

The Paradox of Insider Information and Performance Pay

George-Levi Gayle and Robert A. Miller*

Tepper School of Business, Carnegie Mellon University

August 2008

Abstract

Empirically, managers benefit from their firm's good fortune through their compensation package, and by trading their firm's securities. This practice could easily be eliminated by the board of directors. Theoretically, managers should not profit from changes in the firm's value if there is only private information in the model. Moral hazard explains the paradox of insider information and performance pay. Shareholders permit managers to exploit hidden information in order to incentivize their work activities. The estimated benefits from designing contracts that depend on abnormal returns far outweigh projected savings in lower compensation from paying managers fixed wages.

I. Introduction

Firm executives are much better informed than shareholders about the prospects of the enterprise, and its demands on managerial time, energy and expertise. As opportunities to make the firm more profitable are explored, management gain foresight into which ventures are likely to be successful, and those which will probably fail, putting them in a favorable position to trade on their insider knowledge. If a manager could choose how many firm specific assets to hold without incurring penalties directed at those who engage in insider trading, he might prefer holding more stock and options in his own firm when his private prognosis was more favorable than the market's, and less firm specific assets when his insider knowledge projects a worse outcome than what stockholders and other investors think.

As defined by the Securities and Exchange Commission, insider trading is illegal, but their description of this activity suggests that the Commission is primarily concerned with combatting insider profits from arbitrage. In a pamphlet available on line, the SEC provides information about bounties to those who help expose insider trading:

"Section 21A(e) of the Securities Exchange Act of 1934 ("Exchange Act") [15 U.S.C. 78u-1(e)] authorizes the Securities and Exchange Commission ("") to award a bounty to a person who provides information leading to the recovery of a civil penalty from an insider trader, from a person who "tipped" information to an insider trader, or from a person who directly or indirectly controlled an insider trader. . . . "Insider trading" refers generally to buying or selling a security, in breach of a fiduciary duty or other relationship of trust and confidence, while in possession of material, non-public information about the security. Insider trading violations may also include "tipping" such information, securities trading by the person "tipped" and securities trading by those who misappropriate such information."

How the commission enforces rules against insider trading supports this view. Harris

(2003) describes how the SEC prepares to prosecute cases of alleged insider trading. Large volume transactions accompanied by big price shifts are a signal that information about the firm's prospects may have been exploited by insiders. When alerted to a possible infringement (perhaps by a trader who believes he was exploited by an insider), the SEC compiles a list of investors who traded during the period under consideration, the insiders privy to information that led to the price change, and tries to match parties from both lists.

One simple way of resolving the insider trading problem is for the board of directors to prevent the manager from ever holding any of the firm's assets. Existing regulations in the United States require the manager to frequently report all trading in the firm's assets, so this would be a relatively straightforward requirement to enforce. In the absence of moral hazard and/or the opportunity to benefit from inside trading, is hard to imagine why a manager would prefer to hold financial assets in his own firm compared to the alternative of holding a well diversified portfolio. Managers of nonprofit enterprises and high ranking government officials are routinely required to divest themselves of assets that may cause a conflict of interest between their professional role and personal wealth management. Therefore managers should have no more qualms about agreeing to such a requirement, than agreeing to rules governing company perks, or theft of company property. Thus the board could greatly curb if not entirely eliminate insider trading by issuing a simple easily enforced directive to their executive management.

In reality a large portion of executive compensation is tied to firm specific assets. From an empirical standpoint, trading by corporate insiders appears to be profitable. Seyhun (1986) finds that insiders tend to buy before an abnormal rise in stock prices and sell before an abnormal decline. Earlier studies by Lorie and Niederhoffer (1968), Jaffe (1974), and Finnerty (1976) draw similar conclusions. More recently, Seyhun (1992a) finds compelling evidence that insider trading volume, frequency, and profitability all increases significantly during the 1980s. Over the decade, he documents that insiders earned over 5 percent abnormal returns on average. Seyhun (1992b) determines that insider trades predict up to 60

percent of the total variation in one-year-ahead returns. To summarize, hidden information is an economically important phenomenon in executive compensation.

So it is paradoxical that managers are compensated on the basis of their firm's performance, such as dividends and capital gains, when the profitability of the firm partly depends on how managers assess their own accomplishments and firm's prospects. Bebchuk and Fried (2003) and others have argued that one reason why managers are paid stock options instead of assets that are easier to value, such as cash, is that shareholders underestimate option expenses. Similarly Bertrand and Mullainthan (2000, 2001) argue that the separation of management from ownership in public corporations allows the CEO to gain effective control of the pay-setting process. They argue that skimming is less likely to attract the attention of shareholders when the firm performs well. Consequently granting options should be an excellent vehicle for skimming, costing shareholders nothing when the firm performs poorly.

This paper analyzes the quantitative importance of insider information and performance pay. In Section 2 we briefly describe the data set and then regress the manager's portfolio choices on next period's abnormal returns to the firm, finding the latter are positive and significant, evidence that future returns are a noisy indicator of inside information available to the manager. To quantify the magnitudes of the insider advantage, we construct a simple dynamic portfolio strategy based on changes in asset holdings by managers, and find that this strategy significantly outperforms the market.

In the latter parts of Section 2 we investigate whether managerial compensation varies with idiosyncratic components of the return to his firm. After controlling for the manager's portfolio choices and other observed factors that affect abnormal returns, we find that unexplained variation in abnormal returns is positive and significant. We interpret this result as evidence that managers would be motivated, through their work choices, to raise the mean of unanticipated abnormal returns, if they can. These new findings suggest a second explanation for why shareholders do not prevent managers from personally exploiting their insider knowledge about the firm. If their actions are also hidden and affect firm performance, then

not linking the manager's wealth to the firm's value might create a moral hazard problem.

Section 3 takes up the idea that both insider information and moral hazard might play a role in contracting with managers. In this model shareholders do not observe the manager's activities. Contracts between shareholders and executives must satisfy three conditions, a participation constraint, that assures the manager she will have higher expected utility from employment with her firm rather than another one, an incentive compatibility constraint, that induces her to maximize the value of the firm rather than using the resources of the firm to pursue some other objective, and truth telling constraint that induces the manager to reveal his inside information.

We show that absent moral hazard, the optimal contract is to pay the manager a fixed compensation irrespective of his private information. Second, insider information is not intrinsically linked to the moral hazard problem in sense we make explicit, then although the optimal contract should depend on the firm's abnormal returns, permitting the manager to exploit her insider information is suboptimal. Third, if moral hazard and private information are intrinsically linked, then the gains from private information can be incorporated into the optimal contract. Rather than preclude insider trading, the board might optimally sanction it.

