The Cost of High-Powered Incentive Systems: Gaming Behavior in Enterprise Software Sales *

Ian Larkin
University of California, Berkeley

September 10, 2006

DRAFT

* I am grateful to David Mowery, Steven Tadelis, Catherine Wolfram, Bronwyn Hall, Constanca Esteves, David Reiley, Rob Seamans, Jason Snyder, Mike Toffel and participants in the Haas School of Business Innovation Seminar and Business and Public Policy Research Seminar for helpful feedback. All mistakes are my own. I thank the Ford Foundation and Harry S. Truman Foundation for graciously providing research funding for this project, and the enterprise software vendor for providing the dataset and for many useful discussions. Address for correspondence: Haas School of Business, UC Berkeley, Berkeley CA 94720-1900.
Email: larkin@haas.berkeley.edu

The Cost of High-Powered Incentive Systems: Gaming Behavior in Enterprise Software Sales

Abstract

It is well-known that employees "game" incentive systems designed to motivate effort and retain top performers, sometimes to the detriment of their own employer. However, there are very few detailed empirical studies which document the extent of this gaming, or estimate the effect of gaming on business outcomes. In this paper, I use a proprietary database of deals for a leading enterprise software vendor, together with the incentive system used to compensate salespeople, to demonstrate the scope and business effects of incentive system gaming behavior. The vendor uses a non-linear, accelerating commission schedule to compensate salespeople, which resets every financial quarter. The non-linearity and periodicity, which is not related to underlying demand, give salespeople the incentive to "stuff" deals into a single quarter to maximize their compensation, and to avoid making any sales in other quarters. I empirically demonstrate that the timing of a large set of deals in the database appear to be "gamed" based on compensation concerns of salespeople. This gaming is accomplished by granting excess discounts to customers, to go along with the salesperson's preferred timing. Using matching techniques, I compare deals that were "gamed" to very similar deals that appear not to have been, and estimate that these excess discounts cost the vendor 6-8% of revenue. Salespeople are paid an average of 8% of revenue in commissions, so this "excess discount" result suggests that motivating salespeople costs the vendor approximately twice what it may think it is spending.

1 Introduction

The use of high-powered incentives systems, which base compensation on measures of employee output rather than input, has become increasingly prevalent in recent years. Nearly all companies in the Fortune 500 grant stock options or similar compensation to senior executives, and non-salary compensation such as bonuses and options are now well over half of the average executive's compensation package (Equilar, 2006). Non-executives are often compensated based on high-powered incentives as well; in a recent survey, 95% of salespeople reported their salaries were partly based on commissions and bonuses (Joseph and Kalwani, 1998), while in high technology and service industries, it is not uncommon for companies to use bonuses or other output-based pay for every employee on the payroll (Culpepper, 2006).

The most prevalent rationale for using high-powered incentives is that they closely tie employee effort to the outputs that are considered most critical. However, it is well known that high-powered incentives can have detrimental consequences. Because contracts usually cannot specify all relevant aspects of employee behavior, employees may take unforeseen actions to increase their compensation. Such behavior, commonly referred to as "gaming," can often harm the employees' own company. Perhaps the most celebrated example of detrimental gaming is Sears' experience offering commissions to its auto mechanics based on total charges for parts and labor; mechanics predictably responded to this scheme by ordering unneeded repairs. Sears ended up settling a class-action lawsuit over excessive billing, and the compensation scheme is held to have greatly damaged the reputation of Sears' car mechanic arm (Patterson, 1992; quoted in Baker, 2000).

Gaming of high-powered incentive contracts is often exacerbated by non-linearities in the compensation scheme. For example, in a quota-based sales system, salespeople may receive no extra benefit to closing an incremental deal once their quota is reached, and therefore may not put in effort to close any new deals in that period. By definition, non-linear compensation must be based on some notion of time; almost invariably companies using such systems base them on the firm's financial reporting period. However, this period is usually irrelevant to underlying demand characteristics, so its use in the employee incentive system makes the timing of deals one potential lever for employees to game.

As demonstrated by Oyer (1998), firms selling similar products to the same customer base but with different financial reporting periods have different revenue flow and pricing patterns, with higher revenue but lower prices as the financial reporting deadline nears. This is highly suggestive of the fact that salespeople game the incentive system in which they operate, giving large discounts to customers who would otherwise buy a period later in order to meet this period's quota; Oyer terms this behavior "timing gaming." Oyer's data is at the company level;

he does not have detailed knowledge of the compensation structure for all the companies in his database. Thus, his study is therefore limited to demonstrating that companies which should have similar revenue and pricing flows do not[1]. In the conclusion to his paper, Oyer states that "a potentially informative line of future research would be to use sales data from within companies" to directly estimate the impact of non-linear, period-based incentives on revenue flow and pricing over the course of a financial period (Oyer, 1998), yet few if any empirical studies have actually gone down this path, since such data are hard to obtain.

In this paper, I directly examine the impact of a non-linear incentive system on deal timing and outcomes, using a unique database of direct sales for a leading enterprise software vendor. Like nearly all enterprise software vendors, and indeed many business-to-business (B2B) vendors in technology and other industries with quick rates of innovation, this company compensates salespeople largely on an accelerating commission scale. The salesperson's commission on a deal of a given size can vary by a factor of ten or more depending on how much revenue she has already booked in the corresponding financial period, in this case the financial quarter. Since the sales cycle in enterprise software is often well over a year, the salesperson has some degree of control over the exact quarter in which a deal closes, and is able to manipulate deal timing to maximize her compensation. The salesperson also has some control over discount decisions, which is another mechanism she can strategically use in the attempt to maximize compensation, specifically by granting bigger discounts to customers who purchase on the salesperson's preferred timing schedule.

I test two basic hypotheses in the paper. First, I examine whether the non-linear, quarter-based compensation system affects the timing of deals. I find that the difference in salesperson compensation had the deal closed a quarter earlier or later, versus compensation achieved in the actual quarter of closure, is highly predictive of when the deal closes. This suggests that "timing gaming" is prevalent. Relatedly, I demonstrate that the observed bunching of deals early and late in a financial quarter appears to stem from this "timing gaming." Second, I examine the effect of timing gaming on pricing. OLS specifications suggest that deals closing at the start or end of the quarter, which are likely gamed, receive significantly higher discounts than deals closing in the middle of the quarter. As a robustness check on these results, I use propensity score matching techniques to directly compare the discounts on deals with a similar "gaming" propensities, but where one deal was clearly not gamed (probably because the customer had tight deadlines

[1] By using exogenous changes in financial reporting periods following merger activity, Oyer convincingly demonstrates that the prevalence of higher revenue and lower pricing near the end of a financial period is uncorrelated with underlying demand. This leaves supply-side explanations, such as timing gaming, as the only plausible rationale.

around product procurement); by this method I estimate that timing gaming costs approximately 6-8% of revenue.

This paper makes several contributions. It is, to my knowledge, the first paper to use detailed internal sales, employee and incentive system data to estimate the extent and cost of gaming. In this way, it goes beyond the question of "do incentives matter" to show the extent to which they matter, in a business setting. The estimate of the effect on pricing is also novel, in that many theoretical and empirical studies assume that higher employee wages, not excessive discounts to customers, will be the cost of incentive system gaming. Many empirical studies on the perverse impact of incentives focus on illegal or clearly unethical behavior, while this study examines the negative business impact of employees doing exactly what the incentive system asks them to do: make sales. Finally, the results contain some interesting directions for further research, including the interaction between employee and executive incentive systems, the effect of employee tenure on the likelihood to game, and differences in customer participation in the gaming process.

The paper is laid out as follows. In the next section, I review the relevant literature on the use (and misuse) of incentives in organizations, and incentives in enterprise software, the empirical setting for the study. In section three, I build hypotheses on the effect of non-linear, period-based incentive systems on deal timing and pricing. In section four, I review the data, estimation strategy and empirical results. Finally, in section five I discuss these results in light of the strategic rationale for the incentive system, and briefly examine alternative incentive systems. In the final section I also review the limitations of this study and discuss potential avenues for further research.

2 The use (and misuse) of high-powered incentives: theory and evidence

This section discusses the existing theory and empirical evidence around high-powered incentives in firms, and their use in enterprise software. The use and impact of high-powered incentives on firm performance is a key strategic question facing firms, and many authors have posited that firms using more incentive-heavy compensation structures for everyday employees perform better, and have found some support for this theory in certain empirical settings (Teece, 1986; Zenger, 1994; Zenger and Lazzarrini, 2004).

Although stock options and other executive compensation issues receive considerable attention, the use of high-powered incentives goes far beyond the board room, as companies have increasingly turned to non-salary items such as bonuses, stock options and other instruments to motivate and compensate non-executive employees. The use of high-powered incentives for non-executives is most prevalent in the high technology and professional services industries, where bonuses or other performance-based measures can make up a large percentage of employee compensation in functions such as product development, research and analysis and even administrative support (Zenger, 1994; Culpepper, 2006).

2.1 The rationale for high-powered incentives: agency and other theories

Although output-based incentives[2] such as sales commissions have been prevalent for certain job functions for well over a century, their study in the academic literature is relatively recent. Williamson (1975, 1985) was among the first to discuss the positive motivational effects of tying compensation to outcomes, in that they induced effort. The groundbreaking paper by Holmstrom (1979) on tying incentives to measurable outputs established criteria for improving performance-based contracts when first-best contracts are not available, and became the cornerstone for the subsequent agency literature. Lazear (1986) added an important element to the rationale for output-based incentives: they could effectively sort among worker types, inducing workers of sufficient skill to choose to stay at the company, and those with insufficient skills to leave it. These two agency theory-based rationales – *inducing effort* and *sorting* – have become the leading explanations in the economic literature for the increasing use of output-based incentives (Lazear, 2000a).

