Compensation at the Vendor

Like most enterprise software companies, the Vendor used a system of commission "accelerators" to compensate salespeople during the timeframe of the study.[2] Under a system using "accelerators," the base commission rate on a sale is multiplied as quarterly sales increase. This form of compensation scheme is common throughout many industries. Joseph and Kalwani (1998) find that over 95% of salespeople reported their salaries were partly based on commissions and bonuses. More specifically, Oyer (2000) notes that well over three-quarters of salespeople were compensated by a non-linear scheme, giving them incentives to time sales to maximize their own compensation. Industry executives and analysts in many technology and other industries, including supercomputing, pharmaceuticals, defense equipment, telecom equipment, semiconductors, and large real estate sales, have reported to the author that salespeople in their industry are compensated using some form of accelerating commission.

Table 1 shows a compensation scheme similar to that used by the Vendor. It has been disguised for confidentiality reasons, but the actual compensation scheme underpins all the reported econometrics. A salesperson making a single $250,000 sale in a quarter would only earn the base commission of 3%, but if the sale were made in a quarter with $2 million in other sales, the salesperson would earn a 12% commission, as he would be on the 4x accelerator. The majority of most salespeople's compensation was earned via commissions, as salaries were only $48,000 per year and the average salesperson made over $100,000 in annual commissions. While the Vendor's sales compensation scheme is aggressive, it is by no means an outlier.

One way in which the Vendor's compensation scheme differed from the most common practice was its use of sales quotas. Quotas are most commonly used as an "incentive-free" threshold that salespeople must clear in order to start earning commissions, or sometimes to earn a bonus (Hall, 2002). The Vendor instead chose to pay low base salaries and pay commissions from the first dollar in sales. The Vendor used a quota system, but only to evaluate salespeople and assign territories, which it did every year. All salespeople had the same quota each year, and a salesperson who performed well above quota in a given year would have some of his customers reassigned to other salespeople, in order to equalize the expected sales of all salespeople. While individual-specific quotas are more common, the Vendor's approach is used by many companies (Cichelli, 2010).

[2] "Accelerators" are sometimes called "multipliers."

Most salespeople had between six and twelve customers in their territory, and were not allowed to sell to other customers. Territory reassignments were decided by the head of sales operations in a way designed to be as bias-free as possible. She built relatively sophisticated models of projected sales by each customer based on past sales, the number of staff in certain functions (because this was correlated with the number of licenses sold for many of the products), and customer financials. Salespeople were not privy to these models, and the decisions around territory assignments were not open to appeal. These decisions were therefore the source of great consternation, and the head of sales operations was in general not a popular employee within the organization, although she was regarded as fair.

As in many large-scale procurement environments, sales negotiations in enterprise software are highly driven by discussions of discounts off of list price, rather than price paid.[3] A last important aspect of the incentive scheme is therefore the degree of pricing discretion by salespeople. At the Vendor, salespeople could give a discount of up to 20% off of list price, regardless of deal size. Deals with higher discounts had to be approved by a manager with higher authority; at 30%, 40%, and 60% discounts, the District Manager, Regional Manager and Vice President of Sales had discretion to approve a proposed deal. Any proposed deal with a discount above 60% had to be approved by the CEO.

3 Data

The data for this study were provided by a leading enterprise software vendor, representing all deals closed by 412 salespeople based in North America for at least two quarters between 1997 and 2003.[4] In total the dataset contains 7,912 deals closed over the course of 28 financial quarters. The dataset was also augmented with publicly available information on customers. The final dataset contains five classes of information:

1. Deal outcomes, which include products bought, list price, and price paid.

[3] Litigation brought about by the United States Department of Justice challenging Oracle's proposed acquisition of PeopleSoft on antitrust grounds provided considerable evidence about the degree to which enterprise software negotiations are discount-driven (rather than focused on dollar amounts). Nearly all testimony about pricing was about discount levels, not price paid. For example, testimony by R. Preston McAfee, (McAfee, 2004) uses the word "discount" or its derivatives was used 236 times by McAfee or lawyers in questions to McAfee. The word "price" is used 206 times, and nearly 75% of mentions of "price" are made in reference to discounts. Dollar levels on individual sales were mentioned far less often.

[4] The data exclude a very small number of "site licenses," which granted customers unlimited use of any of the Vendor's software. Most of the data analyzed in the paper, such as discount granted, are not meaningful for site licenses.

2. Deal timing, which is the date of record for the sale (for both compensation and revenue recognition purposes).
3. Salesperson information, which includes a unique salesperson identifier, tenure, age, gender, full sales and compensation history, territory history, and mobility history across sales districts.
4. Customer information, including information taken from COMPUSTAT and other public data sources[5] (number of employees, annual revenues, end of fiscal year), and information maintained by the Vendor (other technology vendors used, previous customer purchases of the Vendor's products, and average price paid for Vendor's products).
5. Deal's contribution to total quarterly compensation for the salesperson, which is the marginal commission the salesperson earned on the sale in the quarter in question. I also calculated what the marginal commission on each sale would have been had it closed one quarter earlier and one quarter later.

I focus only on license revenue in the empirical analysis. Salespeople received a small commission on service revenue sold under the same purchase order (such as consulting or implementation), but their commission was a fixed percentage of service revenue sold, and did not contribute to the non-linear commission system on license revenue studied here.

The data cover all software products sold by the Vendor. As with most large enterprise software vendors, the Vendor's product lines were very complex. The Vendor divided its engineering and product marketing groups into seven distinct product lines. A single "product" had hundreds of price points, depending on the operating environment, the characteristics of customer servers (e.g., the number of processors), the customer's hardware, and other variables. The Vendor's price book contained thousands of individual price points spanning hundreds of pages.

Table 2 shows summary statistics for the dataset; the first set of figures uses the deal as the unit of observation. The average deal size is nearly \$550,000. (As indicated by the standard deviation, this number is heavily skewed; the median deal size is around \$350,000.) Deals are

[5] For private companies and governmental organizations, a research assistant conducted Internet searches for actual or estimated employee numbers, revenues and information on fiscal year end. For governmental organizations, total budget is used instead of revenues. Data could not be found for fewer than 2% of customers, and averages for non-public companies were used for employee numbers and revenues, and December 31 was used for fiscal year end. The paper's results are robust to not including deals when customer employee or revenue figures were not available.

heavily discounted, with an average discount over 30% and a maximum of 95%. Most tellingly, nearly 70% of deals closed on the last day of the financial quarter, suggesting that the presence of the quarterly deadline in the incentive system carries a dramatic effect. The average marginal commission of nearly $40,000 represents a gross average marginal commission rate of 7.1%.[6]

The average realized marginal commission of $39,700 is statistically significantly greater than the commission on the deal had it closed a quarter earlier ($28,800) or later ($25,800), which demonstrates that deal timing is correlated with marginal commission benefit to salespeople. These two hypothetical differences, for which in the rest of the paper I use the shorthand ΔMB_{t-1} and ΔMB_{t+1}, respectively, are crucial to the analysis. They refer to the monetary benefit to the salesperson that the deal closed in the quarter it did, as opposed to the previous (ΔMB_{t-1}) or subsequent (ΔMB_{t+1}) quarter.