Although linking pay to performance can be rationalized within theoretical models of optimal contracting with moral hazard and hidden actions, the practical relevance of moral hazard to managerial compensation is ultimately an empirical phenomenon. In Section 4 we assess its importance by estimating the parameters of a pure moral hazard model, in this way adding to the evidence found by Margiotta and Miller (2000) and Gayle and Miller (2008a) on smaller data sets. These estimates corroborate the earlier work, showing that the losses firms would incur from paying executives a constant wage are much greater than the relatively small amount shareholders pay to incentivize them.

II. Insider Wealth, Abnormal Returns and Compensation

This section contains a brief description of the longitudinal data set compiled for undertaking the empirical work. Then we conduct a linear regression analysis of empirical evidence on insider trading. We first focus on changes in stockholding that occur before the period begins to investigate whether they help predict future returns. Using a model with a simple linear decision rule for insider trading, we test whether managers condition on more information than the market does in forming their expectations about future returns. This leads into some simulations that seek to quantify the magnitudes of the gains to managers from their insider trading opportunities. Finally we check whether, conditional on the information held by the manager, compensation to managers fluctuates with firm returns. If so, this would provide evidence of asymmetric information that goes beyond insider trading opportunities.

A. Data

Our analysis is based on longitudinal data gathered from three main sources: Standard & Poor's ExecuComp, Compustat databases, and Executive Compensation Reports data on firm compensation plan responses to Section 162(m). Our database tracks about 1,500 firms over a nine year panel beginning in 1992 in the S&P 500, Midcap, and Smallcap indices. It contains information on the six highest paid executives for 1,837 unique CUSIP identifiers. For much of our work we partitioned firms by the ten sectors labeled as Energy, Materials, Industrials, Consumer Discretionary, Consumer Staples, Health Care, Financial, Information Technology, Telecommunications Services, and Utilities.

Our data is similar to some of what we used in Gayle and Miller (2008a) to investigate how managerial compensation has changed over the last sixty years. Table 1 shows the average firm size using three measures, sales, equity and assets. The standard deviations are about twice to three times as large as the sample means. We report two measures of income, the return on assets and abnormal returns. The latter are defined for the n^{th} firm at time t as

$x_{nt} = \pi_{nt} - \pi_t$, where π_t denotes the return on the market portfolio in period t and π_{nt} is the firm's financial return. Thus x_{nt} is a relative measure that uses stock market performance as a benchmark. We also experimented with other benchmark performance measures such as industry and sector returns, but they do not significantly affect the results reported below. Note that measures of income are much more dispersed than the measures of firm size with standard deviations about twenty times the respective sample means.

TABLE 1
SUMMARY DATA ON FIRMS
(SALES, EQUITY, AND ASSETS ARE IN MILLIONS OF 2000 \$US)

VARIABLE	MEAN	STANDARD DEVIATION
Abnormal Returns	0.024	0.431
Return on Assets	1.42	25.96
Sales	3023.49	6753.51
Total Equity	1316.66	3198.6
Total Assets	3053.96	6685.625

Total compensation and one of its components, pretax salary and bonus, are summarized in Table 2. The other main components in executive compensation are stock and option grants, vested retirement benefits, as well as gains and losses from abnormal returns on stocks and other financial securities in the manager's portfolio. The reason for including the last component is that outsiders eliminate firm specific risk by holding only a negligible amount of any given firm's securities from their wealth portfolios, in this way guaranteeing the return on the market portfolio, rather a random variable distributed about that return. Thus the fact each manager is so heavily vested in his own firm indicates a professional interest that comes with his job. We report averages for the CEO, as well as the next five highest paid executives in the firm, along with the respective standard deviations.

TABLE 2

SUMMARY DATA ON EXECUTIVE COMPENSATION IN 2000 US\$			
VARIABLE	RANK	MEAN	STANDARD DEVIATION
PRETAX-TAX COMPENSATION			
	All	2,006,466	13,443,230
	CEO	2,313,259	18,456,630
	Non-CEO	1,957,755	12,471,190
PRETAX SALARY AND BONUS			
	All	699,015	656,157
	CEO	919,654	1,069,402
	Non-CEO	655,388	527,863

The main patterns in this data set are reflected in many other samples of executive compensation. The CEO receives less than half of his compensation in salary and bonus, which exhibits much lower variability than the sum of the other components, both between and within industries or sectors. Lower ranked officers are paid less than CEOs and receive a higher proportion of their pay in salary and bonus.

B. Executive portfolio choices and future returns

If managers were more informed than the market and were able to exploit this information for personal gain, it seems reasonable to conjecture that the information impounded in future returns would help predict what choices managers had already taken with respect to their wealth portfolios. In that case abnormal returns would be a noisy predictor of retrospective choices. We now denote the conditional expectation of the abnormal return at the beginning of period $t + 1$ based on all the information available to the manager in period t as:

$$u_{n,t+\Delta} \equiv E_{t+\Delta} [x_{n,t+1}]$$

where $t + \Delta$ indicates the private information the manager has at period t , and let $q_{n,t+\Delta}$ denote stock purchases by the manager in period t . Temporarily assuming that the manager's decision rule for trading is linear in this expectation, we obtain the relation:

$$q_{n,t+\Delta} = \delta_0 + \delta_1 u_{n,t+\Delta} \equiv \delta_0 + \delta_1 x_{n,t+1} + \delta_1 \varepsilon_{n,t+1}$$

where, by the definition of $u_{n,t+\Delta}$:

$$E_{t+1} [\varepsilon_{n,t+1} | u_{n,t+\Delta}] = 0$$

If we impose the additional restriction that $\delta_0 = 0$, then this decision rule may be interpreted as a linear approximation to the optimal rule for a risk averse expected utility maximizer confronted with a favorable gamble. When $\delta_0 = 0$ the rule implies that $q_{n,t+\Delta} \equiv 0$ if and only if $u_{n,t+\Delta} \equiv 0$. From the definition of $u_{n,t+\Delta}$, this is true if and only if $E_{t+\Delta} [x_{n,t+1}]$ and the unconditional expectation, $E [x_{n,t+1}] = 0$, are the same. In that case insider trading is conducted if and only if the manager has insider information about next period's abnormal return. Regressing $q_{n,t+\Delta}$ on $x_{n,t+1}$ we obtain a consistent estimator of

$$\frac{E [q_{n,t+\Delta} x_{n,t+1}]}{E [u_{n,t+1} x_{n,t+1}]} = \delta_1 (1 + E [\varepsilon_{n,t+1}^2])$$

The expression is positive if and only if $\delta_1 > 0$.