There is a vast empirical literature supporting the notion that the use of output-based incentives for employees can have a positive effect on performance. A number of studies examine the change in a performance measure after a switch to high-powered incentives, and nearly every study finds a significant, positive effect. Examples include productivity in installing windshields (Lazear, 2000b), sales productivity for retail stores (Banker, Lee and Potter, 1996), and productivity in collective agriculture concerns in China (McMillian, Whalley and Zhu, 1989). One downside to these and similar studies is that they largely cannot distinguish between the effort and sorting effects of output-based incentives (Pregdergast, 1999).

Another prevalent theory explaining the rise of high-powered incentives is found in the organizational behavior literature: institutional theory. It holds that industry norms, existing practice, management fads and other behavioral-based explanations are critical to understanding compensation and other practices of firms (Zucker, 1987). Empirical work comparing institutional and agency theory in explaining compensation practices has largely found that the two views are complementary, not contradictory (Eisenhardt, 1988). Firms likely have both

[2] This paper uses the terms "high-powered incentives" and "output-based incentives" interchangeably. It should be noted, however, that not all output-based incentives are directly tied to measurable output. Bonuses, for example, are often based on subjective performance measures.

micro-level employee motivational issues and macro-level competitive dynamics issues in mind when deciding compensation and other key firm practices.

The use of high-powered incentives is not constant across industry or job function. In industries marked by high rates of innovation or technical change, high-powered incentives are held to be critically important to attract and retain the best employees, which is among the most critical components of successful innovation (Mansfield et al, 1971; Teece, 1986; Zenger, 1994; Brynjolfsson et al, 1993). The literature on institutional control (Eisenhardt, 1985; Slater and Olson, 2000) holds that output-based incentives are most appropriate in settings where activities are less programmed or products are less mature, which matches many high technology industries. In terms of job function, salespeople are among the most common employees to be compensated by the use of high-powered incentives. The marketing literature has extensively looked at the use of commissions and other output-based incentives for salesforce motivation, and finds a strong agency-based rationale for these practices (e.g. Basu et al, 1985; Lal and Srinivasan, 1993; Shaw et. al, 2000). Salespeople tend to be less risk averse than the average employee (Coughlan and Narasimhan, 1992), and the high prevalence of commission-based compensation for salespeople may be to induce self selection of risk loving types to the sales function[3].

2.2 The unintended consequence of incentives

The use of high-powered incentives has become so prevalent exactly because they so strongly influence actions, and therefore outcomes. However, it is well-known that not all actions or outcomes induced by high-powered incentives are intended or beneficial. One clear problem for the use of high-powered incentives arises when no measurable output neatly corresponds to the principal's goals. The basic logic and several examples of this phenomenon were given in the classic piece by Steven Kerr titled "On the Folly of Rewarding A While Hoping for B" (1975).

The literature on multitasking, inspired by Holmstrom and Milgrom (1991) and Baker (1992), recognizes that job functions are complex, and employees will opportunistically shift their effort towards those tasks that make up parts of their compensation scheme. Since these tend to be tasks with measurable outputs, too little effort may be put into tacit, but important tasks. One interesting empirical application of the multi-tasking model is Johnson, Reiley and Muñoz (2006), which shows that private bus operators in Chile, rewarded solely on the basis of their total number of passengers carried, drive at excessive speeds in the "war for fare" and cause

[3] Camerer et al (1997) make a similar argument around employee self-selection based on compensation plans in the market for taxi drivers.

a disproportionate number of accidents, as compared with drivers of state-run busses, who are paid an hourly wage.

A related stream of research examines non-linear compensation structures, particularly related to deadlines. Deadlines are usually not correlated with underlying demand for a product or service, yet employees can game the timing of a task so they maximize their compensation. In the first detailed empirical work on the topic, Asch (1990) demonstrated that Navy recruiters were very susceptible to a period-based award system the Navy used to recognize and compensate outstanding recruiters. They would strategically stockpile potential recruits until eligible for and likely to achieve an award, resulting in an unsmooth recruitment rate which was not explainable by the underlying demand to enlist.

Oyer (1998) extended this logic to the business setting, looking at revenue streams and resulting margins for companies with similar products and customers but different financial periods. In Oyer's example industries, salespeople are compensated based on a non-linear incentive scheme, most commonly with a set annual sales quota. He found that the revenue flows and margins for such companies depended on their internal financial year-end, with more revenue but worse margins as the year-end approached. This behavior, Oyer demonstrated, was consistent with "timing gaming" by salespeople, who had substantial (but incomplete) control over both deal timing and pricing. He took advantage of exogenous changes to the financial reporting period of some companies to show that demand characteristics did not explain these differences in revenue flow and pricing. Healy (1985) demonstrated that senior executives made similar decisions around revenue recognition when their compensation was non-linear and based on fiscal year firm performance.

A number of micro-level studies examine the perverse impact of non-linear incentives in more detail. Chevalier and Ellison (1997) show that mutual fund managers often make inefficient portfolio choices, because their compensation, which is a linear function of inflows, does not match the pattern of fund flows, which is non-linear at certain points on the risk/reward curve. Leventis (1997) showed that cardiac surgeons in New York, who were penalized if mortality rates exceeded a certain level annually, strategically avoided taking on risky cases.

In sum, there is considerable theoretical and empirical evidence that incentives matter tremendously in organizations, but that their impact can sometimes have unexpected consequences. Researchers have made considerable progress at understanding methods to overcome incentive problems when outputs cannot be easily measured[4]. However, economists

[4] For example, there is a significant literature on the use of subjective performance measures, which partly looks at overcoming the problems of measurability and the link between desired outcome and induced employee action. See, for example, Baker, Gibbons and Murphy (1994).

and business strategists have made "remarkably little progress" in understanding the observed use of non-linear, period-based incentive schemes (Prendergast, 1999).

2.3 Incentives in enterprise software

Enterprise software provides an ideal setting to examine the impact of non-linear incentives on business outcomes. Enterprise software, the large server- and mainframe-based applications and software infrastructure that manage and report the vast information flows corporations need to make strategic decisions, is marked by fast innovation cycles, with major product upgrades occurring every four to seven years. Major product upgrades can cost vendors as much as $3 billion in development costs, and there are virtually no marginal production costs, meaning that the post-development battle for market share is intense. Furthermore, there is very little known about the best way to manage product and business development practices (Cusumano and Selby, 1998). As predicted by the control and innovation literatures, nearly all enterprise software companies respond to these industry dynamics by making output-based pay a prevalent part of employee compensation (Gartner, 2004).

There are several reasons to focus on the sales function within enterprise software. First, the sales function is of critical importance: U.S. companies spend on average 7% of gross revenues on selling-related expenses (Godes, 2003). Secondly, salespeople tend to be compensated using objective compensation criteria based on a measurable output, which eases empirical identification.

2.3.1 Sales dynamics in enterprise software

Large enterprise software vendors sell through several channels, including large direct salesforces which tend to target the biggest corporate customers[5]. Salespeople are usually given a dedicated list of target companies to which they are assigned to sell, which are usually a mix of existing customers and potential targets to whom the company has not sold before.

Major purchases for an enterprise software installation can run into the hundreds of millions of dollars, and the average enterprise software license deal size for Fortune 1000 customers is estimated to be around $1 million (Gartner, 2004). Negotiations for new or upgraded installations typically run 12-24 months; once purchased, it can be several more years until a package in fully installed and ready for use. In addition to the software licenses, customers usually purchase maintenance and support contracts from the software vendor, and may also purchase professional services for installation, training and customization. Industry

[5] Large systems integrators also target big corporations, while smaller IT services firms and other channel partners tend to target mid-market companies.

sources report that the cost of software licenses is usually about 15-40% of the total installation price of a major project, including maintenance, hardware, installation and customization.

2.3.2 Salesforce incentives in enterprise software and hypotheses

Enterprise software vendors have developed aggressive compensation schedules in light of industry economics and competitive dynamics. Table 1 lays out a typical compensation scheme in enterprise software, loosely based on the compensation scheme used at the vendor which provided data for this research, and the author's experience working in the industry[6].

As noted, a salesperson will only make a quarterly base salary of $12,000 if she makes no sales, which is considered a "starvation wage" for the industry. More importantly, the commission she receives on incremental sales rapidly accelerates as her total sales in a quarter rise. For example, on her first $250,000 deal in a quarter, she will make a commission of 2%, or $5,000. However, if she has already closed deals totaling $6 million in a quarter, the same $250,000 deal will result in a commission of 25%, or $62,500. Depending on her revenue flow in the quarter, her sales commission on an identical deal potentially increases by an order of magnitude. As made clear by the compensation schedule, the accelerating commissions reset on a quarterly basis, meaning that all salespeople restart at the lowest commission rate for new deals at the start of every financial quarter.

At first glance, this may seem like a curious compensation system, but enterprise software executives feel that it has a compelling rationale. As predicted by Lazear (2000a), the primary rationale is sorting. Software vendors feel that there are very few good salespeople, many average ones, and even more poor ones[7]. Even if the top salesperson's output is not an order of magnitude greater than an average salesperson's, they feel her salary must be an order of magnitude bigger, or she will choose to work for competitor or even in another high tech industry. This leads to an interesting phenomenon in enterprise software and other high technology industries, where "the most highly compensated employee is often not the CEO, but the top salesperson" (Gartner, 2004). This result not only matches the "informational rent" paid to more able agents in the standard principal-agent literature (Holmstrom, 1979), but also confirms the increasingly prevalent practitioner view that top talent is the greatest bottleneck for

---

[6] It should be noted that all compensation calculations used in the proceeding econometrics use the actual compensation schedule in use at the vendor; the vendor requested that its actual compensation schedule be disguised.