When constructing these variables for each deal in question, the deal was considered the "marginal" deal in the observed quarter, regardless of when it was closed. Since most deals close on the same day, it would be difficult to unpack which deal was in fact the "marginal" deal. If early deals are only assumed to be "marginal" with respect to the preceding quarter, and late deals are assumed "marginal" only with respect to the subsequent quarter, the reported results are stronger. In the spirit of conservatism I therefore maintain the assumption that all deals are "marginal" with respect to moving earlier or later by one quarter.

The bottom part of Table 2 shows summary statistics on the dataset organized at the salesperson and salesperson-quarter levels. The average salesperson makes $900,000 in sales in a quarter,[7] but conditional on a sale, he makes over $1.3 million in sales. This is because 95% of salespeople in the dataset have at least one quarter where they made no sales; in fact, 29% of salesperson-quarters in the data have no sales. These data therefore also reveal preliminary evidence that salespeople manipulate the timing of major sales, since it is unlikely that demand-size factors alone drive the high observed degree of deal bunching.

The data cut at the salesperson level also demonstrate why the Vendor does not require salespeople to make sales in every quarter; the average salesperson only makes about seven sales

[6] This percentage overstates the amount spent on commissions and the actual average commission rate, because it considers each deal "marginal" and the other deals in the quarter "fixed." For quarters with multiple deals, where one deal takes a salesperson past one or more accelerator thresholds, the sum of the marginal commission figures therefore overstates the actual commissions paid. The reason for this approach is explained later in this section.

[7] Again, these averages are heavily skewed. While the average quarterly sales figure is over $900,000, the median is less than $600,000.

per year. Because sales are large and lumpy, and because the sales cycle can extend over several years, it would be difficult for the Vendor to give quarter-based goals to its salespeople. As noted in section 3, the vendor does use an annual quota to evaluate salesperson performance, but that quota is tied to retention decisions, not compensation.

Figure 1, which breaks the financial quarter into 13 weeks and shows the percentage of deals in the database which close in each week, further demonstrates that the closure of deals within a given quarter is not smooth. Not only is there is a large spike in completed deals on the very last day in the quarter, there are spikes of deals early in the quarter and in the last few weeks of the quarter. The line in the figure shows the average discount for deals closing in each week. For the "middle" weeks of the quarter, average discounts hover below 30%; both at the start and the end of the quarter, however, discounts rise to above 35%. The correlations evident in the summary statistics and Figure 1 form the basis of the empirical strategy discussed in the next section.

4 Estimation strategy and results

The estimation strategy takes three steps:

1. Showing that the level of hypothetical salesperson compensation differences across quarters is correlated with the observed timing of deals, both across and within quarters, consistent with the hypothesis that salespeople "pull" and "push" deals due to their compensation.
2. Showing that the level of hypothetical salesperson compensation differences across quarters is negatively correlated with deal pricing (which is equivalent to these differences being positively correlated with discounts given customers).
3. Showing that this pricing correlation does not occur due to differences in customer utility from the product; i.e., ruling out reverse causality (that low observed prices cause hypothetical salesperson compensation differences).

4.1 Timing gaming by salespeople: modeling the deal timing decision

The empirical tests around the timing of deals are inspired by the theoretical model in Oyer (1998) that demonstrates that non-linear incentives can lead to spikes of deals both at the start and the end of financial periods if a salesperson has some influence over the timing of deals. Spikes at the end of a financial period, which appear prominently in the macro-level data in Oyer's research, are consistent with salespeople "pulling" deals forward because they increase their pay compared to the alternative of closing the sale in the next financial period. Although spikes at the beginning of the financial period are not observed in Oyer's data, the model predicts that non-linear incentives will have a "pushing" effect if salespeople increase their pay by delaying a deal until the beginning of the next financial period, compared to the alternative of closing the sale one financial period before.

The critical insight of this model is that "timing gaming" will lead to correlations between salesperson compensation concerns and the timing of deals both *across* and *within* financial periods. "Pulled" deals are more likely to happen in financial periods where the salesperson will earn higher pay for closing a deal in one period, compared to his hypothetical pay if the deal closed in the subsequent period. The model also shows that these deals are more likely to occur at the end of the period. Conversely, "pushed" deals are more likely to occur where the salesperson will earn higher pay for closing the sale in one period, compared to his hypothetical pay had the deal closed in the *previous* period. "Pushed" deals are more likely to occur at the *beginning* of the financial period. I take these predictions to the data.

Specifically, I construct a variable representing when a deal was closed within a quarter – late in (C_L), early in (C_E), or in the middle of a quarter (C_M) – in order to test whether compensation concerns drive the bunching of deals early and late in the quarter, but do not drive the timing of middle deals. "Middle deals" therefore represent a kind of counterfactual of deals where timing within or across quarters was not driven by salesperson compensation concerns. As seen in Figure 2, "middle deals" are associated with the smallest discounts.

The definition of these variables is somewhat arbitrary; Figure 1 suggests a natural definition of "early" is deals closing in the first two weeks of a quarter; a natural definition of "late" is deals closing in the last week of the quarter; and a natural definition of "middle" is the other weeks in the quarter. The empirical results reported in this paper are based on this definition of the timing variable. The data on deal timing are coded to the exact day, allowing the test of alternative definitions of these variables. The results were effectively the same for all reasonable definitions.

It is useful to examine whether basic statistics suggest that the grouping of deals according to when they closed within a financial quarter provides an accurate comparison. If middle deals differ substantially from the rest of deals in the database, for example because they are much smaller in size, then it may not be valid to compare this group to late- and early-closing deals. Table 3 shows the average values for key deal characteristics for early, late, and middle deals as defined above. There is no evidence that there are underlying differences in deal characteristics across the three types of deals. Middle deals are not significantly smaller, nor are they sold to significantly smaller customers, than deals closing early or late in a quarter.

Additionally, Table 3 breaks down deals by customer type across the three "within quarter" timing dummy variables. The only noticeable pattern is that government purchasers are slightly more likely to close middle deals than other customers, but none of the differences are close to statistically significant. Notably, customers who had previously purchased early or late from the Vendor are no more likely to do so again than brand new customers to the Vendor. This is likely because almost all enterprise software vendors use period-based, non-linear compensation, so all corporate customers are used to bargaining under these dynamics.[8]

These facts, and evidence from customer and salesperson interviews, suggest that middle deals are not dramatically different from other deals, except that customers sometimes have budgetary or other constraints which make their timing preferences tight. Notably, Table 3 gives further preliminary evidence of timing gaming, since the change in hypothetical commission had the deal closed a quarter earlier goes down dramatically for early deals, while the change in hypothetical commission had the deal closed a quarter later goes down for late deals. The differences in actual and hypothetical commissions across quarters for middle deals are small, further supporting their use as a quasi-counterfactual.

Formally, I model the probability that deals will close early, late, or in the middle of a quarter as a function of the change in marginal salesperson benefit if the deal closed in the preceding or subsequent quarter. The regression specification is:

$$\Pr(C_i = J) = f(\Delta MB_{i,t-1}, \Delta MB_{i,t+1}, \Omega_i, \varepsilon_i) \qquad (1)$$
$$J \in \{E, L, M\}$$

where C represents the observed timing of the deal within the financial quarter; the subscript *i* refers to the deal in question; the subscript *j* refers to the timing of a deal within a quarter; E, L, and M refer to early, late and middle, respectively; ΔMB represents the change in marginal commissions had the deal closed a period earlier or later (indicated by subscripts *t-1* and *t+1*, respectively); Ω represents a vector of controls; and ε represents the error term.