The results from running this regression are reported in Table 3. The coefficient on lead abnormal return, δ_1 , is positive and significant in the sample as predicted by this simple model of insider information. Also consistent with the simple linear model δ_0 , the constant term, is insignificant. In the same regression we also included the ratio of (contemporaneous) salary and bonus to total compensation to investigate whether the manager takes a lower salary and bonus in return for more claims that are contingent on the firms' return. Although the sign of δ_2 is negative, it is not statistically significant. The lack of significance should

not, however be interpreted as evidence against the model, since the manager is free to draw from his own outside wealth to invest in his firms' stock when promising prospects arise.

TABLE 3

REGRESSION ON CHANGES IN MANAGERS STOCK HOLDINGS		
VARIABLE	ESTIMATE	STANDARD ERROR
Ratio of Salary and Bonus to Total Compensation	-0.768	2.13
Lead Abnormal Return	2.304	1.108
Constant	80.34	50.21

Managers are required to report all their trading activity to the SEC within a month, and their reports are available for public scrutiny. Consequently our finding that managers appear to exploit inside information when investing in their own firm raises the possibility that others might be able to benefit from their serendipitous choices. Table 4 presents our findings from regressing abnormal returns on the manager's lagged trading activity, providing some evidence of how well their trading activity is a useful predictor of abnormal returns. (Dummy variables for the sectors were also included in the regression, but not reported here.) The estimated coefficients in question are positive and significant in both regressions, consistent with the hypothesis that managers exploit insider information. The estimates also show there is a negative relationship between abnormal returns of the firm and the ratio of salary and bonus to total compensation, but again the relationship is statistically insignificant, reinforcing our earlier suggestion that resources used for insider trading need not come at the expense of other components in the compensation package, but could simply reflect an adjustment in the manager's asset portfolio.

TABLE 4

REGRESSION ON ABNORMAL RETURNS		
REGRESSORS	ESTIMATE	STANDARD ERROR
Lagged Change in Manager's Stock Holdings	0.0002911	0.0000796
Ratio of Salary and bonus to Total Compensation	-0.008193	0.523
Lagged Return on Book Value of Assets	-0.0040613	0.0004682
Lagged Dividends per Share	-0.0347653	0.0094995
Lagged Return on Market Value of Equity	-0.000423	0.0000588
Lagged Earnings per Share	$3.75e - 06$	0.000115

Much of the evidence from Tables 3 and 4 supports the notion that managers exploit their superior knowledge about their own firm's performance on the stock market, but not all. As above, suppose the manager follows the linear decision rule for insider trading, and has access to the other regressors listed in Table 4, which we now call z_{nt} . In this case the inverse of the coefficient on lagged changes in the manager's stock holdings is δ_1 , and the coefficients values on all the other variables are zero because

$$x_{n,t+1} = \delta_1^{-1} q_{n,t+\Delta} - \delta_0 \delta_1^{-1} - \varepsilon_{n,t+1}$$

and the manager's forecast error satisfies the conditional expectation $E[\varepsilon_{n,t+1} | z_{nt}] = 0$. Hence, our finding that several coefficients are significant, constitutes evidence against the linear model. We also note that an estimate of $E[\varepsilon_{n,t+1}^2]$ can be obtained by subtracting 1 from the product of the estimated coefficient on $x_{n,t+1}$ in Table 3 and the estimated coefficient on $q_{n,t+\Delta}$ in Table 4. The estimated variance is negative, casting further doubt on the linear specification.

C. Gains from insider trading

To gauge the magnitude of the gains from insider trading, we conducted a simulation exercise to retrospectively evaluate how lucrative it would have been to base a portfolio investment strategy on data from these reports over the 9 year period covered by the new data set. The simulations generated the outcomes of three self financing strategies. The first strategy is an outsider strategy, to invest in the market portfolio. The third strategy is only feasible if the inside investor perfectly anticipates the one period ahead abnormal return of the companies; an investor privy to perfect inside information pertaining to the n^{th} firm invests all her wealth in its shares in period t if $\pi_{n,t+1} > \pi_{t+1}$ and all of it in the market portfolio if $\pi_{n,t+1} \leq \pi_{t+1}$, reaping a certain return for the period of

$$\pi_{n,t+1}^{(3)} \equiv \max \{ \pi_{n,t+1}, \pi_{t+1} \}$$

Note there is an upper bound to the gains from perfect foresight because it is self financing strategy after the initial outlay.

The second strategy allocates a fraction of the manager's discretionary wealth, λ_{nt} , to the market portfolio in period t , and the remaining proportion $(1 - \lambda_{nt})$ to stock in the n^{th} firm for a return of

$$\pi_{n,t+1}^{(2)} = \lambda_{nt}\pi_{t+1} + (1 - \lambda_{nt})\pi_{n,t+1}$$

where λ_{nt} reflect the historical portfolio choices of the n^{th} manager as observed in the data. Here discretionary wealth is defined as the difference between the maximum observed wealth observed the executive in the firm observed over the sample period, denoted by \overline{W}_n , and the minimum, denoted by \underline{W}_n . Thus λ_{nt} is defined by

$$\lambda_{nt} \equiv \frac{W_{nt} - \underline{W}_n}{\overline{W}_n - \underline{W}_n}$$

We compared the outcomes of these three investment strategies, to see whether following

the reports managers submit would have been profitable, and how much of the potential gains from clairvoyance managers are able to extract. The market return averaged almost 1.089 per year in this period (with standard deviation 0.097), but if an executive could have perfectly anticipated returns in her own firm, this number increases to 1.192 (with standard deviation 0.268). More surprising is our result that almost all these gains are realized by following the second strategy we defined, which produced an average annual return of 1.196 (with standard deviation 0.336).

As a final check we investigated whether the cumulative gain from following these different strategies are statistically significant from each other, by testing the null hypothesis:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \left[\prod_{t=1}^T \left(\pi_{nt}^{(i)} \right) - \prod_{t=1}^T \left(\pi_{nt}^{(j)} \right) \right] = 0$$

for various $(i, j) \in \{1, 2, 3\}$ where $\pi_{nt}^{(1)} \equiv \pi_t$ is the market return and $n \in \{1, \dots, N\}$ is the sample population of executives. The null hypothesis is strongly rejected for all the various combinations. To summarize, while perfect foresight beats everything, building an investment strategy based on the manager's stock holding is also significantly more profitable than specializing in the market portfolio.

D. Evidence for moral hazard

The evidence presented above favors the view that managers undertake insider trading, exploiting privy information to trade in their firm's stock at the expense of shareholders. We argued in the introduction that these activities are tacitly or explicitly approved by their respective boards of directors because insider trading by managers could be greatly curbed or even eliminated. Boards could require managers to refrain from owning financial assets of the firms they manage. After all certain positions in the public sector, such as elected officers, require the occupant to divest herself of assets in firms that might create a conflict of interest between his professional duties and the incentives of the firms' shareholders. One reason why boards might be reluctant to discourage insider trading is that compensation

from insider trading might help align incentives between shareholders and the manager. If so, executive compensation packages might also depend on those components of abnormal returns that are not anticipated by inside knowledge.