[7] This phenomenon is not unique to software or even to high tech, although the quick innovation cycles and experience good nature of the product exacerbate the differences between good and bad salespeople in software and tech environments.

most organizations, and that the resulting "war for talent" results in top employees achieving pay that is much larger than their marginal products (Michaels et. al, 2002)[8].

The tying of the non-linear commission schedule to the financial quarter has two rationales. First, given the increasingly quick innovation cycles in the industry, executives feel that basing commissions on annual sales would risk missing critical upgrade cycles. Oyer (1998) shows that salespeople in annual quota systems appear to shirk for a number of months at the start of the financial year, and vendors feel they cannot risk this type of behavior. Just as importantly, the senior executives at enterprise software vendors are compensated (and retained) based largely on the stock market performance of the company's shares, which is largely tied to the ability of their companies to hit "Wall Street" financial targets, which are set quarterly.

There is one other relevant aspect of the incentive system: the degree of salesperson control over price. As is common in large B2B procurement environments with intense price competition (Bhardwaj, 2001), enterprise software vendors give salespeople a great deal of flexibility to control discounts. However, as the discount negotiated with the end customer increases, the level of authority needed to authorize the deal increases as well. Table 2 shows a disguised example of the deal approval process for the vendor which provided data for this research. The key idea is that salespeople do have a great degree of control over pricing, but the likelihood of getting a high discount approved goes down as the level of discount, and therefore the level of needed approval, goes up.

3 Theory development and hypotheses

In building a theory around "timing gaming" and the effect of this gaming on business outcomes, it is most useful to consider the factors influencing the actions of two actors: salespeople and customers. While other participants, including executives, salesforce management personnel and shareholders are relevant, I will treat the actions of these actors as exogenous when building hypotheses, and empirically control for the effects of their actions when possible.

In terms of salespeople, the key holding of agency theory is that they will take advantage of any aspect of their incentive system which will increase their compensation. If the salesperson faces a non-linear, period-based commission system, she has two mechanisms by which to increase her compensation: influencing the timing of a deal, and influencing the price paid. Influencing timing can increase compensation due to the non-linearity in compensation: making more deals happen in a single period will leave her better off than having a smooth flow of the same deals across periods. Salespeople can influence timing furtively, by deliberately

[8] The final section in this paper briefly considers other potential compensation systems which would also effectively sort; this section is only concerned with the details and rationale for the system used at the time of the study.

slowing negotiations in a "bad" quarter, or openly, by promising better deals to customers if deals are closed on the salesperson's preferred timeline. In terms of pricing, a salesperson would prefer to sell at a higher price (all else held equal) because it results in a higher commission. However, she would be willing to sell at a lower price if doing so would result in higher overall compensation, due to the non-linearity in compensation.

On the customer side, I assume utility is based on two factors: price paid, and timing of purchase. I assume that there are two customer types: one cares greatly about timing, and one has weak timing preferences. Both customer types care greatly about price. Even a customer with weak timing preferences faces disutility as a deal's timing moves away from its preferred timing.

The assumptions about customer timing preferences match anecdotal evidence in the enterprise software industry. For most projects, customers have relatively flexible timing preferences; it takes 12-24 months to negotiate a deal, and another 2-5 years to implement a software package, so timing a deal to the exact day is usually not important. However, for a subset of projects customers do have strict timing requirements. This is usually because there are exogenous budgeting or administrative deadlines which require that allocated funds be spent by a certain date, although there are instances where senior management imposes a "hard" deadline due to a perceived competitive need to start implementing a package quickly. The same customer may have different timing preferences depending on the exact project in question, so it is actually more useful to think about timing preferences over projects, not customers[9].

The assumptions about the salesperson's mechanisms to maximize compensation, and the elements of customer utility, suggest strongly that most deals will be timed, in order to maximize salesperson compensation. This is essentially because salespeople have a much stronger motivation to affect timing than customers. A salesperson who has already generated a large amount of revenue in a quarter will try to "pull" forward some deals which would otherwise close later. Conversely, a salesperson who has not or expects not to close many deals in a quarter will try to "push" out deals which would otherwise close in that quarter. Formally, I hypothesize:

> **H1: Deals for which customers do not have exact timing preferences will close in the quarter which maximizes salesperson compensation, subject to the degree to which customers are willing to speed or delay purchase.**

---

[9] Data for the enterprise software vendor show that repeat customers who have previously purchased only according to the vendor's timing, e.g. at the deadline, are just as likely to purchase later in the middle of a period. This supports the notion that project, not customer, characteristics are important in determining timing preferences.

Again, this hypothesis recognizes that customers have some underlying time preference, and the salesperson's ability to game therefore cannot stretch infinitely across quarters[10].

Hypothesis 1 focuses on the timing of deals across quarters, but a natural extension of it is that the incentive system will also affect the timing of deals within quarters. Deals that are "pulled" from later quarters will naturally close late in a quarter, as the salesperson attempts to convince the customer to purchase earlier than it would like, and so as to not inflict too large a disutility on these customers. The converse statement also would appear to have merit: that deals "pushed" from earlier quarters will naturally close early in a quarter. While it may be natural to think that a salesperson would prefer never to close a deal early to have an "option" around pushing the deal into even later quarters, it is important to remember that customers face an increasing disutility by moving their purchase date away from their preferred timing. Salespeople who expect to have a big quarter will, therefore, will push a deal so that it enters their planned big quarter, but not so far as to risk losing the deal altogether. This effect will lead to a prevalence of early deals. Formally, I hypothesize:

> **H1a: The "timing gaming" of deals will lead to a natural bunching of deals at the beginning and end of the financial period.**

Again, some deals will occur away from the beginning and end of the period, for customers with strong timing preferences.

Customers with weak timing preferences are only willing to change their preferred date if it positively affects their utility via a lower price. Of course, all else held equal, salespeople would prefer not to lower price, since it lowers their commission. However, if giving a customer a few more points off of list price helps convince the customer to purchase in a period where the salesperson has already closed many other deals, the "discount effect" on compensation is easily swamped by the "commission effect" inherent in the non-linear compensation schedule. To the extent that customers derive utility more from paying less than from having exact control over price, they will hold out for a good deal in a period where the salesperson has the incentive to discount heavily. Of course, this is again subject to the provision that the customer does not have strong preferences around exact deal timing, which is not true in all instances. From this discussion comes the paper's second hypothesis:

> **H2: Deals whose timing was strategically manipulated result in significantly higher discounts for end customers than deals whose timing was not.**

---

[10] In fact, some customers may have very exact timing preferences, due to budgetary or other constraints. This fact will be critical to the identification strategy discussed in the next section.

Again, the identification of a set of deals whose timing appears not to have been manipulated is critical to the empirical strategy for identifying this effect.

4. Data, estimation strategy and results

4.1 Data

The data for this study was provided by a leading enterprise software vendor, representing all deals closed by 175 salespeople, selected randomly from all salespeople employed by the company for at least two quarters between 1997 and 2002. In total the dataset contains 2,938 deals closed over the course of 22 financial quarters. The database excludes two types of deals booked by salespeople: deals under $50,000, which are usually small add-on purchases sold by a telephone representative and are not the result of negotiations; and "site license" deals, which give the customer the right to use as many licenses as it wishes for a particular product. Site license contracts were not available, and much of the data used later in identification is not relevant to them, since, for example, there is no notion of the level of discount granted[11]. Still, deals in the database account for nearly 90% of total direct sales revenue for the salespeople in question. The dataset was also augmented with publicly-available information on customers. The final dataset contains five classes of information:

1. Deal outcomes, which includes products bought (licenses, maintenance and services), list price, and price paid
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 across sales districts
4. Customer information, including name, number of employees, revenues, market capitalization, some information on IT use, and previous customer purchases of the vendor's products[12]
5. Deal's contribution to total quarterly compensation for the salesperson, which is the marginal commission the salesperson earned on the sale in the

[11] This study therefore does not investigate the timing of site license deals. However, I do use their contribution to commission in the quarterly measure of salesperson compensation, so the incentive effect of site license deals on deals in the database *is* taken into account.

[12] I only observe the products bought and total number of licenses for repeat customers in periods outside of the database. I do not observe pricing or discounts.

quarter in question[13]. I also calculated what the marginal commission on each sale would have been had it closed one quarter earlier and one quarter later.

One relatively unfortunate aspect of the data is that the commission schedule stayed constant over the course of the study, meaning there is not straightforward experiment utilizing changes in incentives. List prices, however, were largely increasing during the period in question, and average discounts stayed about the same, meaning overall salaries rose. Also, list prices on some products changed dramatically. This impact of this variation is an interesting source of future research.

I focus only on license revenue in the empirical analysis. Salespeople do get a commission on service revenue generated, but their commissions are a set percentage of amount sold, are therefore not based on any kind of financial period, and do not count as revenue generated towards their non-linear commission schedule used for licenses. Therefore, commissions earned on service revenue are disconnected from the incentive system under investigation here. I do control for service spend when assessing pricing on licenses, since salespeople may give better deals on licenses to customers who buy more services.