[8] Larkin (2011) contains anecdotal evidence that IT employees at software customers talk to other customers on Internet discussion boards about the use of "accelerators" and the resulting ability to drive high discounts.

To estimate the model given in equation (1), I run a multinomial logit with the set of three deal timing dummies (close early, close late, close middle) as the dependent variable, the calculated change in compensation had the deal closed a quarter earlier and a quarter later as the main explanatory variables, and a full set of controls. I would expect large changes in marginal benefit across quarters to be correlated with a deal's closing early or late in a quarter.

The controls include product line, customer industry, and salesperson region dummies; the deal's total purchase price; salesperson tenure; and whether the deal comes in the Vendor's and/or customer's final quarter of their fiscal year. Product line dummies were constructed based on three factors: the vendor's seven dedicated product lines, the operating system (UNIX vs. non-UNIX), and whether the product was less than two years old or more than two years old. There were therefore 28 effective product lines controlled for.

The need for customer and sales region dummies is clear; as with product lines, these variables could be related to a deal's timing, regardless of the incentive effects on salespeople. I control for purchase price in case salespeople are averse to attempting to game larger deals because of the risk of losing them. I control for salesperson tenure as there may be differences in a salesperson's propensity and/or ability to engage in timing gaming as he becomes more experienced. I control for the final quarter of the Vendor's fiscal year because executive pay largely depends on fiscal year-end stock price, leading executives to make it more difficult for salespeople to "push" deals out of the fourth quarter, and motivating them to "pull" more deals into this quarter. I control for customer fiscal year-end effects as well, since demand-side IT budgeting cycles may cause customers to rush to purchase by the end of their fiscal year.

In this and all regressions reported in the paper, I cluster standard errors at the level of the salesperson, as the data are panel-like in structure and the main explanatory variables are calculated by taking into account variables from previous or future periods.

Table 4 reports the marginal effects from estimation evaluated at the mean value of the variables. These coefficients can usefully be interpreted as the change in probability of choice *j* due to a one-unit change in the value of the independent variable. These changes in probabilities are reported in comparison to the baseline choice, which is to close the deal in the middle of a quarter.

The results reported in Table 4 suggest the existence of a strong relation between changes in marginal benefit to salespeople across quarters and the probability that the deal in question closes early, late, or in the middle of a quarter. The coefficient in column A on the ΔMB_{t-1} variable suggests that a \$1,000 reduction in commission had the deal closed a quarter earlier is associated with a 0.8% greater likelihood of the deal closing early in the quarter. Similarly, as reported in column B, a \$1,000 reduction in expected commission had the deal closed a period later is associated with a 1.9% greater likelihood of the deal closing late in the quarter. Both estimates are significant at standard statistical levels.

Due to the large spread in the ΔMB data, it is useful to examine the model's predicted likelihood of "pushing" or "pulling" across a wider range of these variables. Consider a deal where the ΔMB_{t-1} or ΔMB_{t+1} variables took values of \$50,000, which would occur, for example, if a salesperson had a \$400,000 deal for which he engaged in "timing gaming," \$2 million in other sales in quarter *t*, and no other sales in the previous or subsequent quarters. The model predicts that the probability of this deal being "pushed" reaches nearly 30%, an increase of almost 300% from the observed average probability. The model predicts the likelihood of the deal being "pulled" reaches 90%, an increase of over 15 percentage points from the observed average probability. Both of these point estimates are significant at the 1% level. The kinks in the commission curve apparently have a large effect on the timing of deals both across and within quarters.

Table 4 contains some other interesting results. Closing around deadlines does not appear to be correlated with deal size or any customer observable, including whether the deal is in the final quarter of the customer's fiscal year. I therefore do not find evidence of demand-side effects causing bunching. More highly-tenured salespeople are more associated with deals closed around the deadline; these salespeople may be more adept at "timing gaming," or their superiors may allow them to engage in a greater amount of "timing gaming" through the discount approval process.

Of course, these results do not demonstrate a causal link between differences in salesperson compensation and the timing of deal closing. Some unobserved factor in underlying demand could produce this pattern of timing. However, it is difficult to think of factors affecting demand which would correlate so strongly with salesperson incentives. Furthermore, interviews with customers suggest that the timing of demand would be random across customers since *all* customers attempt to manipulate timing to achieve bigger discounts, and since the reasons not to do so are therefore idiosyncratic to the deal, not the customer. Table 3 contains high-level evidence corroborating these assertions. Finally, Oyer (1998) used the natural experiment of exogenous merger and acquisition activities to show that a similar, macro-level revenue timing result was not an artifact of unobserved differences in customer demand.

A second alternative explanation is that the deadline in the incentive system causes salespeople to work harder near the end of the quarter, leading to a prevalence of deadline deals. Of course, this result would be an interesting validation of agency theory in itself. However, more importantly, it would not explain the changes in probability of *early* deals as marginal benefit in the previous quarter declines.[9] Finally, it would not explain why salespeople choose not to sell at all in nearly 30% of quarters.

4.2 Giving bigger discounts to deals: modeling deal outcomes

Having demonstrated that gaming by salespeople appears to affect the flow of deals both within and across financial quarters, I next turn to the question of how this gaming affects deal outcomes. A key question is the variable to use as the outcome measure. One outcome measure would be unit price paid per software license. However, the complexity of the Vendor's product lines and pricing structure mean that it is infeasible to use this measure to study pricing variance across deals.

However, there is significant evidence that this very complexity in pricing drives customers to negotiate discounts, not prices. A market expert stated that, "Discounting has long been a fixture of the enterprise software business, where list prices exist only in theory" (Riccuiti, 2004). Since the discount measure normalizes all deals and gives a direct unit of comparison, it is also useful as an outcome measure. Indeed, the Vendor's deal approval system focuses on discounts as the key decision variable, regardless of deal size.[10] While differences in discount propensities obviously occur, for example for new products, product-line fixed effects control for these differences. I therefore use total discount as the outcome measure in the empirical deal outcome tests.

I specifically test whether salespeople grant bigger discounts to customers when they close deals in the quarter which maximizes their compensation. The hypothetical compensation difference variables are a direct measure of how strongly a salesperson prefers the deal to close in a given quarter, so it is a natural explanatory variable to regress against discount given. For

[9] One might argue that some deadline deals "spill over" into the early part of the next quarter, because the salesperson had too much to do at the deadline. This would not explain why early deals are so correlated with the commission benefit to salespeople for having the deal close in the next quarter, not at the deadline of the previous quarter.

[10] Larkin (2011) reports an incidence of an enterprise software vendor waiving deal approval requirements for small deals. A savvy salesperson reacted to this policy change by carving up a $24 million deal into hundreds of increments, and selling them at an 85% discount.

simplicity and ease of interpretation, I use Ordinary Least Squares (OLS) in the estimation.[11] Notationally, the regression equation is:

$$Y_i = \beta_1 + \varphi_{t-1} * \Delta MB_{i,t-1} + \varphi_{t+1} * \Delta MB_{i,t+1} + \beta_2 * \Omega_i + \varepsilon_i \quad (2)$$

where the subscript *i* again refers to the deal observation, Y refers to the discount given, the ΔMB variables are as defined earlier, Ω_i is a vector of deal controls, and ε_i is the error term. Standard errors are robust and clustered at the salesperson level.