Recall $x_{n,t+1}$ is the abnormal return in the upcoming period $t+1$ and $\varepsilon_{n,t+1}$ is the residual of abnormal returns that the manager of firm n does not anticipate. Let $w_{n,t+1}$ denote his compensation paid at the beginning of the next period $t+1$. If insider trading does not fully resolve the conflicts of interest between shareholder and management objectives, then the board of directors should make $w_{n,t+1}$ depend positively on $\varepsilon_{n,t+1}$.

Since $\varepsilon_{n,t+1}$ is unobserved, we regressed $w_{n,t+1}$ on a estimate of $\varepsilon_{n,t+1}$, simultaneously controlling for other variables that managers use in forming their expectations about $x_{n,t+1}$. Based on the identity

$$\varepsilon_{n,t+1} \equiv x_{n,t+1} - u_{n,t+\Delta}$$

and recalling $u_{n,t+\Delta} \equiv E_{t+\Delta}[x_{n,t+1}]$ is his conditional expectation in period t about $x_{n,t+1}$, we formed:

$$\hat{\varepsilon}_{n,t+1} \equiv x_{n,t+1} - \hat{u}_{n,t+\Delta}$$

from the estimated expectation function presented in Table 4. Then we regressed $w_{n,t+1}$ on $\hat{\varepsilon}_{n,t+1}$ as well as the variables used in estimating $\hat{u}_{n,t+\Delta}$. Our estimates in Table 5 show that managers are rewarded (punished) when the unanticipated component of abnormal returns is higher (lower) than they expected. This suggests that shareholders are not only less informed about the economic prospects of their firm, but also that shareholders do not fully monitor the activities of their management. We take up this idea in the next section.

TABLE 5

REGRESSION ON TOTAL COMPENSATION			
REGRESSORS		ESTIMATE	STANDARD ERROR
UNANTICIPATED CHANGE IN ABNORMAL RETURN			
	$\hat{\varepsilon}_{n,t+1}$	725.95	88.97
INFORMATION OF MANAGER			
	lagged change in stock holdings	11.58	1.15
	lagged return on assets	16.85	6.66
	lagged dividends per share	60.39	115.17
	lagged return on equity	0.56	0.85
	lagged earnings per share	-0.95	1.65
	constant	2503.77	98.00

III. Generalized Moral Hazard

To illustrate the interactions between insider information, moral hazard and executive compensation we borrow an example analyzed by Gayle and Miller (2008b) in greater detail. After paying the manager for his work in the previous period, at the beginning of each period the board of directors proposes a compensation plan to the manager, which depends on the realization of the firms abnormal returns as well as accounting information to be later provided by the manager. Based on the board's proposal the manager decides whether to remain with the firm or leave and picks real consumption expenditure for the period. Having accepted the contract offer, the manager observes the firms prospects, provides some accounting information, and chooses a work routine that is not observed by the directors. The return on the firms assets are realized at the end of the period. It depends on how well the firm was managed during the period, the private information available to the manager, as well as other unanticipated factors. The objective of the manager is to sequentially maximize

his expected lifetime utility, and the goal of the firm is expected value maximization.

A. The model

More specifically, at the beginning of period t the manager is paid compensation denoted w_t for his work in period $t - 1$ according to the schedule the shareholders had previously committed, and his managerial contract is up for renewal. He makes his consumption choice, a positive real number denoted by c_t , and the board proposes a new contract. At that time the manager chooses whether to be engaged by the firm or be engaged outside the firm, either with another firm or in retirement. Denote this decision by the indicator $l_{t0} \in \{0, 1\}$, where $l_{t0} = 1$ if the manager chooses to be engaged outside the firm and $l_{t0} = 0$ if she chooses to be engaged inside the firm. If $l_{t0} = 0$, the prospects of the firm are then fully revealed to the manager but partially hidden to the shareholders.

We assume throughout that managers privately observe $s_t \in \{1, 2\}$ in period t , information that affects the distribution of the firm's abnormal returns. The board announces how managerial compensation will be determined as a function of $s'_t \in \{1, 2\}$, what he tells them about the firm's prospects and its subsequent performance, as measured by abnormal returns x_{t+1} revealed at the beginning period $t + 1$. The manager truthfully declares or lies about the firm's prospects by announcing $s'_t \in S$, effectively selecting a schedule $w(s'_t, x_{t+1})$ indexed by his announcement s'_t . He then makes his unobserved labor effort choice, denoted by $l_{tj} \in \{0, 1\}$ for $j \in \{1, 2\}$ in each period t . There are two possibilities, to work diligently for the firm by pursuing the shareholders objectives of value maximization, and indicated by setting $l_{t2} = 1$, or to be employed by the firm but shirk, following different objectives than maximizing the firm's value, and here denoted by $l_{t1} = 1$.

At the beginning of the period $t + 1$ abnormal returns x_{t+1} for the firm are drawn from a probability distribution which depends on the true state s_t and the manager's action l_t . We denote the probability density function for abnormal returns in period t when the manager works diligently and the state is s by $f_s(x_{t+1})$, and let $f_s(x_{t+1})g_s(x_{t+1})$ denote

the probability density function for abnormal returns in period t when the manager shirks, bounded below by the same real number ψ . Note that $g_s(x)$ is the likelihood ratio for abnormal returns from shirking versus working diligently in state s . We assume that the firm's losses from shirking does not depend on the state, meaning $f_1(x)g_1(x) \equiv f_2(x)g_2(x)$. Consequently higher profits from being in the better state can only be realized if the manager is diligent.

Preferences over consumption and work are parameterized by a utility function exhibiting absolute risk aversion that is additively separable over periods and multiplicatively separable with respect to consumption and work activity within periods. Lifetime utility is expressed as:

$$-\sum_{t=0}^{\infty} \sum_{j=0}^2 \beta^t \alpha_j l_{tj} \exp(-\rho c_t)$$

where β is the constant subjective discount factor, ρ is the constant absolute level of risk aversion, and α_j is a utility parameters with consumption equivalent $-\rho^{-1} \log(\alpha_j)$ that measures the distaste from working at level $j \in \{0, 1, 2\}$. We assume $\alpha_2 > \alpha_1$ meaning that compared to the activity called shirking, diligence is more aligned to the shareholders' interest than the manager's interests, and without loss of generality scale utility so that $\alpha_0 = 1$. This simply means that α_j values the nonpecuniary features of engaging in activity $j \in \{1, 2\}$ within the firm relative to the total utility value from leaving the firm.