Table 3 shows summary statistics on the dataset using the deal as the unit of observation, and reveals some interesting deal characteristics. First, the deals are large, with an average size of over $850,000. Second, they are heavily discounted, with an average discount over 35%, and a large spread on discounts; some discounts reached 90%. Most tellingly, nearly 75% 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 commission of $71,000 represents a gross average commission rate of about 8%[14]. There is preliminary evidence of timing gaming, as the achieved commission on the deal of $71,000 is statistically significantly greater than the commission on the deal had it closed a quarter earlier ($64,000) or later ($62,000).

Table 4 shows summary statistics on the dataset cutting the data by salesperson-quarter, and again reveals preliminary evidence that salespeople carry out timing gaming. Salespeople make no sales at all for a full third of the salesperson-quarters in the dataset which, in combination with the prevalence of deadline deals, is difficult to rationalize given demand-side characteristics alone. Over 95% of salespeople employed for at least 4 quarters in the dataset have at least one quarter in which they booked no sales.

---

[13] For this calculation, I used the actual compensation structure used by the vendor, similar to but not the same as the schedule given in Table 1.

[14] This figure is close to that reported by Godes (2003), who stated that sales expenses were, on average, 7% of revenues for large corporations. Of course, the vendor incurs other selling expenses beyond commissions, although commissions make up the bulk of selling cost.

Even beyond the fact that 74% of deals close on the last day of the quarter, deal timing over the course of the quarter is not smooth. Table 5 breaks the financial quarter into 13 weeks, and shows the total number and percentage of total deals in the database which close in each week[15]. Except for the beginning and end of the financial quarter, there are around 40 or 50 sales in the average week. However, there is a large spike not only on the very last day in the quarter, but also in the weeks leading up to the end of the quarter. Even more strikingly, there is a spike of deals in the beginning weeks of a quarter. The final column in the table shows the average discount for deals closing in each week. For the "middle" weeks of the quarter, average discounts hover around 30%; however, both at the start and the end of the quarter, discounts rise to 35-37%.

These data appear very consistent with the timing gaming hypothesis. As the end of the quarter approaches, salespeople attempt to close as many deals as possible in order to take advantage of the convexity of their commission schedule. Since, as seen in Table 4, so many salesperson-quarters have no sales at all, it is logical to think that salespeople are likely to "pull" forward deals that would naturally close in the proceeding quarter. To do this, they are willing to trade off giving a higher discount, because the "discount effect" on their salary is swamped by the "commission effect." It is comforting that the corresponding case at the start of the quarter is evident in the data, since the incentive story is the same. Salespeople who are in the midst of a poor quarter, or who expect many deals to close the next quarter, may "push" out deals that would otherwise close a quarter earlier. Again, they are able to motivate customers to go along by giving them higher discounts. The critical insight here, which is key to the identification strategy discussed in the next section, is that the expected marginal commission *across* quarters influences when a deal is closed *within* a quarter. Middle deals are least likely to have been influenced by the periodicity in the incentive system, and are therefore most likely to be deals around which the customer had strong timing preferences.

Finally, as suggested by interviews with salespeople and customers, there are some projects for which timing is not fungible, and therefore the incentive effects on deal timing do not enter into the discounting equation. Because the vast majority of days in the quarter occur away from the deadline, these deals are most likely to fall in the "middle" weeks of the quarter, where discounts are lowest. Again, as discussed below, these deals will be critical to the identification strategy, since they arguably represent a counterfactual set of deals which are not influenced by the incentive system.

---

[15] Week 13 represents the last week in the financial quarter, except the last day. Week 14 represents the last day of the financial quarter. All the other weeks correspond to five-business day weeks.

4.2 Estimation strategy and results

Timing gaming by salespeople: modeling the deal timing decision

The hypothesis about timing gaming (H1) holds that salespeople will influence the timing of deal closure to maximize their compensation, subject to underlying customer preferences around demand. The hypothesis hinges on the assumption, confirmed in industry interviews, that most customers tend to have relatively weak preferences around the specific date a deal closes.

As noted in the previous section, an implication of Hypothesis 1 is that the difference in salesperson compensation *across* quarters when closing a potential deal will influence when a deal closes *within* a quarter. Deals that are "pulled" forward a quarter, because the salesperson expects a greater salary benefit, will much more likely close very near the quarter's end. Conversely, deals that are "pushed" out a quarter are likely to close near the start of a quarter[16]. I make an inherent assumption here that salespeople have rational expectations regarding their ability to close deals. Interviews suggest that salespeople spend significant time "planning the pipeline" and assessing the likelihood of deal closure and preferred customer timing; indeed, the vendor provides salespeople with a sophisticated sales planning tool which helps them with this task, and even allows them to compare their commissions across different scenarios for deal size and closure rates.

We can therefore 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. Formally, I model:

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

$$J \in \{E, L, M\}$$

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 to the change in marginal benefit had the deal closed a period earlier or later (notated by subscripts *t-1* and *t+1*, respectively); Ω represents a vector of controls; and ε represents the error term.

Given that I am modeling the salesperson's choice of three discrete, non-ordinal periods in which to close a deal, the most natural estimation technique is the multinomial logit. Horowitz and Savin (2001) provide a good review of the assumptions underlying this technique.

[16] It may seem more natural to believe a salesperson will avoid closing any deal early, so that she has a real option to close the deal in a quarter where her sales are actually high. However, customers do have some degree of timing preference, and salespeople also report huge time constraints at the end of quarters. Therefore, interviews suggest they often rely on expectations of a big quarter when deciding to "push" a deal to the proceeding quarter.

The most common problem in estimating this type of model, known as the Independence of Irrelevant Alternatives (IIA) assumption, is more of a concern for consumer choice models when there are a large number of closely corresponding goods from which to choose, and is likely not violated here.

Another concern around this estimation method is around the rather arbitrary definition choice for deals closed early, late or in the middle of a financial quarter. A casual glance at table 5 suggests a natural definition of "early" would be deals closing in weeks one or two; a natural definition of "late" would deals closing in weeks 12 or 13, or the very last day of the quarter; and a natural definition of "middle" would be 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, so I carry out robustness checks to ensure the results are not an artifact of the choice of timing variable.

Before turning to the results, 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. The empirical technique described above essentially sets up middle deals as a counterfactual set of deals to which early- and late-closing deals can be compared. Again, the assumption here is that deals closing in the middle of a quarter are likely to be those around which the customer has very strong timing preferences, due to budget constraints or other factors, and are therefore not subject to timing gaming. If the middle deals look substantively different than the rest of deals in the database, their robustness as a counterfactual for comparison would be questionable.

Table 6 shows the average values for key variables for early, late and middle deals as defined above. The table does not suggest that there are underlying differences in deal characteristics if selection into middle deals were not randomly assigned across deals and customers. Middle deals are not significantly smaller, nor are they sold to significantly smaller customers, than deals closing early or late in a quarter. These facts, and evidence from customer and salesperson interviews, suggest that middle deals are not dramatically different than other deals, except that customers sometimes have budgetary or other constraints which make their timing preferences tight. Notably, Table 6 gives further preliminary evidence of timing gaming, since the change in marginal benefit had the deal closed a quarter earlier goes down dramatically for early deals, while the change in marginal benefit had the deal closed a quarter later goes down for late deals. Again, supporting the assumption that middle deals close without large influence from the incentive system, the differences in marginal benefit across quarters for middle deals are not nearly as large.

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 marginal benefit had the deal closed a quarter earlier and a quarter later as the main explanatory variables, and a full set of controls. The controls include product, operating system, customer industry, and salesperson region dummies; the deal's total purchase price; salesperson tenure; and whether the deal comes in the vendor's final quarter of the fiscal year. The need for product, customer and region dummies is clear. I control for purchase price in case salespeople are averse to attempting to game larger deals, in the fear of losing them. I control for salesperson tenure to control for differences in the propensity to engage in timing gaming as salespeople become more experienced. Finally, I control for the final quarter of the vendor's fiscal year because executive pay largely depends on fiscal year-end stock price, and executives may make it more difficult for salespeople to "push" deals out of this quarter, and may motivate them to "pull" more deals into this quarter.

Table 7 reports the estimation results. The estimated coefficients in the multinomial logit specification do not carry economic meaning (Horowitz and Savin, 2001), necessitating the calculation of marginal effects at the average value of the independent variable. These values can usefully be interpreted as the change in probability of choice *j* due to a one-unit change in the value of the dependent variable. These changes in probabilities are most usefully reported in comparison to a "baseline" choice, which in this case is to close the deal in the middle of a quarter.

The results reported in Table 7 suggest a strong correlation 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 2.1% 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 5.2% greater likelihood of the deal closing late in the quarter. Both estimates are significant at standard statistical levels. Although the multinomial logit is a non-linear model and therefore marginal effects should not be extrapolated too far, it is useful to compare these figures with the average salesperson's quarterly commission of $71,100. If a deal closing a quarter earlier would reduce a salesperson's total commissions by 10% of total salary, or $7,000, the deal is 14% more likely to close early in the subsequent quarter. A 10% reduction in average salary for a deal closing in the subsequent quarter is associated with over a 35% higher likelihood of the deal closing at the deadline of the preceding quarter.

Of course, these results do not demonstrate a causal link between differences in salesperson compensation and the timing of deal closing. It could theoretically be the case that some unobserved factor in underlying demand leads to the pattern of timing observed. However, it is difficult to think of factors affecting demand which would correlate so strongly with salesperson incentives. Furthermore, interviews with customers strongly suggest that the timing of demand would be random across customers absent their ability to use the vendor's incentive system to get higher discounts, since customers have different underlying budgeting, financial and human resource constraints. 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.

The second likely 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. First, 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 goes down. Finally, it would not explain why salespeople choose not to sell at all in nearly 40% of quarters. In short, while causality cannot be definitively nailed down, it is difficult to come up with plausible alternative explanations which explain the full set of results.