This regression directly compares discounts given to customers when salespeople face a large change in commission earned had the deal closed a quarter earlier or later, to discounts given when the commissions do not differ across quarters. I therefore expect the signs on both the φ coefficients to be positive.

The most important control variable is the use of salesperson fixed effects. This is because salespeople have different abilities and different status levels in the Vendor's organization, and may also sell to different types of customers. By using salesperson fixed effects, the model focuses on within-salesperson changes in discount due to the presence of financial concerns.

Other control variables are largely the same as in the deal timing model: a full set of product line and customer industry dummies; deal size; salesperson tenure; and basic customer information such as size and revenue. I also introduce quarter dummies, to control for the possibility that the Vendor may be more or less lenient in approving discounts in certain quarters of the fiscal year. In addition, I introduce controls on the customer's previous purchases of the Vendor's products. In an effort to induce customers to initially buy a package, the Vendor will often grant very large discounts, in the hope of charging quasi-monopoly prices later as customers upgrade.[12] I therefore control for whether the customer is new to the Vendor and/or new to the product line in question.

Model (A) of Table 5 presents the results of the estimation, which strongly indicate that deal outcomes are correlated with salesperson salary concerns. The estimates of φ_{t-1} and φ_{t+1} are easily interpreted; a deal for which the salesperson would have received a commission \$1,000 lower had

---

[11] I have also carried out this model using log-odds, since the dependent variable represents a percentage. The results are extremely similar, but coefficients in the linear model are much simpler to interpret.

[12] For a discussion of these dynamics, and an empirical investigation on the depth of product lock-in and the premiums vendors can charge, see Larkin (2008).

the deal closed a period earlier is discounted 0.26 percentage points more; the discount for a deal with a $1,000 lower commission if closed a quarter later is 0.44 percentage points higher. Both coefficients are significant at the 1% level. Evaluated at the ΔMB variables' average values of 10.9 and 13.9, respectively, the average "pushed" deal is discounted by 2.8 percentage points more than a deal where a salesperson is indifferent about the quarter in which the deal is closed, while the average "pulled" deal is discounted by 6.2 percentage points more. Again, we find no evidence that demand-side concerns are driving discounts; the customer's fiscal year end is insignificant.

One simple alternative explanation for the results in Model (A) is that deadlines drive the larger discounts, not salesperson financial concerns. The timing model results from Table 4 demonstrate that timing is highly correlated with salesperson compensation concerns. There are other explanations for the deadline effect: that the Vendor presses salespeople to close deals so that the Vendor meets its financial targets; a belief on the customer side that waiting until late in the quarter is the optimal strategy; or even the psychological effect of deadlines, which has been well documented experimentally (e.g., Roth, Murnighan, and Schoumaker, 1988).

In Model (B) of Table 5, I test a simple deadline effects model by replacing the sales compensation concern variables with the C_E and C_L deal timing dummy variables. As expected, the coefficients on these dummy variables are both highly significant. In Model (C), I use both the salesperson compensation variables and the deal timing dummies as explanatory variables, which in effect runs a "horse race" between two distinct hypotheses: that waiting until the end of the quarter (or buying at the start of the quarter) is in itself enough to drive higher discounts, or that salespeople compensation concerns drive discounts. Put another way, this model tests whether simply waiting until the end of the quarter is enough to drive a discount, which matches customer belief.

The results of Model (C) suggest that compensation concerns, not deal timing, drive the differences in discounts. The coefficients on the timing variables fall out of significance, while the coefficients on the compensation concerns variables stay positive and highly significant, and in fact are very similar to those of Model (A). This finding is critical because it suggests that the common customer strategy of purchasing at the end of the quarter is not in itself useful in achieving a lower price. In fact, the results suggest that salespeople are adept at keeping prices at non-distorted levels even at the end of quarters, provided they do not have compensation riding on the outcome.

It is informative to examine tenure effects. However, there is little variation in tenure levels within an individual salesperson, compared to variation across salespeople. The use of salesperson fixed effects therefore complicates the tenure analysis. Tests using tenure as an explanatory variable are shown in an earlier version of the paper (Larkin, 2012), and indicate that within-person tenure changes do not affect the propensity to discount when faced with compensation concerns, but across-salesperson variation in tenure *is* correlated with compensation-based discounting. These results may suggest that salespeople who are prone to game are more likely to stay employed at the Vendor, while those who do not wish to game choose to leave the Vendor.

4.3 Salesperson compensation or price discrimination?

These results may be biased due to the omission of critical explanatory variables in the regression model. Most notably, underlying customer valuation or utility from the product is not observable, and therefore cannot be accounted for in the regression. The counterfactual from the discount regressions implicitly assumes all deals in the dataset would still happen without the observed differences in salesperson compensation, and in fact the counterfactual in the estimation is precisely a deal where ΔMB is zero. However, if low-value customers would not buy at all unless large discounts were offered, then the implied counterfactual is invalid. In this case, the results in Table 5 would suffer from reverse causality: large salesperson compensation concerns would be driven by low-value customers "bunching" into quarters.

This model may therefore suffer from the classic "Heckman problem" (Heckman, 1979) — it is unclear whether the differences in deal outcomes stem from the explanatory variables and controls used in the model, or instead reflect unobserved customer heterogeneity. This is a common problem in studies similar to this one. Most of these studies make the same assumption made in the preceding model: that customer utility is random. They therefore do not attempt to distinguish between "the economic, institutional and spurious statistical factors" that lead to heterogeneous deal outcomes (Greenstein, 1993). Because there is such a clear reason that customer utility may be non-random with respect to salesperson compensation concerns, it would be inappropriate to not attempt to correct for this potential problem.

Models (B) and (C) partially alleviate this concern, by showing that deadline effects alone do not drive higher discounts, especially because industry experts commonly report that customers usually do not know the value they receive from large software purchases (e.g., Ricciuti, 2004). Therefore the price discrimination hypothesis is arguably less likely than the hypothesis that compensation causally drives larger discounts. Of course, this anecdotal evidence requires additional scrutiny in order to definitively prove causality.

With non-experimental data, it would be ideal to use an instrumental variable correlated with salesperson compensation concerns but not customer utility. However, since salesperson compensation concerns are so strongly driven by the non-linear nature of the compensation curve, it is difficult to find variables that are sufficiently correlated with the ΔMB variables but are not subject to the same critique around potential endogeneity.[13] There are several quasi-exogenous changes which affect salesperson compensation, such as changes to product list prices and customer territory assignments. However, the correlation between the compensation differences brought about by these changes and the ΔMB variables is so weak that these variables fail as instruments. Also, it is not clear that these changes are uncorrelated with differences in customer utility for a product.

I therefore attempt a less ideal approach to test the alternative hypothesis of price discrimination: using a series of robustness checks to look for evidence that customers with potentially lower utility are more likely than other customers to purchase when salesperson compensation concerns are high. This approach cannot definitively prove causality, but it can provide important corroboratory evidence about the most likely causal mechanism behind the results reported in Table 5 on the positive correlation between ΔMB and discounts.

Historical variance in pricing across specific products is arguably a good proxy of heterogeneity in customer valuation. Higher variance in historical prices may be indicative of price discrimination across customers based on variance in utility; indeed, variance in pricing has been used in empirical research as a proxy for the ability to price discriminate (e.g., Bennett, 2013).