The manager's wealth is endogenously determined by his consumption and compensation. We assume there are a complete set of markets for all publicly disclosed events, effectively attributing all deviations from the law of one price to the market imperfections under consideration.

B. The optimal contract

Drawing upon the work of Fudenberg, Holmstrom and Milgrom (1990), we can show that the optimal long term contract solved by shareholders can be implemented by a sequence of short term contracts. Appealing to Myerson (1982), the revelation principle implies that,

without loss of generality, we can restrict the set of feasible contracts to those that respect the participation, incentive compatibility and truth telling constraints we now define. The participation constraint states that the manager prefers working one more period and then leaving to not working for the firm at all. The incentive compatibility constraint restricts short term contracts to those payment schedules in which the manager prefers to work diligently rather than shirk. The truth telling condition requires shareholders to write contracts that induce the manager to select a compensation schedule that reveals the firm's prospects. We assume throughout that legal considerations induce the manager not to overstate the firm's prospects but that incentives must be provided to persuade the manager from understating them.

This leads to the following formulation of the optimization problem shareholders solve, to maximize:

$$\begin{aligned}
& \sum_{s=1}^2 \varphi_s \int_{\psi}^{\infty} \log v_s(x) + \eta_0 \left[\alpha_2^{1/(1-b_t)} - v_s(x) \right] f_s(x) dx \\
& + \sum_{s=1}^2 \varphi_s \eta_s \int_{\psi}^{\infty} v_s(x) \left[(g_s(x) - (\alpha_2/\alpha_1)^{1/(b_t-1)}) \right] f_s(x) dx \\
& + \varphi_2 \eta_3 \int_{\psi}^{\infty} [v_1(x) - v_2(x)] f_2(x) dx
\end{aligned}$$

with respect to $v_s(x) \equiv \exp[-\rho w_s(x)/b_{t+1}]$, which measures how utility is scaled up by compensation if abnormal returns x are realized at the end of the current period t when state s is announced, where b_{t+1} is the bond price at $t+1$, while η_0 is the shadow value of relaxing the participation constraint, η_1 and η_2 are the shadow values for relaxing the incentive compatibility constraints, and η_3 is the shadow value of relaxing the truth telling

constraint. The first order conditions for this problem are:

$$\begin{aligned} v_1(x)^{-1} &= \eta_0 + \eta_1 [(\alpha_2/\alpha_1)^{1/(b_t-1)} - g_1(x)] - \eta_3 h(x) \\ v_2(x)^{-1} &= \eta_0 + \eta_2 [(\alpha_2/\alpha_1)^{1/(b_t-1)} - g_2(x)] + \eta_3 \end{aligned}$$

where $h(x) \equiv \varphi_2 f_2(x) / \varphi_1 f_1(x)$ and φ_s denotes the probability the state is s . It is straightforward to show that $\eta_0 = \alpha_2^{1/(b_t-1)}$. If the truth telling is not binding, then $\eta_3 = 0$ and the optimization problem reduces to the pure moral hazard problem solved in Margiotta and Miller (2000). Otherwise $\eta_3 > 0$, and we substitute the first order condition into the incentive compatibility and truth telling constraints, yielding the following system of three equations to solve for the remaining three unknowns η_1, η_2 , and η_3 :

$$\begin{aligned} & \int_{\psi}^{\infty} \frac{1}{\alpha_2^{1/(b_t-1)} - \eta_3 h(x) + \eta_1 [(\alpha_2/\alpha_1)^{1/(b_t-1)} - g_1(x)]} f_2(x) dx \\ &= \int_{\psi}^{\infty} \frac{1}{\alpha_2^{1/(b_t-1)} + \eta_3 + \eta_2 [(\alpha_2/\alpha_1)^{1/(b_t-1)} - g_2(x)]} f_2(x) dx \\ 0 &= \int_{\psi}^{\infty} \frac{g_1(x) - (\alpha_2/\alpha_1)^{1/(b_t-1)}}{\alpha_2^{1/(b_t-1)} - \eta_3 h(x) + \eta_1 (\alpha_2/\alpha_1)^{1/(b_t-1)} - \eta_1 g_1(x)} f_1(x) dx \\ &= \int_{\psi}^{\infty} \frac{g_2(x) - (\alpha_2/\alpha_1)^{1/(b_t-1)}}{\alpha_2^{1/(b_t-1)} - \eta_3 + \eta_2 (\alpha_2/\alpha_1)^{1/(b_t-1)} - \eta_2 g_2(x)} f_2(x) dx \end{aligned}$$

The firm solves similar maximization problems for two of the remaining combinations of effort level, shirking in the first state but working diligently in the second, shirking in the second but not the first, and selects the value maximizing contract.

If there is moral hazard, it is easy to see from the first order conditions that compensation varies with the firm's abnormal returns, exposing the manager to uncertainty. Consequently the firm must pay a risk premium to meet the participation constraint if her compensation is uncertain and depends on the firm's abnormal returns, because the manager is risk averse.

Absent moral hazard, the optimal compensation is a fixed wage award of:

$$w_{j,t+1} = \rho^{-1}(b_t - 1)^{-1} b_{t+1} \log(\alpha_j)$$

for working at effort level j , which just offsets the alternative use of his time. Setting $j = 1$ gives the shirking contract. A rule prohibiting any trading in the firm stock is optimal in this case, and can easily be implemented if the manager's trades are publicly disclosed.

More generally, the compensation schedule should not depend on the manager's private information if, conditional on the manager's effort, the distribution of abnormal returns is independently distributed of the state. To prove this second claim, consider a model where there is only one state, by setting $\varphi_2 = h(x) = 0$. Let $w^*(x)$ denote the optimal contract for the one state model where η_1^* is the associated multiplier for the incentive compatibility constraint. Now suppose $\varphi_2 \neq 0$ but assume $f_1(x) = f_2(x)$ instead. For example the states might be revealed to the manager after he has committed to his effort level but before the end of the period when abnormal returns are realized. In this case he would still personally benefit from insider trading if he was permitted. By assumption, the shirking distributions are the same in both states, meaning $f_1(x)g_1(x) = f_2(x)g_2(x)$, so it now follows that $g_1(x) = g_2(x)$. Hence, upon inspection of the first order conditions and the solution equations for the multipliers, $w^*(x)$ is also the optimal contract for the specialization $f_1(x) = f_2(x)$, and is supported by the multipliers $\eta_1^* = \eta_1 = \eta_2$ with $\eta_3 = 0$. This demonstrates the manager should not be compensated for his hidden information in this case. Unless it is intrinsically tied to the moral hazard problem of motivating them to work diligently, in this model managers should not be allowed to trade on their inside information.