Trading off discounts for the salesperson's preferred timing: modeling deal outcomes

Having demonstrated that timing gaming appears to affect the flow of deals during the course of a financial quarter, I next turn to the question of how this gaming affects outcomes. I model deal outcomes in two ways: by using OLS to directly examine the effect of deal timing within a quarter on discounts, and by matching deals with very similar "gaming" propensities and comparing discounts only on these deals. The first test has more power, but may be biased since I can only use a proxy – when the deal closed in a quarter – to represent whether or not a deal is gamed. Matching methods correct for the fact that certain deals closing at the beginning or end of the quarter are in fact those for which customers do have strong timing preferences (which happen to coincide with those periods), and certain deals closing in the middle of the quarter are in fact "gamed" (and there is some reason the usual bunching result does not hold); however, matching results in some deals being discarded, and thus has less power.

A key question for both of these approaches is the variable to use as the outcome measure. One natural outcome measure would be unit price paid per software license. However, this measure, while observed in the data, cannot be used as a dependent variable due to the complex price discrimination schemes used by enterprise software vendors. The exact same product often has hundreds or even thousands of price points, depending on the operating environment, server characteristics such as the number of processors, specific hardware used, and other IT-related variables. A typical enterprise software book contains tens of thousands of SKUs and is hundreds of pages long. This leads to huge variations in both list and achieved price for the same product, and entering a full set of product controls would quickly render any econometric test powerless.

However, there is significant evidence that customers are just as confused about list prices, and negotiate discounts, not prices. A market expert recently stated that "Discounting has long been a fixture of the enterprise software business, where list prices exist only in theory" (Riccuiti, 2004), and results of customer interviews suggest that discounts, not unit or even total price paid, is the key element of negotiations around pricing. Since the discount measure normalizes all deals and gives a direct unit of comparison, it is also useful as an outcome measure. As demonstrated by the vendor's own deal approval guidelines, the vendor itself is also oriented around thinking about pricing in terms of discounts. While differences in discount propensities obviously occur, for example for new or high-priority products, these can be easily controlled with broad product-level and other control variables. I therefore use total discount as the outcome measure in all the empirical tests on outcome.

OLS approach

In my first approach to examine outcomes, I therefore regress discount given against the deal's timing within a quarter, and a full set of controls. This implicitly uses the deal's timing within a quarter as a proxy for whether the deal was gamed; later I correct for this assumption using matching techniques. Since discounts are a continuous, linear variable, OLS estimation is the natural technique to use. Notationally, the regression equation is:

$$Y_i = \lambda_E * C_{E,i} + \lambda_L * C_{L,i} + \beta * \Omega_i + \varepsilon_i \qquad (2)$$

where the subscript *i* again refers to the deal observation, Y refers to the discount given, $C_{E,i}$ is a dummy variable equal to 1 if the deal closes in the early portion of the quarter, $C_{L,i}$ is a dummy variable equal to 1 if the deal closes in the late portion of the quarter, Ω_i is a vector of deal controls, and ε_i is the error term.

The dummy on deals closing in the middle of the quarter is the excluded variable. The key coefficients of interest are given by λ_E and λ_L, which are estimates of the effect on discounts for deals closing early or late in a quarter, controlling for all observable deal characteristics. While not all deals closing early or late are gamed, and not all gamed deals necessarily close early or late in the quarter, the results of the deal timing model suggest deals closing or late are significantly more likely to have done so as a result of gaming. Therefore this measure does have some error, but does not bias the coefficients since it is strongly correlated with the underlying variable of interest.

The important control variables are largely the same as in the deal timing model: a full set of product, customer industry, sales region and operating environment dummies; deal size; salesperson tenure; and basic customer information such as size and revenue. I also introduce quarter controls, to strip out the effect of institutional strategies about being more or less lenient on deal approval in certain quarters. (In an alternative specification, I control only for the vendor's final fiscal quarter each year, since interviews suggest this is the quarter in which deal approval policies change the most dramatically.) In addition, I introduce controls on the customer's previous purchases of the vendor's products. This is due to the intertemporal nature of consumption and the vendor's dependence on the sales of product upgrades to existing clients. 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[17]. I therefore control for whether the customer is new to the vendor and/or new to the product line in question.

Table 8 presents strong evidence that deal outcomes are correlated with deal timing. The estimates on λ_E and λ_L are easily interpreted; all else held equal, deals closing early in a quarter receive a discount of 1.96 percentage points more than deals closing in the middle of a quarter, while the corresponding estimate for late deals is 4.92 percentage points. Both coefficients are significant at the 1% level.

There are a number of other interesting results from this regression estimate. First, deals closing in the vendor's final fiscal quarter receive nearly an additional percentage point off of list. This supports the notion that the incentive system for executives, which is based on annual stock performance, affects deal outcomes. Deals which are discounted exactly at the highest level possible without seeking the next level of approval receive a *lower* discount than deals not on a band. This suggests that the deal approval system may be effective in constraining discounts to a certain extent. Finally, mirroring the results of the deal timing model, which suggested that higher-tenured salespeople are more likely to engage in timing gaming, salespeople of higher tenure appear to grant higher discounts.

Again, due to data limitations, I conduct no natural experiment, nor do I use instrumental variables to nail down causality; these results are therefore open to alternative explanations. The main alternative explanation is likely that the vendor uses financial deadlines to engage in price discrimination, punishing impatient customers by making them pay higher prices, and charging a lower price to customers who are willing to wait. Again, however, this explanation does not explain the significantly bigger discounts given to customers which close deals *early* in the

[17] 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 (2006).

financial quarter. Also, interviews with customers suggest their willingness to patiently wait for the quarter end is exactly to take advantage of the vendor's incentive system, not because they cannot get a price which they are willing to take earlier in the quarter. The "price discrimination" which is occurring does not appear to be due to demand or willingness-to-pay differences, but because both customers and salespeople enter into a mutually beneficial agreement around closing the deal to maximize salesperson utility, in return for improving customer utility by lowering price. Put another way, the price discrimination at work here does not appear to be based on customer valuation or utility differences, but simply stems from a willingness to wait, which itself is an artifact of using quarterly deadlines.

It is informative to think about the cost of the incentive system prevalent in enterprise software, in terms of foregone vendor revenue. While this exercise cannot represent a full cost-benefit analysis, it would provide a useful benchmark for the scope of sorting or other benefits the vendor feels it receives from the incentive system's use.

One high-level estimate of the cost is found in the estimates on the coefficients λ_E and λ_L given in Table 8. Timing gaming alone results in discounts nearly two percentage points bigger for "pushed" deals, and nearly five percentage points bigger for "pulled" deals. "Pushed" deals represent 9.8% of the database, while "pulled" deals make up 80.2%. This means that the average deal is discounted 4.1 percentage points than it would be had the deal closed at the customer's preferred date. However, this calculation is off of list price, not realized price; since the average deal results in revenue capture of only 64.6% of list price, the cost to the vendor in terms of foregone revenue is 4.1/.646, or 6.4% of revenue. For a vendor doing over a billion dollars in revenue a year, the estimated cost is therefore substantial.

Matching approach

It may be possible to improve on the outcome estimates from the OLS technique. For one, OLS estimation makes strong assumptions about linear effects for all control variables used. Secondly, matching can eliminate the "noise" introduced by using the timing of a deal within a quarter as a proxy for whether or not a deal is gamed. Using matching, I can directly estimate the propensity of a deal to be gamed, and discard "early" or "late" deals which appear to not have been gamed, and "middle" deals which appear to have been gamed. To do this, I match early or late deals to middle deals with similar "gaming" propensity, and discard those deals which do not have a close match.

Matching techniques are used prevalently in the program evaluation literature to test whether voluntary participation in programs results in better performance, or simply is an artifact that better performers select to participate (for a good example of this literature and a more detailed discussion of matching techniques, see Toffel, 2006). Such studies typically compare the outcome variable of a participant and the non-participant most like it in terms of propensity to participate in the program. In this study, I am interested in whether deals with very similar underlying incentive effects, but which differed in whether the incentive effects came into play due to exogenous customer differences, lead to significantly different deal outcomes, measured by discounts.

The implementation of matching in this case is in theory very simple. I have already sorted the deals into a candidate group of "control" observations, the middle deals where the incentive effect of the deadline and accelerating commissions may not have influenced timing, and two groups of candidate "treatment" observations, where the incentive system likely did affect timing for most observations. The hypothesis is that early deals would naturally close the quarter before, but are "pushed" to the next quarter. Using matching, I search for deals which closed in the middle of a quarter, which a salesperson would have been statistically just as likely to "push" out a quarter, compared to a deal that was actually pushed. In effect, I look for middle deals where the incentive effects suggest they were closed one period earlier than the salesperson would have liked. The analogous situation holds for "late" deals: I search for middle deals that show a similar statistical propensity for the salesperson to have "pulled" them forward a quarter, but where the salesperson did not do so. Again, the hypothesis is that there were exogenous, customer-driven reasons that the salesperson could not game the timing of these deals, and matching exploits this exogenous change to directly compare deals otherwise very similar in incentive effects.

This methodology therefore creates a "gaming propensity" score for all deals, by comparing the quarter in which a deal closed to the quarter which maximizes salesperson revenue. Middle deals have two "gaming" propensity scores – a propensity to be "pulled" forward a quarter, and the propensity to be "pushed" out a quarter. To operationalize this idea, I split my sample into a "push" candidate pool of early and middle deals, and a "pull" candidate pool of late and middle deals. I then run two sets of probit equations, estimating the likelihood of a deal closing in the salesperson's preferred quarter[18], as a function of the change in marginal benefit to the salesperson across the two quarters in question.