The Vendor provided a measure of pricing variance for all products in the database from 1991 to 2003; it therefore was a measure of 13 years of historical variance and included data from the database used in this study. It did not contain the number of deals, so I was unable to usefully compare pre-1997 deals to observed pricing variance in the dataset. Using only variance in the deals in the database yields very similar results. In the historical data, the average product is discounted by 32.3 percentage points, very similar to the figure in the database, and the average within-product standard deviation was 11.6. The largest observed standard deviation was 42.1 (on a product with very low sales in the database), and the lowest was 3.8. Historical price variation only

[13] For example, many studies use data on other observed units, which in the context of this study would be compensation concerns of other salespeople, to construct an instrument. In this case, all salespeople are paid via the same commission scheme, so using data on other salespeople does not avoid the endogeneity problem.

weakly correlated with discount levels (p=0.04), suggesting that differences in variation were driven approximately equally by high-paying and low-paying customers.

In Model (D) of Table 5, I introduce the measure of historical pricing variance for products, and this measure interacted with salesperson compensation concerns. Note that the price variance measure was at the detailed product level; i.e., it was calculated separately for different iterations of the product, and therefore is not collinear with the product line fixed effects, which are broader. The regression results confirm the weak correlation between historical variance and discounts; a ten percentage point increase in the standard deviation in pricing increases discounts by 0.7 percentage points, but the coefficient is only 10% significant. More notably, the coefficients on the interaction terms between historical variance and the ΔMB variables are insignificant, indicating products with higher historical variance are no more likely to be associated with discounts in deals with large salesperson compensation concerns than products with lower historical variance.

I next analyze whether customers with higher historical discounts are more associated with discounting in quarters where the salesperson has strong compensation concerns. The Vendor provided me with historical data from 1992 to 1996 on average discount at the level of the individual customer. Data from 1997 to 2003 were not used in this analysis, since discount given is the dependent variable in the regressions, and including deals from 1997-2003 would therefore treat information on observed discounts as both dependent and independent variables in the same regression. Therefore, this analysis necessitated dropping the 40% of deals that involved customers with no deals between 1992 and 1996. For this variable, the average customer received a 39.0% discount, with a standard deviation of 12. Both the maximum and minimum average discount were for customers with two purchases, and therefore are not meaningful. If customer utility is stable across time, under the price discrimination hypothesis we would expect customers with higher historical discounts to be more likely to purchase in quarters with high salesperson compensation concerns compared to other customers.

Model (E) of Table 5 introduces average historical customer discount as a proxy for customer utility. Average historical discount is much more highly associated with higher discounts on observed deals; customers with an average discount that is 10 percentage points higher than average until 1997 receive a 3.2 percentage point increase in discount on observed deals between 1997 and 2002. But these customers are not getting differentially higher discounts in quarters with large salesperson compensation concerns. The estimates on the baseline ΔMB variables come down by approximately 25%, but remain highly significant. Therefore, differential utility may explain some of the observed correlation between discounts and compensation concerns, but a significant portion remains even after accounting for the potential for price discrimination.

I conducted several other tests involving proxies of customer valuation, which are reported in a previous version of the paper (Larkin, 2012). For example, I examine customers that the Vendor chooses to reassign to a new salesperson because the customer had not bought from the Vendor in some time, since these customers may have low utility for the Vendor's products. I also examine pricing for new products, which industry analysts maintain are only bought by high-value customers, compared to pricing for older products. In none of these or other similar tests do I find evidence that proxies for customer utility are driving the observed correlation between discounts and salesperson compensation concerns. While not definitive proof of causality, the two leading explanations of the result – the simple "deadline effect" and the price discrimination hypothesis – appear to be invalid.

4.4 What is the cost of the incentive system in terms of foregone revenue?

It is informative to think about the cost of this incentive system in terms of foregone Vendor revenue. It is important to note that this calculation does not represent a full cost-benefit analysis of the incentive system. One of the important lessons of basic agency theory is that all incentive systems carry benefits (e.g., increased effort, beneficial sorting) and costs (e.g., wage premiums). The cost estimates in this section are only concerned with measuring one type of incentive system cost that has not been previously explored in detail: the cost of mispricing due to "gaming."

One estimate of the cost can be calculated using the estimated coefficients for φ_{t-1} and φ_{t+1}, the effect of marginal changes in compensation across quarters on discount levels, shown in Tables 5 and 6. For example, I noted previously that applying the coefficients from Model (A) of Table 5 to the average values of the ΔMB variables suggests the average "pushed" deal is discounted 2.8 percentage points more than a deal where a salesperson is indifferent about the quarter in which the deal is closed, while the average "pulled" deal is discounted by 6.2 percentage points more. As proxied by deals closing early in a quarter, 8.8% of deals in the database are "pushed," while 74.2% of deals are "pulled," as proxied by deals closing late. Therefore, the average gamed deal is discounted by 4.8 percentage points more than it would have been had it not been "gamed." However, this calculation is based on list prices, not realized prices; since the average deal results in revenue capture of only 65.2% of list price, the cost to the Vendor in terms of foregone revenue is 7.4%.

When applying this logic to the four relevant models in Table 5,[14] the cost of the incentive system is calculated to be between 5.69% (in Model E) and 8.65% (in Model C). These estimates suggest that the actual costs of a non-linear, period-based incentive scheme may be up to twice what they appear to be, if one only considers the direct cost of commissions, which average about 6%.

It is important to note that this figure represents an upper bound on the pricing cost of gaming, because it implicitly assumes that all "gamed" deals would have been carried out at a higher price. However, because there is no evidence that customer utility differences drives the "gaming" result, this assumption appears largely realistic. In addition, this cost estimate does not include the potential cost of excess compensation paid to salespeople due to gaming.

5 Discussion

Following the classic paper by Kerr (1975), many of the papers in the literature on incentives quote the old adage, "Firms get what they pay for;" this research demonstrates the extent to which this adage holds true. An incentive system ideally motivates value-creating effort, but this study suggests that salespeople put significant effort into manipulating the timing of deals, which actually hurts the Vendor due to mispricing. Interestingly, the Vendor spends approximately 6% of revenue on sales commissions, a number very similar to the average across industries (Godes, 2003). Given that the pricing cost of the incentive scheme is estimated to be 6-8%, compensating salespeople is therefore twice as expensive as it appears at first glance.[15]

One puzzle from the results is why salespeople cut price at all, when the majority of customers seemed to be primed to wait until the end of the quarter to agree to deals, and when these customers appear to agree to deals at higher prices when the salesperson does not have compensation concerns. Put another way, salespeople are very adept at using discounts to incentivize the closing of deals in quarters where they have strong financial concerns, but also appear adept at "toeing the line" on pricing in other quarters, even when they sell deals at the

[14] Model (B) – the "deadline effects" model – is excluded because it incorrectly omits salesperson compensation concerns. However, the cost estimates from Model (B) are actually very similar to those reported. In addition, in the alternative robustness checks reported in a previous version of the paper (Larkin, 2012), no estimated cost was below 6%.

[15] Vendor executives reported to me that they knew timing gaming was potentially costly, but felt they managed it well through the discount approval system. They were therefore very surprised at the size of the estimates.

quarter deadline. Salespeople earn lower commissions when they cut price, so *ceteris paribus* they prefer to keep prices high.