Finally if $\eta_3 > 0$, meaning the truth telling constraint is binding, it follows from arguments in Gayle and Miller (2208b), that the manager is paid more on average in the second state than he would be in the pure moral hazard case, and less in the first state. Intuitively shareholders must bribe the manager to truthfully reveal the second state when they cannot

observe it directly, in order to incentivize his effort. However this gain is offset by a penalty paid in the first state, because competition amongst executives for the position before the state is revealed equalizes their expected utility to the level attained in the pure moral hazard case. The optimal contract exposes managers to uncertainty from two sources, namely the state s , and conditional on the state, the draw of the abnormal return x . Thus the additional uncertainty increases the expected cost of the compensation package to shareholders for the higher risk premium.

IV. Estimating the Costs of Moral Hazard

Our theoretical framework demonstrates the manager does not profit from changes in the value of the firm if she signs an optimal contract unless there is a moral hazard problem. As we remarked in the introduction, the disclosure rules of the SEC make it relatively easy for boards to write contracts with managers that prohibit any trading in the firm's securities. Yet our reduced form empirical evidence shows that managers benefit significantly from their firm's good fortune. Given the risks that insider trading pose for shareholders, is moral hazard a sufficiently important economic factor for firms to incentivize managers?

To address this question we estimated the structural parameters of a pure moral hazard model, and computed the costs and benefits of incentivizing managers to their firms. Our empirical analysis applies estimation techniques for estimating parametric models of optimal contracting described in Margiotta and Miller (2000) and Gayle and Miller (2008a), extending our earlier work to all industries, rather than just a select few, in order reach conclusions about the importance of moral hazard for much broader industry spectrum of publicly traded firms. A companion paper, Gayle and Miller (2008b), analyzes identification, testing and estimation in nonparametric optimal contracting models where there is both moral hazard and private information.

A. *Parameterizing the model*

There are five parameters to account for systematic differences in executive compensation. They are the probability distribution of abnormal returns conditional on working, $f(x)$, the probability distribution of abnormal returns conditional on shirking, $f(x)g(x)$, the risk aversion parameter, ρ , the nonpecuniary benefit from shirking versus working, captured by parameter ratio α_2/α_1 , and the nonpecuniary benefit of working versus retiring or accepting employment outside the firm, α_2 .

Our empirical analysis assumes $f(x_t)$ and $f(x)g(x)$ are truncated normal with support bounded below by ψ , where (μ_2, σ^2) and (μ_1, σ^2) respectively denote the mean and variance of the parent normal distributions. Thus

$$f(x) = \frac{\exp[-(x^2 - 2x\mu_2 + \mu_2^2)/2\sigma^2]}{\Phi[(\mu_2 - \psi)/\sigma] \sigma\sqrt{2\pi}}$$

and

$$g(x) = \frac{\Phi[(\mu_2 - \psi)/\sigma]}{\Phi[(\mu_1 - \psi)/\sigma]} \exp[(\mu_2^2 - \mu_1^2 - 2x\mu_2 + 2x\mu_1)/2\sigma^2]$$

where Φ is the standard normal distribution function. As indicated in the previous section, we cannot reject the null hypothesis of restricting the mean of abnormal returns to zero conditional on working in the data, a restriction we impose in the estimation of the parameter μ_2 . This leaves the truncation point ψ , the mean of the parent normal distribution under shirking μ_1 , the common variance of the parent normal σ , the risk aversion parameter ρ , the ratio of nonpecuniary benefits from working to shirking α_2/α_1 , and the ratio of nonpecuniary benefits from working to quitting α_2/α_0 , to estimate.

The parameters of the distribution of returns are estimated separately for each sector, whereas the taste parameters α_2/α_1 and α_2 are specified as mappings of executive rank. To accommodate other factors that might affect compensation not included in the model we assumed that total compensation, denoted \tilde{w}_t , is the sum of optimal contract compensation w_t plus an independently distributed disturbance term ε_t , assumed orthogonal to the other

variables of interest.

B. Parameter Estimates

Table 6 presents the estimates of ψ_i for $i \in \{1, 2, \dots, 10\}$, the minimal abnormal return defining the lower support point of the truncated normal distribution in the i^{th} sector. The estimators are the minimum difference of the firm return and the market return across all observations in the sector. Because the estimators converge faster than the square root of sample size, so their standard errors (not reported) have no impact on the asymptotic properties of the other parameter estimates.

TABLE 6

TRUNCATION POINTS OF ABNORMAL RETURNS DISTRIBUTIONS

PARAMETER	SECTOR	ESTIMATE
ψ_1	Energy	-0.8198
ψ_2	Materials	-0.9812
ψ_3	Industrials	-2.1423
ψ_4	Consumer Discretionary	-1.4905
ψ_5	Consumer Staples	-1.0323
ψ_6	Health Care	-1.0301
ψ_7	Financial	-1.0184
ψ_8	Information Technology	-1.1362
ψ_9	Telecommunication Services	-0.8911
ψ_{10}	Utilities	-0.8097

Loosely speaking, the reported values represent the abnormal return that trigger delisting from the exchange. Our estimates suggest the points at which creditors instigate bankruptcy proceedings, are bought by private investors, or are amalgamated, differ by sector, but are dispersed around the value where the equity value of the firm is close to zero. Since the difference between the firm's financial returns and the return on the market is

identically the abnormal return it follows that if the return on a diversified portfolio was r , then an abnormal return of $-r$ would reduce shareholder value to zero, and we previously noted the return on the market portfolio over this period was 1.089.

TABLE 7

ABNORMAL RETURNS DISTRIBUTION GIVEN DILIGENCE			
PARAMETER	SECTOR	ESTIMATE	STANDARD ERROR
σ_1	Energy	0.898	0.032
σ_2	Materials	0.333	0.005
σ_3	Industrials	1.743	0.022
σ_4	Consumer Discretionary	0.626	0.006
σ_5	Consumer Staples	0.420	0.008
σ_6	Health care	42.815	0.775
σ_7	Financial	0.373	0.004
σ_8	Information Technology	1.849	0.069
σ_9	Telecommunication Services	0.579	0.029
σ_{10}	Utilities	0.289	0.004
μ_{21}	Energy	-0.5591	0.0592
μ_{22}	Materials	-0.0017	0.0003
μ_{23}	Industrials	-0.5652	0.02452
μ_{24}	Consumer Discretionary	-0.0158	0.0011
μ_{25}	Consumer Staples	-0.0087	0.0012
μ_{26}	Health Care	-1608.1984	29.0809
μ_{27}	Financial	-0.0037	0.0004
μ_{28}	Information Technology	-2.2483	0.2108
μ_{29}	Telecommunication Services	-0.0989	0.0207
μ_{210}	Utilities	-0.0024	0.0003

In Tables 7 and 8 we report estimates of the three remaining parameters $(\mu_{1i}, \mu_{2i}, \sigma_i)$ that define the truncated normal distribution for each sector $i \in \{1, 2, \dots, 10\}$. As indicated in the previous section, we cannot reject the null hypothesis of restricting the mean of abnormal returns to zero conditional on working in the data, a restriction we impose in the estimation equations. This explains why μ_{2i} is, without exception, negative and significant.