The results of these estimates are found in Table 9; beyond the statistical significance of the coefficients, this test is useful only in creating a set of "gaming propensity" scores for each deal. To find this propensity, I use the estimated coefficients from the probit model to calculate a fitted likelihood score for each deal in the dataset. Note that middle deals are associated with two separate scores: their propensity to be "pushed," and their propensity to be "pulled."

[18] For "early" and "late" deals, this is the quarter in which the deal actually closed. For "middle" deals, this is the quarter *after* the deal closed for "pushed" deals, and the quarter *"before"* the deal closed for "pulled" deals.

The next step in the exercise is to match treatment and control observations with similar propensity scores. In so doing, I create a set of directly comparable deals: one set gamed deals, and one set of "non-gamed" deals that the salesperson would have liked to game just as much as those deals in the first set.

In carrying out the matching of observations based on propensity scores, there is a tradeoff between selecting matches with the closest propensity scores regardless of underlying characteristics, and imposing ex-ante restrictions on matches so that matched pairs are similar in obvious ways, even if this means ignoring potential matches with closer propensity scores. Because of the wide disparity in products, customers and deal sizes in this study, I use the latter approach, imposing ex-ante restrictions on potential matches so that matched deals are visibly similar. The restrictions used for the reported estimates are the following: the same customer industry, same product class, and a deal size within 20%. Within this class, I match all "treatment" deals to "control" deals within 0.3 in propensity score. More and less restrictive assumptions were tried as robustness checks, without changing the basic result[19]. Deals in both samples with no match meeting the criteria, or off the common support, were excluded. I allowed replacement for both treatment and control observations, so one deal outcome could have multiple matches if it met the criteria outlined above.

For the "push" model looking at deals closing early in the financial quarter, I successfully match 114 "early" deals to 84 "middle' deals, for an average of 1.4 "treatment" deals per control" deal. For the "pull" model, I successfully match 752 "late" deals to 168 "middle" deals, for an average of 4.5 "treatment" deals per "control" deal. I next compared the key characteristics of these deals, including deal size, salesperson tenure, and customer size, to ensure that the matching did not result in deal subsets that were substantially different. T-tests on all of these variables confirmed that the matched groups' averages in each category were not significantly different. This supports the hypothesis that the matched sample of middle deals provides a valid counterfactual to deals closing in the shadow of the incentive system.

Finally, I assessed the effect on deal discounts of deal closing subject to the gaming of the incentive system, by re-estimating equation (2) on the matched samples, but replacing the deal timing dummies with a dummy representing that the deal was part of the treatment group:

$$Y_i = \delta * K_{i,} + \beta * \Omega_i + \varepsilon_i \qquad (3)$$

[19] If extremely restrictive matching criteria are used, the pool of matched deals can get so small that it is impossible to get statistical significance when comparing outcomes.

K_i is a dummy variable taking the value 1 if the deal is in the "treatment" pool[20], and the controls are the same as in equation (2). Note that an observation will show up more than once in the sample if it has more than one match. Also note that I run equation (3) on two sets of matched data: the "pushed" match sample, made up of early deals matched to middle ones, and the "pulled" match sample, made up of late deals matched to middle ones.

The estimation results are reported in table 10. The coefficient on the treatment effect is positive and statistically significant for both types of treatment deals. As expected, the matching exercise resulted in estimates which were remarkably similar to those of the original OLS model: early deals are estimated to receive discounts bigger by 2.42 percentage points (versus 1.96 percentage points in OLS), while late deals are estimated to receive discounts bigger by 6.04 percentage points (versus 4.92 percentage points in OLS). The slight rise in the discount difference is explained by the fact that the matching technique discarded some deals where timing within the quarter did not correlate with whether the deal was gamed; therefore, the matching results are arguably more robust[21]. Using the same calculation reported previously, the results of the matching exercise suggest that the total cost in foregone revenue to the vendor is 7.9% of revenue, very close to the OLS estimate of 6.4%.

5. Discussion

Many of the papers in the literature on incentives quote the old adage "Firms get what they pay for," and this research demonstrates the extent to which this adage holds true. In implementing a sales system designed to reward top performers by an order of magnitude more than the average, and to motivate large deals to cover its huge development costs, this vendor has ensured that salespeople use any means they can to close as many sales as possible in a single financial quarter. Unfortunately for the vendor, this appears to mean that a large amount of effort is put into artificially manipulating deal timing to achieve this end, rather than generating new sales to new or existing customers. Finally, as demonstrated, the cost to the vendor in terms of foregone revenue appears quite significant; it spends approximately 8% of revenue on commissions, and the incentive system appears to cost it a remarkably similar amount.

Of course, without a valid counterfactual to examine, it is impossible to say whether the system is suboptimal. Any real-life incentive system is likely to have faults, and the theory of the second best says that discrete changes in situations which appear sub-optimal may in fact lead to worse results. It is, however, revealing to compare the ~6-8% cost to the chief rationale for the incentive system cited by the vendor's executives and other industry observers: that it

[20] Mechanically, this means the deal closed in the beginning or end of a quarter.
[21] The tradeoff in this process, of course, is that using matching restricted the sample size, leading to somewhat less precise estimates.

allows vendors to attract and retain the top salespeople, in an industry environment where talented salespeople are widely held to be few and far between. The academic literature also has evidence that retention is a primary rationale for output-based salesperson incentives (Joseph and Kalwani, 1992), so it useful to briefly examine the retention performance for the vendor in the timeframe of the study.

To carry out this analysis, I ranked each salesperson employed by the company for over four quarters during the period in the dataset by total sales[22], on a yearly basis. I then broke the salespeople into quintiles based on their rank using this criterion. Finally, I examined whether the salesperson exited the company within a set number of years after attaining that performance, not counting internal promotions or other internal mobility issues as a departure.

The results of this analysis are shown in Table 11. The results suggest that the incentive system is very adept at ridding the company of the worst performers, but results in the very top performers leaving at a higher rate than the middle performers[23]. While alternative compensation systems may be even worse at retaining the cream of the crop, it is evident that this system has mixed performance in this regard. The easiest way to increase retention for top salespeople – by increasing the convexity of the commission schedule – will make the incentive system even more prone to gaming.

It is therefore worthwhile to think about alternative compensation mechanisms which have the same incentive effects in terms of sorting, but would reduce the incentives to engage in timing gaming. Many economists and other academic observers would look at the compensation structure used by this vendor and question why the vendor does not use simple linear commissions. Linearity would make the deadline irrelevant for determining compensation, so the incentive to engage in timing gaming would disappear. The problem, of course, is that a linear commission schedule also greatly reduces the sorting incentives of the scheme, in that the compensation of salespeople who sell two or three times more product would only increase by two or three times. Since there is a widespread view that top salespeople are so rare, the system must pay much more than their marginal product, and Table 11 shows preliminary evidence that even this vendor's aggressive acceleration may not be strong enough in sorting incentives.

Another candidate frequently mentioned in the theoretical and empirical literature on compensation would be to use more subjective compensation measures such as bonuses, since salespeople are so good at figuring out how to game any objective measure to their advantage. There is a substantial literature demonstrating the benefits of subjective measures (e.g. Baker,

[22] Note I used total revenue generated to rank the salespeople, not total compensation; this is because the vendor cares about the former, not the latter. Of course, the correspondence between the two measures is very close.

[23] When examining these results, the vendor was quick to note that the Internet bubble occurred during part of the timeframe in question, making it difficult for any technology company to retain workers, especially high-performing ones.

Gibbons and Murphy, 1994); relatedly, many researchers have examined the use of tournaments (e.g. Lazear and Rosen, 1981), which can be based on objective or subjective measures, to build wide variations in compensation in ways that do not motivate gaming, sabotage or other detrimental employee actions[24].

The objection to the introduction of subjective performance measures and/or formal tournaments in a sales setting is simple: salespeople do not like them. It is widely held that salespeople are used to and expect ex-ante contracts containing clear, measurable performance goals, and their salaries based on these performance goals (Churchill et. al, 1981). This likely exactly why subjective performance measures are so uncommon in sales environments, even outside enterprise software. As noted by institutional theories of control, industry norms, existing practice and employee expectations are powerful forces determining incentive and other business systems, making subjective measures an unlikely candidate for dampening the incentive to engage in gaming behavior.

If neither the clear tie to output nor the non-linearity of commissions can be changed, the final candidate is the use of financial deadlines in the incentive system. As already discussed, a non-linearity in the compensation schedule means a deadline has to enter the incentive system at some point, but firms have a clear choice regarding the period around which to base this deadline. As mentioned, the industry rationale for using such short-term deadlines for deals with such long sales cycles is two-fold. First, vendors are afraid that long-lived deadlines will cause shirking behavior. They are afraid that losing a few months of salesperson effort will lead to problems building quick momentum after the release of a major upgrade can lead to a product's failure, risking the forfeiture of a huge amount of sunk cost. Second, and perhaps more prominently, vendors are afraid they will lose organizational focus on "making the quarterly numbers," which drives the company stock price and therefore the senior executive compensation.

While it is impossible to empirically assess the downside of moving to a system with longer periods, such as one based on the financial year, intuitively it would leave fewer deals open to timing gaming. This is because customers do have underlying time preferences; a three month window is simply too short to affect those preferences for the vast majority of customers.