It could be that reductions in price and commission of an average of 6-8% is a worthwhile "insurance premium" for salespeople to pay when faced with the prospect of a commission reduction much larger than that due to the effects of the accelerators. Since accelerators can increase commissions by well over 10x, it is easy to envision a salesperson deciding that a small reduction of his commission due to a larger discount is a worthwhile insurance premium to pay to receive such a large increase in commission. Indeed, it would be fascinating in future research to calculate the magnitude of this implied insurance cost.

As noted previously, there is no evidence that customer valuation differences significantly account for the link between discounts and compensation. Customer valuation differences surely exist, but discounts that occur from them do not appear to be driving a significant percentage of the price reductions that occur due to salesperson compensation concerns. The largest reduction in effect size in the historical customer discount model reduced the estimated effect of sales compensation by less than 25%. The estimated pricing cost of the incentive system in this model was close to 6% of revenue.

In fact, the results are consistent with an alternative interpretation around the use of non-linear incentive systems: perhaps they are a method by which to make hesitant customers unsure about the purchase more comfortable about making a major purchase decision. For example, conventional wisdom says car purchasers should wait until the end of the month to buy a car, because prices will be lower since salespeople who "haven't made their quota" will be desperate to close a deal. Some industry insiders say this conventional wisdom was carefully promulgated by car dealers as a sales technique to make customers believe they received a better deal than they did (Snowe, 2007). This is a different take on the price discrimination critique that suggests one cannot be sure that low valuation customers would make observed purchases without the discounts offered by financially-motivated salespeople. Under this alternative explanation, <u>*all*</u> customers are less likely to purchase without the psychological comfort that comes from feeling like they got a good deal. So the incentive scheme may drive bigger discounts, but this could be to take advantage of a behavioral bias of customers. Again, this is an interesting area for future research.

This study has a number of limitations. Most prominently, it looks at a single vendor in a single institutional setting over a relatively short span of time. That said, it goes further than other research in estimating the depth of incentive system gaming and the resulting effect on business outcomes. Furthermore, as demonstrated by Oyer (1998) and others, non-linear, period-based deadlines are commonplace in business settings, which extends the generalizability of these findings. Simply put, deadlines and non-linear incentives are the rule, not the exception for a wide category of employees and business situations. The Vendor is widely held to be representative of the industry, and many high technology and service industries use similar high-powered compensation schemes for employees. It would, however, be useful to extend the study to other vendors, industries and time periods.

Secondly, the study does not involve a natural experiment utilizing randomized changes in incentives for salespeople. Although the results appear robust to the most logical alternative explanation, the research cannot definitively prove causality. It is important to note that randomized changes in incentives with firms as large as the Vendor are extremely rare, exactly because incentives are so important. Furthermore, this study goes further than many other correlative studies of incentives in thinking about causality.

Finally, it is important to reiterate that the pricing cost identified in the study does not suggest that the incentive system is suboptimal. It is highly likely that there are effort and sorting benefits that are brought by the incentive system, which this study does not attempt to measure.

The results of the study suggest several other avenues for further research. First, as noted, there are interactions between tenure and gaming. It is unclear whether salespeople "learn to game" over time, or whether salespeople who enjoy or are adept at gaming are more likely to stay at the company. Understanding the causal factors for the correlation between tenure and gaming, and better estimating the costs and benefits of higher tenure would broaden our understanding of internal labor markets. Second, there is an interaction between salesperson, sales manager and senior executive compensation structures, which at times can compete with each other. This aspect of "gaming" remains under-explored. Also, the delegation system used by the Vendor, as represented by its escalating deal approval process as proposed discount goes up, has clear effects on outcome, since the majority of the Vendor's deals are sold at a discount representing the exact limit of an employee's approval authority.

Finally, the study is silent on the natural and important question of why such distortionary systems are so prominently used in the first place, and the benefits to such schemes which must exist given their prominence in the presence of such large apparent gaming costs. A behavioral explanation around customer purchasing decisions was briefly explored above. Recent research (Larkin and Leider, 2012) suggests that behavioral biases of salespeople may explain non-linear incentive schemes. If salespeople are overconfident, they may base recruiting and effort decisions on highly-compensated parts of the incentive curve, even if they are unlikely to reach them. This explanation would help explain why it is uncommon in sales environments for companies to offer differential compensation curves based on observed skill levels; doing so would reduce the positive effects on the attraction, retention and motivation of overconfident workers. Still, this explanation underlies the fundamental point that more research should be done on the benefits of non-linear incentive schemes given their prominence and the significant gaming costs they incur.

References

Asch, Beth J. 1990. Do incentives matter? The case of navy recruiters. *Industrial and Labor Relations* Review 43(3): 89S-106S.

Basu, Amiya K., Rajiv Lal, V. Srinivasan, and Richard Staelin. 1985. Salesforce compensation plans: An agency theoretic perspective. *Marketing Science* 8(3): 324-342.

Bennett, Victor. 2013. Organization and bargaining: Sales process design at auto dealerships. *Management Science* (forthcoming).

Chevalier, J. and G. Ellison. 1997. Risk taking by mutual funds as a response to incentives. *Journal of Political Economy*. 105(6):1167-2000.

Cichelli, David J. 2010. *Compensating the Sales Force: A Practical Guide to Designing Winning Sales Reward Programs*. 2nd ed. New York: McGraw-Hill.

Godes, David. 2003. In the eye of the beholder: An analysis of the relative value of a top sales rep across firms and products. *Marketing Science* 22(2):161-187.

Greenstein, Shane M. 1993. Did installed base give an incumbent any (measurable) advantages in federal computer procurement? RAND *Journal of Economics* 24(1):19-39.

Hall, Brian J. 2002. Incentive strategy within organizations. Harvard Business School Course Note 902-131.

Healy, Paul M. 1985. The effect of bonus schemes on accounting decisions. *Journal of Accounting and Economics* (7)1-3:85-107.

Heckman, James J. 1979. Sample selection bias as a specification error. *Econometrica* 47(1):153–61.

Holmstrom, Bengt. 1979. Moral Hazard and Observability. *Bell Journal of Economics*, 10:74-91.

Joseph, Kissan, and Manohar U. Kalwani. 1998. The role of bonus pay in salesforce compensation plans. *Industrial Marketing Management* 27:147-159.

Kerr, Steven. 1975. On the folly of rewarding A, while hoping for B. *Academy of Management Journal*, 18(4):769-783.

Lal, Rajiv, and V. Srinivasan. 1993. Compensation plans for single- and multi-product salesforces: An application of the Holmstrom-Milgrom model. *Management Science*, 39(7): 777-793.

Larkin, Ian. 2008. Bargains-then-ripoffs: Innovation, pricing and lock-in in enterprise software. *Best Paper Proceedings of the Academy of Management.*

Larkin, Ian. 2011. Arck Systems (TN) (A), (B), (C), (D), (E) and (F). Harvard Business School Case Teaching Note 911-074.

Larkin, Ian. 2012. The cost of high-powered incentives: Employee gaming in enterprise software sales. Working Paper. http://www.people.hbs.edu/ilarkin/HPI.pdf (accessed February 20, 2013).

Larkin, Ian, and Stephen Leider. 2012. Incentive schemes, sorting and behavioral biases of employees: Experimental evidence. *American Economic Journal: Microeconomics*. 4(2):184-214.