TABLE 8
SHIRKING RETURNS DISTRIBUTION

PARAMETER	DESCRIPTION	SECTOR	ESTIMATE	STANDARD ERROR
μ_{11}	Mean return from	Energy	-0.7591	0.0592
μ_{12}	shirking	Materials	-0.037	0.0033
μ_{13}		Industrials	-0.6652	0.0352
μ_{14}		Consumer Discretionary	-0.0458	0.0211
μ_{15}		Consumer Staples	-0.027	0.0312
μ_{16}		Health Care	-1901.19	40.02
μ_{17}		Financial	-0.0097	0.0024
μ_{18}		Information Technology	-4.433	0.4108
μ_{19}		Telecommunication	-0.2989	0.0307
μ_{110}		Utilities	-0.0324	0.0083

The theory predicts that the support for abnormal returns distribution conditional on shirking is contained in the support for abnormal returns conditional on working diligently. Otherwise a first best contract could be achieved by paying the manager a fixed wage supplemented with a sufficiently high fine whenever abnormal returns stray into the region outside the latter, obviating the need for incentive pay that at best can produce a second best contract. In our empirical specification, the two supports are the same, and they share a common sector specific parameter, σ_i , differing only in the mean of the parent distributions, μ_{1i} and μ_{2i} . Thus the mean of the truncated distribution for shirking is less than the mean

of the truncated distribution for working diligently if and only if $\mu_{1i} < \mu_{2i}$. Our estimates confirm this is the case in every sector, as Table 8 shows.

Abnormal returns in the health sector behave very differently than the others; although its lower truncation point is the same order of magnitude as in the other sectors, a very low μ_{26} coupled with a very high σ_6 imply the probability density function for abnormal returns in that sector is estimated to be very flat so that it can capture some high returns that occurred in some firm/years. Similarly our estimate of μ_{1i} is orders of magnitude lower than the counterparts for the other sectors, signifying an even flatter density for the shirking distribution.

The preference parameter estimates are presented in Table 9. Our estimate of ρ , the risk aversion parameter, implies utility is concave increasing as required by the model, and lies between results in reported in Margiotta and Miller (2000) and Gayle and Miller(2008a). It implies that a manager would be indifferent between accepting a lottery offering even odds of winning or losing one million dollars versus losing an amount of \$103,259 for sure. Our estimate of ξ , the parameter explaining variation not captured by the model, is similar too. (More precisely, the variance of the measurement error is $2b_{t+1}\rho^{-2}\xi$.)

We estimated α_2 , the parameter determining competitive opportunities in the labor market for executives, and α_2/α_1 , a measure of nonpecuniary benefits from shirking versus working diligently, by executive position for the top six ranked executives. All our estimates of α_2 are greater than one numerically, but only in the upper ranks is the null hypothesis that $\alpha_2 = 1$ rejected in favor of the one sided alternative $\alpha_2 > 1$. Recalling that the exponential utility function is negative, these results are weak evidence that the nonpecuniary benefits of the job relative to outside opportunities in the labor market, decline with promotion. It appears that financial remuneration, rather than power, prestige or perks, is necessary to motivate executives to climb the corporate ladder. Our estimates of α_2/α_1 show that chief executive officers would benefit significantly from taking actions that are not in the shareholder's interests if they are not incentivized, whereas the opportunities afforded to

lower ranked executives if their pay is not tied to performance are more limited, presumably because discretion about work activities and job duties increase with rank while the degree of supervision declines. These issues are investigated more thoroughly in Gayle, Golan and Miller (2008), who estimate the dynamic life cycle aspects of executive promotion and turnover.

TABLE 9

UTILITY PARAMETERS

PARAMETER	DESCRIPTION	EXECUTIVE RANK	ESTIMATE	STANDARD ERROR
ρ	Risk aversion		0.208	0.102
ξ	Measurement error		2.03	0.505
α_2	Preference for diligence relative to quitting firm	CEO	1.292	0.0162
		2nd ranked	1.523	0.126
		3rd ranked	1.420	0.118
		4th ranked	1.48	0.375
		5th ranked	1.373	0.504
		6th ranked	1.849	0.969
α_2/α_1	Preference for diligence relative to shirking	CEO	1.356	0.129
		2nd ranked	1.034	0.034
		3rd ranked	1.012	0.045
		4th ranked	1.023	0.078
		5th ranked	1.01	0.678
		6th ranked	0.987	0.567

C. The Costs of Moral Hazard

We characterize the importance of moral hazard three ways, the gross loss shareholders would incur before accounting for managerial compensation from the manager tending his

own interests, the benefits accruing to the manager from tending her own interests instead of her shareholders, and how much the shareholders are willing to pay to eliminate the problem of moral hazard altogether.

TABLE 10
GROSS LOSS TO FIRMS FROM NOT CONTROLLING MORAL HAZARD

IN MILLIONS OF 2000 \$US	
SECTOR	ESTIMATE OF τ_1
Energy	1,290
Materials	1,468
Industrials	1,679
Consumer Discretionary	1,235
Consumer Staples	987
Health Care	2,877
Financial	1,568
Information Technology	1,457
Telecommunication	1,078
Utilities	569

The first measure, denoted τ_1 , is the expected gross output loss to the firm switching from the distribution of abnormal returns for diligent work to the distribution for shirking, which has probability density function $f(x)g(x)$. Stated differently, τ_1 is the difference between the expected output to the plant from the manager pursuing the firm's goals versus his or her own, before netting out expected managerial compensation:

$$\tau_1 = -v \int xg(x) f(x) dx \equiv -vE[xg(x)]$$

this formula exploiting the identity that the expected value of abnormal returns is zero when the manager pursues the interests of the firm. Table 10 displays the estimated average over

all firms in each sector for withdrawing the incentives for the managers to work diligently. Comparing these numbers with the size of firms reported in Table 3, we find that the value of equity would decline precipitously if managers were not incentivized to align their personal objectives with those of the firms they manage. This result essentially replicates the findings of Margiotta and Miller (2000) and Gayle and Miller (2008b) for a much smaller select group of narrowly defined industries (aerospace, chemicals and electronics).