This research, therefore, gives vendors a useful benchmark against which to compare the costs of moving to an alternative compensation scheme, which, based on the above discussion, would likely involve lengthening the period under which the non-linear incentive system operates. It also provides a useful comparison to the supposed sorting benefits stemming from

---

[24] There are also very large literatures on tournament theory in organizations within both the social and behavioral psychology literatures.

the non-linearity in commissions; in short, the study suggests enterprise software vendors spend double what they may think they do on salesperson compensation, and they should be weighing this full cost when considering the system's benefits.

This study has a number of obvious limitations. Most prominently, it looks at a single vendor in a single institutional setting over a relatively short span of time. That said, it still goes further than essentially any other research in terms of 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 real business settings, which extend the generalizability of these findings. Also, as noted, 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's framework of analysis both to other vendors and similar industries. Secondly, the study does not involve a natural experiment or other research techniques which clearly address endogeneity and other questions around research design. It is, however, difficult to come up with alternative explanations which fit all the patterns of the data in the many empirical tests carried out by the study. Still, more work should be done on this point. One promising opportunity lies in the effect of the large variations in list price for the same product over the course of the period in question, since list price does affect revenue capture and commissions paid, and therefore incentives.

Finally, the results of the study suggest some clear avenues for further research. First, as noted, there are clear experience or tenure effects at play; the longer a salesperson works at the company, the more adept she appears to become at gaming. Understanding the causal factors for this result and better estimating its strength would greatly broaden our understanding of internal labor markets. Second, there is a clear interaction between salesperson and executive compensation structures, which at times can compete with each other. This part of the paper remains under-explored. Finally, 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. Delegation is another aspect of internal incentive systems which deserves deeper treatment.

References

Asch, B. (1990). "Do Incentives Matter? The Case of Navy Recruiters," *Industrial and Labor Relations Review*, 43 (3): 89S-106S.

Baker, G. (1992). "Incentive Contracts and Performance Measurement," *Journal of Political Economy*, 100 (3): 598-614.

Baker, G. (2000). "The Use of Performance Measures in Incentive Contracting," *American Economic Review (Papers and Proceedings)*, 90 (2): 415-420.

Baker, G., R. Gibbons and K. Murphy (1994). "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics*, 109 (4): 1125-1156.

Banker, Y., S. Lee, G. Potter and D. Srinivasan (1996). "Contextual Analysis of Performance Impacts of Outcome-Based Incentive Compensation," *Academy of Management Journal*, 39 (4): 920-938.

Basu, A., R. Lal, V. Srinivasan and R. Staelin (1985). "Salesforce Compensation Plans: An Agency Theoretic Perspective," *Marketing Science*, 8 (3): 324-342.

Bhardwaj, P. (2001). "Delegating Pricing Decisions," *Marketing Science*, 20 (2): 143-169.

Brynjolfsson, E. W. Malone, V. Gurbaxani and A. Kambil (1994). "Does Information Technology Lead to Smaller Firms?" *Management Science*, 40 (12): 1628-1644.

Camerer, C., L. Babcock, G. Loewenstein and R. Thaler (1997). "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics*, 112 (2): 407-441.

Chevalier, J. and G. Ellison (1997). "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy*, 105 (6): 1167-2000.

Churchill, G., N. Ford and O. Walker (1981). *Salesforce Management*. Irwin: Homewood, IL.

Coase, R. (1972). "Durability and Monopoly," *Journal of Law and Economics*, 15 (1): 143-49

Coughlan, A. and C. Narasimhan (1992). "An Empirical Analysis of Salesforce Compensation Plans," *Journal of Business*, 65 (1): 93-121.

Cusumano, M. and R. Selby (1998). *Microsoft Secrets: How the World's Most Powerful Software Company Creates Technology, Shapes Markets and Manages People*, Free Press: New York.

Culpepper Group (2006), *Global Compensation and Benefit Survey*, Alphretta, GA.

Eisenhardt, K. (1985). "Control: Organizational and Economic Approaches," *Management Science*, 31 (2) 134-149.

Eisendardt, K. (1988). "Agency- and Institutional-Theory Explanations: The Case of Retail Sales Compensation," *Academy of Management Journal*, 31 (3): 488-511.

Equilar Inc. (2006). *Executive Compensation Trends*. San Mateo, CA.

Gartner Group (2004). *2004 Enterprise Software Review*. Palo Alto, CA.

Godes, D. (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.

Gibbons, R. (1998). "Incentives in Organizations," *Journal of Economic Perspectives*, 13 (2): 324-340.

Hay Group (1985). "Compensation and Motivation in the Technology Industries," mimeo.

Healy, P. (1985). "The Effect of Bonus Schemes on Accounting Decisions," *Journal of Accounting and Economics*, (7) 1-3: 85-107.

Heckman, J., H. Ichimura and P. Todd (1998). "Matching as an Econometric Evaluation Estimator," *Review of Economic Studies*, 65 (2): 261-294.

Holmstrom, B. (1979). "Moral Hazard and Observability," *Bell Journal of Economics*, 10: 74-91.

Holmstrom, B. and P. Milgrom (1991). "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics and Organization*, 7: 24-52.

Horowitz, J. and N. Savin (2001). "Binary Response Models: Logits, Probits and Semiparametrics," *Journal of Economic Perspectives*, 15 (4), 43-56.

Johnson, R., D. Reiley and J. Muñoz (2006). "The War for Fare: How Driver Compensation Affects Bus System Performance," working paper, The University of Arizona.

Joseph, K. and M. Kalwani (1992). "Do Bonus Payments Help Enhance Salesforce Retention?" *Marketing Letters*, 3 (4): 331-341.

Joseph, K. and M. Kalwani (1998). "The Role of Bonus Pay in Salesforce Compensation Plans," *Industrial Marketing Management*, 27: 147-159.

Katz, M. (1991). "Game-Playing Agents: Unobservable Contracts as Precommitments," *RAND Journal of Economics*, 22 (3): 307-328.

Kerr, S. (1975). "On the Folly of Rewarding for A While Hoping for B," *Academy of Management Journal*, 18 (4): 769-783.

Lal, R. 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, I. (2006). "Product Lock-In and Innovation Cycles In Enterprise Software: An Empirical Analysis," working paper, UC Berkeley.

Lazear, E. (1986). "Salaries and Piece Rates," *Journal of Business*, 59 (3): 405-431.

Lazear, E. (2000a). "The Power of Incentives," *American Economic Review (Papers and Proceedings)*, 90 (2): 410-414.

Lazear, E. (2000b). "Performance, Pay, and Productivity," *American Economic Review*, 90 (5), 1346-1361.

Lazear, E. and S. Rosen (1981). "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, 89 (5): 841-864.

Leventis, A. (1997). "Cardiac Surgeons Under the Knife," working paper, Princeton University.

Mansfield, E. J. Rapoport, J. Schnee, S. Wagner and N. Hamburger (1971). *Research and Innovation in the Modern Corporation*, Norton: New York.

McMillian, J., J. Whalley and L. Zhu (1989). "The Impact of China's Economic Reforms on Agricultural Productivity Growth," *Journal of Political Economy*, 97 (4): 781-807.

Michaels, E., Handfeld-Jones, H. and Axelrod, B. (2002). *The War for Talent*, Harvard Business School Press: Boston.

Oyer, P. (1998). "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," *Quarterly Journal of Economics*, 113 (1): 149-185.

Oyer, P. (2000). "A Theory of Sales Quotas with Limited Liability and Rent Sharing," *Journal of Labor Economics*, 18 (3), 405-426.

Patterson, G (1992). "Distressed Shoppers, Disaffected Workers Prompt Stores to Alter Sales Commissions," *Wall Street Journal*, July 1, 1992: B1.

Prendergast, C. (1999). "The Provision of Incentives in Firms," *Journal of Economic Literature*, 37 (1): 7-63.

Ricciuti, M. (2004). "Eyes on Enterprise: SAP-Microsoft Talks Underscore Harsh Market Reality," CNET News, July 6, 2004: retrieved from http://news.com.com/Eyes+on+enterprise/2009-1001_3-5250319.html.

Shaw, J. N. Gupta and J. Delery (2000). "Emprical Organizational-level Examinations of Agency and Collaborative Predictions of Performance-Contingent Compensation," *Strategic Management Journal*, 21: 611-623.

Slater, S. and E. Olson (2000). "Strategy Type and Performance: The Influence of Sales Force Management," *Strategic Management Journal*, 21: 813-829.

Teece, D. (1986). "Profiting From Technological Innovation: Implications for Integration, Strategic Partnering, Licensing and Public Policy," *Research Policy*, 16: 285-305.

Toffel, M. (2006). "Resolving Information Asymmetries in Markets: The Role of Certified Management Programs," working paper, Harvard Business School.

Williamson, O. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*, Free Press: New York.

Williamson, O. (1985). *The Economic Institutions of Capitalism*, Free Press: New York.

Zenger, T. (1994). "Explaining Organization Diseconomies of Scale in R&D: The Allication of Engineering Talent, Ideas and Effort by Firm Size," *Management Science*, 40: 708-729.

Zenger, T. and S. Lazzarini (2004). "Compensation for Innovation: Do Small Firms Offer High-Powered Incentives That Lure Talent and Motivate Effort?" *Managerial and Decision Economics*, 25: 329-345.

Zucker, L. (1987). "Institutional Theories of Organizations," *Annual Review of Sociology*, 13: 443-464.