Lazear, Edward P., and Paul Oyer. 2012. Personnel economics. In *The Handbook of Organizational Economics*, ed. Robert Gibbons and John Roberts. Princeton, NJ: Princeton University Press.

Leventis, Andrew. 1997. Cardiac surgeons under the knife. Working paper, Princeton University.

McAfee, R. Preston. 2004. "Testimony of R. Preston McAfee, United States vs. Oracle (Oracle – Peoplesoft acquisition). http://vita.mcafee.cc/PDF/OracleTestimony.pdf (accessed June 20, 2010).

Oyer, Paul. 1998. Fiscal year ends and nonlinear incentive contracts: The effect on business seasonality. *Quarterly Journal of Economics*. 113(1):149-185.

Oyer, Paul. 2000. Theory of sales quotas with limited liability and rent sharing. *Journal of Labor Economics*. 18(3):405-426.

Prendergast, Canice. 1999. The provision of incentives in firms. *Journal of Economic Literature*. 37(1):7-63.

Ricciuti, Mike. 2004. Eyes on enterprise: SAP-Microsoft talks underscore harsh market reality. CNET News. http://news.com.com/Eyes+on+enterprise/2009-1001_3-5250319.html

Roth, Alvin E., J. Keith Murnighan, and Francoise Schoumaker. 1988. The deadline effect in bargaining: Some experimental evidence. *The American Economic Review*. 78(4):806-823.

Shaw, Jason, Nina Gupta, and John E. Delery. 2000. Emprical organizational-level examinations of agency and collaborative predictions of performance-contingent compensation. *Strategic Management Journal*. 21:611-623.

Snowe, Daniel. 2007. The used-car salesman's training manual: 25 tricks they use to charge you more. *Inside CRM Newsletter*. http://www.insidecrm.com/features/used-car-salesman-training-manual-121807/ (accessed February 18, 2011).

Table 1: Illustrative enterprise software salesperson quarterly commission schedule

Sales revenue generated	Commission accelerator	Incremental commissions earned
First $500,000 in sales	1x	3% of sales (max of $15,000)
Next $500,000 in sales	2x	6% of incremental sales (max of $30,000)
Next $1,000,000 in sales	3x	9% of incremental sales (max of $90,000)
Next $2,000,000 in sales	4x	12% of incremental sales (max of $240,000)
Next $2,000,000 in sales	6x	18% of incremental sales (max of $360,000)
Any sales above $6,000,000	8x	24% of incremental sales

Source: Disguised example from company providing data for this research.
Note: does not include base salary of $12,000 per quarter.

Table 2: Summary statistics for key variables

Variable	Unit	Mean	Std. Dev	Min	Max
DEAL DATASET (N=7,912 deals)					
Basic deal characteristics					
Total list price	\$1,000	856	2,305	10	36,125
Total price paid	\$1,000	557	942	10	14,450
Total discount given	percent	34.8	19.8	0	95
Discount exactly on approval band (20%, 30%, 40%, 60%)	1=yes	0.64	0.48	0	1
New product (<2 years old)	1=yes	0.16	0.37	0	1
Deal timing characteristics					
Deal closed on last day of quarter	1=yes	0.67	0.47	0	1
Deal closed in last week of quarter (but not day quarter ended)	1=yes	0.04	0.20	0	1
Deal closed in middle eleven weeks of a quarter	1=yes	0.21	0.31	0	1
Deal closed in first two weeks of quarter	1=yes	0.08	0.24	0	1
Deal closed in Vendor's 4^{th} quarter	1=yes	0.30	0.46	0	1
Salesperson characteristics					
Tenure at time of deal closing	# of quarters	9.90	7.5	1	**
Customer characteristics					
New to vendor	1=yes	0.27	0.44	0	1
New to product line	1=yes	0.52	0.50	0	1
Annual revenue of customer	\$ billion	6.5	17.9	0	**
# of customer employees	1,000	5.2	14.0	0.03	**
Deal closed in customer's 4^{th} quarter	1=yes	0.28	0.45	0	1
Compensation characteristics					
Marginal commission on deal (quarter t)	\$1,000	39.7	74.6	0.2	1,130
Marginal commission had the deal closed a quarter earlier	\$1,000	28.8	50.5	0.2	1,050
Hypothetical commission difference versus deal closing in previous quarter (ΔMB_{t-1})	\$1,000	10.9	37.6	-396.8	444.1
Marginal commission had the deal closed a quarter later	\$1,000	25.8	52.5	0.2	1,050
Hypothetical commission difference versus deal closing in subsequent quarter (ΔMB_{t+1})	\$1,000	13.9	43.5	-381.2	465.1
SALESPERSON DATSET (N=412 salespeople; 4,609 salesperson-quarters)					
# of employed quarters, 1997-2003	quarter	11.19	5.62	2	28
Average number of sales per quarter	count	1.72	1.25	0.33	3.91
Average quarterly sales revenue	\$1,000	951.7	906	46.0	4,172
Average quarterly commissions earned	\$1,000	68.2	118.3	0.1	492.3
Conditional on having a sale in a quarter:					
Average number of sales per quarter	Count	2.43	1.68	0	5.12
Average quarterly sales revenue	\$1,000	1,354	974	46.0	5,564
Average quarterly commissions earned	\$1,000	96.5	134.7	0.1	656.4

Note: ** represents that the data is not reported per agreement with the provider of the dataset (to protect its identity or identity of customers).