The second measure, τ_2 , is the nonpecuniary benefits to management from ongoing shirking, that is successive managers pursuing their own goals within the firm each period. Suppressing the time subscript, and supposing that the bond price b is constant, let w_2 denote the manager's reservation wage to work under perfect monitoring or if there were no moral hazard problem, that is her certainty equivalent of the current compensation package, and let w_1 denote the manager's reservation wage to shirk. Then the present value of the compensating differential for these two activities, can be expressed as the difference $\tau_2 \equiv b(w_2 - w_1)$.

Our estimates of τ_2 for the top executive position is \$24,690,192, for the second in command \$4,460,774, tiny compared to the expected losses a firm would incur; our model predicts there are enormous gains from having managers act in the interests of shareholders. (Estimates of the lower ranked executives are considerably lower than for the second in command.) From the manager's perspective, however, the annuity implied by τ_2 is quite substantial, and for a sizeable proportion of the sample population exceeds annual compensation. Of course if a manager decided to shirk to receive these sizable nonpecuniary benefits then her expected compensation would fall drastically because her inside wealth would lose much of its value.

Finally we estimated the maximal amount shareholders are willing pay to eliminate the moral hazard problem, the value of a perfect monitor. Absent moral hazard, the firm would pay the manager the fixed wage w_2 , instead of according to the compensation schedule $w(x)$. Another way of expressing τ_3 is the equilibrium risk premium paid to an executive for taking

a job that offers an uncertain income. The firms' willingness to pay for eliminating the moral hazard problem for one period is therefore $\tau_3 \equiv E[w(x)] - w_2$. We computed this measure for the CEO and second highest highest ranked executive only, but for each of the sectors separately, because the distribution abnormal returns conditional on diligence and shirking vary across the sectors.

TABLE 11

COST OF MORAL HAZARD IN 2000 \$US		
SECTOR	EXECUTIVE	ESTIMATE OF τ_3
Energy	CEO	10,450,320
	2nd ranked	1,345,098
Materials	CEO	11,450,450
	2nd ranked	1,745,067
Industrials	CEO	14,670,350
	2nd ranked	1,675,067
Consumer Discretionary	CEO	8,210,950
	2nd ranked	3,245,067
Consumer Staples	CEO	4,210,950
	2nd ranked	545,068
Health care	CEO	30,410,580
	2nd ranked	10,450,000
Information Technology	CEO	12,410,580
	2nd ranked	4,550,134
Telecommunication	CEO	15,670,892
	2nd ranked	4,550,134
Utilities	CEO	6,590,872
	2nd ranked	450,674

Confirming our previous work, the estimates in Table 11 demonstrate that the risk premium paid to executives is a very important part of their pay package. Elsewhere we have argued, in Gayle and Miller (2008b), that changes in this component are largely responsible for expected compensation and its volatility, increasing faster than real wages over the last 60 years. Here we simply add that there is notable variation between the costs of moral hazard across the sectors, with the health care sector registering as an outlier worthy of special attention in a future study.

V. Conclusion

The disclosure rules of the SEC make it relatively easy for boards to write contracts with managers that prohibit any trading in the firm's securities. Yet our empirical evidence shows that managers significantly benefit from their firm's good fortune. Consistent with previous work, in this area we find that managers exploit insider information about the profitability of their own firm for direct personal gain. But this is not sufficient to prove that executive compensation contracts are defective. Our theoretical framework demonstrates that in an optimal contract the manager should not profit from changes in the value of the firm unless there is a moral hazard problem. This feature might explain the paradox of inside information and performance pay. Optimal contracting in models of generalized moral hazard with both private information and hidden actions reward managers for truthfully revealing the state of the firm. Shareholders permit compensation schemes that correlate firm performance with executive pay because the profitability of the firm depends on how managers assess their own accomplishments and firm's prospects, as well as what managers do, which is organizing human resources in creative ways that add value to their firm. Rewarding managers for revealing hidden information about the firm's profitability helps the board set contracts that incentivize the manager's work activities. If moral hazard is anywhere near as costly as our estimated values, then de-coupling managerial compensation from changes in shareholder wealth would be very costly indeed.

REFERENCES

- Bebchuk, L., and J. Fried (2003), "Executive Compensation as an Agency Problem", *Journal of Economic Perspectives* **17**, 71-92.
- Bertrand, M., and S. Mullainathan (2000), "Agents with and without Principals", *American Economic Review* **90**, 203-208.
- Bertrand, M., and S. Mullainathan (2001), "Are CEOs Rewarded for Luck? The Ones without Principals Are", *The Quarterly Journal of Economics* **116**, 901-932.
- Fudenberg, D., B. Holmstrom and P. Milgrom (1990) "Short-Term Contracts and Long-Term Agency Relationships", *Journal of Economic Theory* **51**, 1-31.
- Finnerty, J. (1976), "Insiders and Market Efficiency", *The Journal of Finance* **31**, 1141-1148.
- Gayle G., L. Golan and R. Miller (2008) "Promotion, Turnover and Compensation in the Executive Market", Working paper, Tepper School of Business, Carnegie Mellon University.
- Gayle, G., and R. Miller (2008a), "Has Moral Hazard Become a More Important Factor in Managerial Compensation?", *American Economic Review*, forthcoming.
- Gayle, G., and R. Miller (2008b), "Identifying and Testing Generalized Moral Hazard Models of Managerial Compensation", Working paper, Tepper School of Business, Carnegie Mellon University.
- Harris, L. (2003), *Trading and Exchanges: Market Microstructure for Practitioners*, Oxford University Press, Oxford.
- Jaffe, J. (1974) "Special Information and Insider Trading", *Journal of Business* **47**, 410-428.
- Lorie, J., and V. Niederhoffer (1968), "Predictive and Statistical Properties of Insider Trading", *Journal of Law and Economics* **11**, 35-53.

Margiotta, M., and R. Miller (2000), "Managerial Compensation and the Cost of Moral Hazard ", *International Economic Review* **41**, 669-719.

Myerson, R. (1982), "Optimal Coordination Mechanisms in Generalized Principal-Agent Models", *Journal of Mathematical Economics* **10**, 67–81.

Seyhun, H. (1992a), "The Effectiveness of Insider Trading Sanctions", *Journal of Law and Economics* **35**, 149-182.

Seyhun, H. (1992b), "Why Does Aggregate Insider Trading Predict Future Stock Returns?", *The Quarterly Journal of Economics* **107**, 1303-1331.

Notes

*We would like to thank seminar participants, at SOLE 2006, Econometric Society summer meetings 2006, Carnegie Mellon University, Cornell University, Duke University, University of Essex, Georgia State University, London School of Economics University of Maryland, New York University, University of Pennsylvania, Princeton University and the Venice Summer Institute Workshop on Executive Pay.