Table 1: Illustrative enterprise software application salesperson quarterly compensation scheme

Income source	Incremental compensation
Base salary	$ 12,000
Commissions on incremental sales	
-- on first $250,000 in sales	2% of sales (max of $5,000)
-- on next $250,000 in sales	5% of incremental sales (max of $12,500)
-- on next $500,000 in sales	8% of incremental sales (max of $40,000)
-- on next $1,000,000 in sales	12% of incremental sales (max of $120,000)
-- on next $2,000,000 in sales	15% of incremental sales (max of $300,000)
-- on next $2,000,000 in sales	20% of incremental sales (max of $400,000)
-- amount above $6,000,000	25% of incremental sales

Source: Disguised example from company providing data for this research

Table 2: Deal approval system at enterprise software application vendor

Discount requested	Approval level needed	Approximate approval rate
Up to 20%	Individual salesperson's discretion	100%
Up to 30%	District manager	90%
Up to 40%	Regional manager	75%
Up to 60%	Country head of sales	40%
Above 60%	CEO	20%

Source: Disguised example from company providing data for this research

Table 3: Deal dataset; Summary statistics for key variables, N=2,938

Variable	Unit	Mean	Std. Dev	Minimum	Maximum
Basic deal characteristics					
Total list price	$100,000	1,327	1,418	62	19,725
Total price paid	$100,000	851	704	50	7,890
Total discount given	%	35.9	12.3	5	90
Discount exactly on approval band (20%, 30%, 40%, 60%)	1=yes	0.70	0.46	0	1
Includes new product (<2 years old)	1=yes	0.11	0.32	0	1
Service spend as % of deal size	%	0.28	0.09	0	0.54
Deal timing characteristics					
Deal closed at quarter deadline	1=yes	0.737	0.44	0	1
Deal closed in last four weeks of quarter (but not day quarter ended)	1=yes	0.098	0.29	0	1
Deal closed in middle five weeks of quarter	1=yes	0.074	0.25	0	1
Deal closed in first four weeks of quarter	1=yes	0.088	0.28	0	1
Deal closed before last week of quarter	1=yes	0.232	0.41	0	1
Fourth quarter deal	1=yes	0.30	0.46	0	1
Deal signed in quarter directly after quarter where salesperson had no sales	1=yes	0.41	0.49	0	1
Salesperson characteristics					
Tenure at time of deal closing	# of quarters	12.5	8.5	1	**
Multi-salesperson deal	%	0.04	0.20	0	1
Customer characteristics					
New to vendor	1=yes	0.24	0.44	0	1
New to product	1=yes	0.66	0.47	0	1
Bought multiple products on PO	1=yes	0.11	0.32	0	1
Direct switch from competitor	1=yes	0.10	0.31	0	1
Annual revenue of customer	$ bn	19.1	20.8	**	**
Five-year cash flow change of customer	%	10.1	7.1	**	**
Compensation characteristics					
Marginal commission on deal	$1,000	71.4	112.1	1	**
Marginal commission had the deal closed a quarter earlier	$1,000	64.7	108.3	1	**
Marginal commission had the deal closed a quarter later	$1,000	61.8	107.8	1	**

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

Table 4: Salesperson-quarter panel dataset; Summary statistics for key variables, N=175 salespeople, 2,160 total salesperson/quarters

Variable	Unit	Mean	Std. Dev	Minimum	Maximum
# of employed quarters	quarter	12.34	5.16	2	22
Employed quarter w/$0 sales	1=yes	0.26	0.44	0	1
Employee employed for full dataset?	%	0.33	0.48	0	1
Average quarterly sales	$100,000	990.9	758	0	**
Average quarterly sales conditional on making at least one sale in quarter	$100,000	1,335.9	842	64	**
Average quarterly commission (does not include the $12K base)	$1,000	$71.1	118.3	0	**
Average quarterly commission conditional on making at least one sale in quarter	$1,000	$132.5	178.1	1.3	**

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

Table 5: Observed pattern of deal timing and average discount over the course of the financial quarter; N=2,938 deals

Week	# of deals closed	% of total deals	Average discount given (%)
1	94	3.2%	35.1%
2	81	2.8	34.9
3	56	1.9	32.1
4	57	1.9	29.8
5	50	1.7	30.2
6	52	1.8	28.5
7	40	1.4	27.9
8	39	1.3	29.3
9	43	1.5	30.4
10	33	1.1	27.6
11	39	1.3	29.7
12	61	2.1	33.6
13*	91	3.1	36.8
14*	2,165	73.7	37.8

Appear to be deals "pushed" into quarter from previous quarter

Appear to be deals "pulled" into quarter from next quarter

* 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 6: Average values for key variables by timing of deal closing

Variable	Unit	Early deal average (C_E)	Middle deal average (C_M)	Late deal average (C_L)
Total price paid	$100,000	831	818	857
Total discount given	%	33.4	29.2	37.3
Salesperson tenure	# of quarters	12.4	11.3	12.7
Annual revenue of customer	$bn	18.9	19.2	19.2
Marginal salesperson commission	$1,000	70.4	68.9	71.8
Marginal salesperson commission had the deal closed a quarter earlier	$1,000	59.9	66.0	65.4
Marginal salesperson commission had the deal closed a quarter later	$1,000	68.8	64.8	60.3

Table 7: Deal timing model, results after multinomial logit estimating probability of timing of deal close
Dependent variable = timing of deal close (C_E, C_L ,or C_M); N=2,938; robust standard errors in parentheses
Columns (A) and (B) report the difference in marginal effects after multinomial logit, compared to Pr (C_M); standard error of comparison in parentheses

	(A) Pr (C_E) - Pr (C_M)	(B) Pr(C_L) – Pr (C_M)	(C) X (average variable value)
ΔMB_{t-1}	-.021 (-.009)**	.009 (.005)*	6.7
ΔMB_{t+1}	.006 (.008)	-.052 (-.020)***	9.4
Log deal size	.081 (.079)	-.045 (-.050)	6.75
Salesperson tenure	.012 (.005)**	.028 (.014)**	12.5
Quarter 4 deal	.0350 (.026)	.101 (.041)**	0.30
Controls not reported	Product, operating system, industry, sales region		

Table 8: Deal outcomes model, OLS results
Dependent variable = discount received; N=2,938; robust standard errors in parentheses

Variable	(A)	(B)
Constant	8.41 (3.13)***	7.55 (3.41)**
Basic deal characteristics		
Log deal size	2.18 (0.82)***	1.99 (0.75)***
Discount exactly on band		-1.88 (-1.02)*
New product		1.51 (0.76)*
% services spend		1.04 (0.58)*
Deal timing characteristics		
C_E	1.96 (0.69)***	1.81 (0.68)***
C_L	4.92 (1.61)***	4.61 (1.60)***
4^{th} quarter deal		0.99 (0.55)*
Salesperson characteristics		
Log tenure	2.17 (0.93)**	1.79 (0.68)***
Customer characteristics		
New to vendor	-0.38 (0.31)	-0.10 (-0.16)
New to product	3.94 (1.79)**	1.58 (0.82)*
Multi-product deal		0.45 (0.36)
Log # of employees	0.40 (0.23)*	0.33 (0.17)*
Log total 2002 revenue	0.14 (0.15)	0.08 (0.07)
Product fixed effects	**Y**	**N**
Quarter fixed effects	**Y**	**N**
Customer industry fixed effects	**Y**	**Y**
Sales region fixed effects	**Y**	**Y**
Operating system fixed effects	**Y**	**Y**
R-squared	0.270	0.254

Table 9: First-stage matching results: probit to estimate the likelihood of gaming

Dependent variable = timing of deal close (C_E, C_L ,or C_M); N=2,938; robust standard errors in parentheses
Columns (A) and (B) report the difference in marginal effects after multinomial logit, compared to Pr (C_M); standard error of comparison in parentheses

	(1) Likelihood of deal being "pushed" from previous quarter	(2) Likelihood of deal being "pulled" from subsequent quarter
ΔMB_{t-1}	.042 (.020)**	
ΔMB_{t+1}		.098 (.032)***
Log deal size	.113 (.096)	.346 (.298)
Salesperson tenure	.022 (.009)*	.037 (.020)*
Quarter 4 deal	.006 (.005)	.067 (.035)*
Controls not reported	Product, operating system, industry, sales region, customer characteristics	Product, operating system, industry, sales region, customer characteristics
N =	512	2613
Wald test	567.8***	885.1***
Pseudo R-squared	.09	.13

Table 10: Deal outcomes model using matched sample, OLS results
Dependent variable = discount received; robust standard errors in parentheses

Variable	(1) "Pushed" deal matched sample	(2) "Pulled" deal matched sample
Constant	6.78 (3.21)**	7.12 (3.05)**
Sample characteristics		
Treatment dummy	2.42 (1.16)**	6.04 (2.67)**
Basic deal characteristics		
Log deal size	3.02 (1.41)**	2.15 (1.05)**
Discount exactly on band	-2.45 (-1.91)	-2.03 (-1.18)*
% services spend	0.75 (0.50)	1.40 (0.78)*
Salesperson characteristics		
Log tenure	1.67 (0.85)*	2.14 (0.78)***
Customer characteristics		
New to product	4.14 (2.31)*	2.55 (1.03)**
Log # of employees	0.28 (0.15)*	0.38 (0.18)**
Product fixed effects	**Y**	**Y**
Quarter fixed effects	**Y**	**Y**
Customer industry fixed effects	**Y**	**Y**
Sales region fixed effects	**Y**	**Y**
Operating system fixed effects	**Y**	**Y**
R-squared	0.13	0.19
N	228	1504

Table 11: Average attrition rates by performance quartile (all salespeople employed by the company for at least one year in North America; N not reported for confidentiality reasons)

Quartile	One year attrition rate	Two-year attrition rate
1 (top 20% in total revenue generated)	28%	39%
2	24%	33%
3	19%	31%
4	29%	38%
5 (bottom 20% in total revenue generated)	43%	67%