Figure 1: Deal timing and discounts within financial quarter

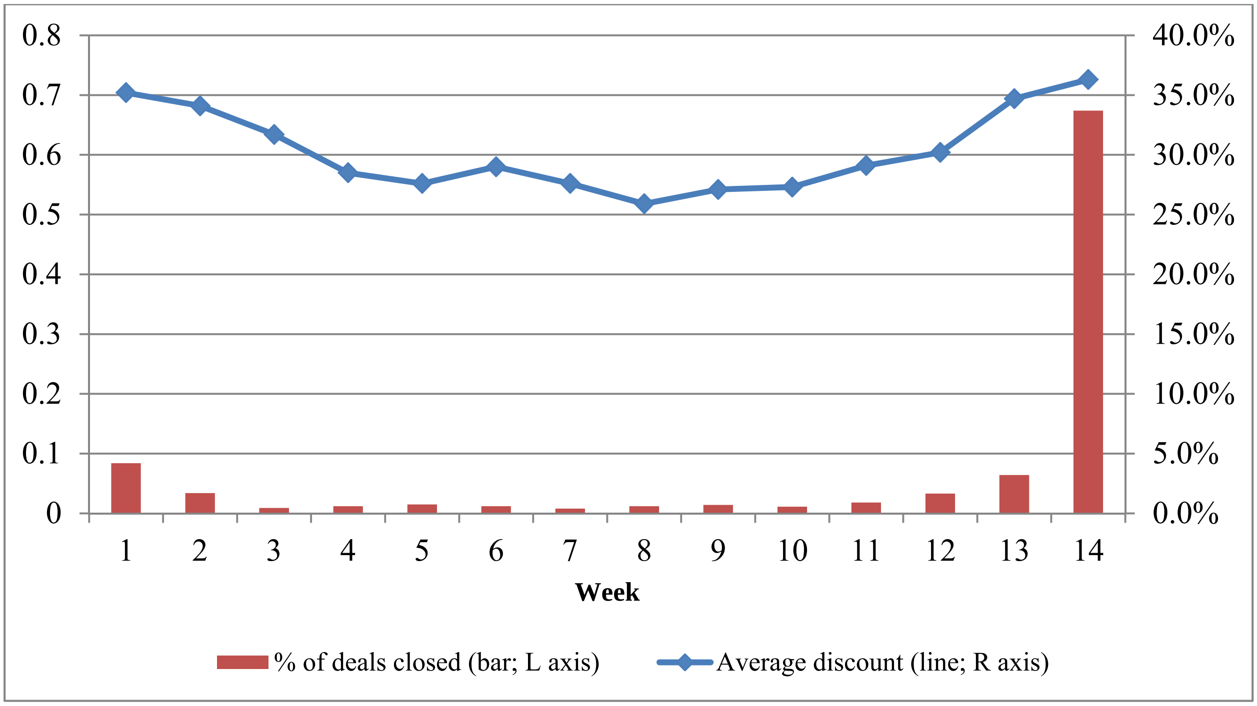

Note: X-axis refers to week of deal closing in the Vendor's financial quarter. Data are aggregated across the 28 quarters in the study. Week 13 refers to the last week of the financial quarter EXCEPT for the last day in the financial quarter. Week 14 refers to the last day of the financial quarter.

Table 3: Summary statistics and customer analysis of early, late and middle deals

Customer type	Unit	Early deals (C_E)	Middle deals (C_M)	Late deals (C_L)
Variable averages				
Total price paid	$1,000	548	532	564
Total discount given	Percent	34.6	28.8	36.2
Salesperson tenure	Number of quarters	9.5	9.2	10.1
Annual revenue of customer	$billion	6.0	6.0	6.7
Marginal salesperson commission	$1,000	41.1	29.4	41.9
Hypothetical marginal commission had the deal closed a quarter earlier	$1,000	24.3	30.2	29.0
Hypothetical marginal commission had the deal closed a quarter later	$1,000	33.4	33.7	23.1
Percentage of deals by customer type				
All customers	Percent	8.8	17.0	74.2
Private sector customers	Percent	8.6	16.5	74.9
Government and education customers	Percent	9.3	17.8	72.9
Customers with previous purchases early or late in the financial quarter	Percent	8.3	16.6	75.1
Customers with no previous purchases from the vendor	Percent	8.8	17.1	74.1

Note: Early deals = weeks 1 and 2; late deals = week 13 and last day of quarter; middle deals = weeks 3-12.

Table 4: Deal timing model, marginal effects after multinomial logit

	(A) Pr (C_E) - Pr (C_M)	(B) Pr(C_L) – Pr (C_M)	(C) X (average variable value)
Hypothetical Commission Δ_{t-1}	.008 (.004)**	.000 (.001)	10.9
Hypothetical Commission Δ_{t+1}	-.006 (.008)	.019 (.007)***	13.9
Log deal size	.081 (.079)	-.045 (-.050)	6.32
Salesperson tenure	.014 (.005)***	.031 (.011)***	9.90
Vendor Quarter 4 deal	.030 (.021)	.065 (.016)***	0.30
Log customer revenues	-.002 (.002)	.003 (.002)	0.81
Customer Quarter 4 deal	.010 (.008)	-.036 (.024)	0.28
Controls not reported	Product line, operating system, industry, sales region		
Log psuedolikelihood	-2001.3		
Pseudo R-squared	.112		

Note: Dependent variable = timing of deal close within quarter (C_E=early, C_L=late, or C_M=middle); N=7,912; robust standard errors clustered by salesperson in parentheses.

Columns (A) and (B) report the difference in marginal effects after multinomial logit, compared to Pr (C_M).

Hypothetical Commission Δ_{t-1} and Hypothetical Commission Δ_{t+1} refer to the hypothetical commission difference had the deal closed a quarter earlier or later, respectively.

* $p < .10$.

** $p < .05$.

*** $p < .01$.

Table 5: Deal outcomes model, OLS results

Variable	(A) Baseline salesperson fixed effects model	(B) Simple deadline effects model	(C) Deadline effects model with compensation concerns	(D) Historical product pricing variance model	(E) Historical customer discount model
Constant	-2.98 (3.37)	-4.32 (2.31)*	-1.66 (1.43)	-0.96 (1.04)	-2.15 (1.76)
Hypothetical Commission Δ_{t-1} (ΔMB_{t-1})	0.26 (0.09)***		0.29 (0.10)***	0.22 (0.06)***	0.19 (0.07)***
Hypothetical Commission Δ_{t+1} (ΔMB_{t+1})	0.44 (0.16)***		0.53 (0.15)***	0.47 (0.16)***	0.34 (0.12)***
Closes early in quarter (C_E)		1.90 (0.71)***	-0.67 (0.80)		
Closes late in quarter (C_L)		3.78 (1.44)***	-1.01 (1.03)		
Historical product pricing variance				0.07 (0.04)*	
Historical Product Price Variance * ΔMB_{t-1}				-0.00 (0.01)	
Historical Product Price Variance * ΔMB_{t+1}				0.01 (0.02)	
Historical customer discount					0.32 (0.12)***
Historical Customer Discount * ΔMB_{t-1}					-0.08 (0.10)
Historical Customer Discount * ΔMB_{t+1}					0.09 (0.07)
Log deal size	11.63 (3.41)***	9.87 (3.19)***	10.54 (3.22)***	9.43 (3.16)***	10.11 (3.08)***
Log deal size squared	-0.95 (0.30)***	-0.67 (0.21)***	-1.03 (0.35)***	-0.45 (0.13)***	-0.63 (0.20)***
Customer new to vendor	-0.05 (0.10)	-0.78 (0.92)	-0.42 (0.54)	0.04 (0.13)	-0.08 (0.15)
Customer new to product	3.98 (1.45)***	4.31 (1.37)***	4.17 (1.49)***	4.64 (1.60)***	4.45 (1.53)***
Log customer # of employees	0.64 (0.31)**	0.59 (0.26)**	0.62 (0.30)**	0.57 (0.24)**	0.49 (0.23)**
Log customer revenues	-0.89 (0.69)	-0.15 (0.20)	-0.41 (0.38)	-0.39 (0.28)	-0.50 (0.37)
Customer quarter 4 deal	-0.39 (0.43)	-0.55 (0.51)	-0.27 (0.35)	-0.63 (0.58)	-0.51 (0.46)
Individual salesperson fixed effects	Y	Y	Y	Y	Y
Product line fixed effects	Y	Y	Y	Y	Y
Quarter fixed effects	Y	Y	Y	Y	Y
Customer industry fixed effects	Y	Y	Y	Y	Y
R-squared	0.390	0.375	0.394	0.399	0.411
N	7,912	7,912	7,912	7,912	4,770

Note: Dependent variable = discount received; N=7,912; robust standard errors clustered by salesperson in parentheses.

Hypothetical Commission Δ_{t-1} and Hypothetical Commission Δ_{t+1} refer to the hypothetical commission difference had the deal closed a quarter earlier or later, respectively. In the interaction terms, these variables are referred to as ΔMB_{t-1} and ΔMB_{t+1}, respectively.

* $p < .10$.

** $p < .05$.

*** $p < .01$.