

Firm Performance and Executive Compensation in the Savings and Loan Industry

Benjamin E. Hermalin and Nancy E. Wallace¹
Walter A. Haas School of Business
University of California at Berkeley

October 30, 1997

¹The authors thank Paul Joskow, Canice Prendergast, Mike Riordan, Nancy Rose, Jeff Zwiebel, and seminar participants at Berkeley, UCLA, NBER, and the AFA meetings for their helpful comments on earlier drafts. Although the first author is always grateful to Albert Hermalin, both authors thank him for his suggestions concerning an estimation strategy. The authors thank the Population Studies Center at the University of Michigan for supplying the computer program, *GENMOD*, used in estimation. Financial support from the University of California Committee on Research, the Walter A. Haas School, the Fisher Center for Real Estate, and the Harold Furst Chair in Management Philosophy and Values is gratefully acknowledged.

Abstract

Previous empirical analyses of the relationship between executive compensation and firm performance are often interpreted as suggesting that this relationship is weak. Although an absolute term like “weak” is ambiguous in this context, relative terms, such as “stronger,” are meaningful. We argue that a stronger relationship can be found if a more appropriate specification is used in estimation. Specifically, an implicit assumption in the previous literature is that all firms use the same compensation scheme. Theoretically, this is a difficult assumption to accept. Moreover, we show that it is rejected empirically as well. When we allow different firms to use different compensation schemes, we indeed find a relationship between executive pay and firm performance that is about 2.8 times larger than that found using previous methods.

1 Introduction

Agency theory predicts that a manager’s compensation should be positively correlated with his firm’s performance.¹ What agency theory does not do, however, is tell us how strongly correlated. Consequently, a sizable empirical literature has emerged that seeks to measure this relationship (see Gibbons, 1997, for a survey). The most notable example of which is arguably Jensen and Murphy (1990), which found that CEO compensation in many industries seemed fairly insensitive to firm performance (there is statistically significant positive correlation but the coefficients are “small”). This has also been found in papers that have focussed on regulated firms (Carroll and Ciscel, 1982; Joskow *et al.*, 1993). As a consequence of these studies, some have argued that pay for performance is too small to be economically significant. Others (e.g., Haubrich, 1994) have argued that Jensen and Murphy’s results are actually consistent with plausibly parameterized principal-agent models. Our purpose in this paper is not to take sides in this debate—an absolute term like “small” is at best ambiguous in this context—but rather to argue that the underlying “facts” could be misleading due to potential econometric mis-specification.

One potential failing of existing empirical analyses of the pay-performance relationship is that their specifications do not sufficiently account for inter-firm heterogeneity, which, in theory at least, could be an important determinant of the pay-performance relationship.² Our focus, therefore, is on controlling for heterogeneity in an empirical analysis of the pay-performance relationship. We begin the next section by estimating the “standard” specification of a pay-on-performance regression using a panel of 86 publicly traded savings and loans (thrifts). This regression treats the coefficients on the performance measures as being the same across firms; that is, it assumes that all firms use precisely the same compensation scheme. For a number of reasons, this is a suspect assumption. More importantly, it could result in a downward bias on the estimated coefficients. To correct this, we develop an alternative specification that does not treat these coefficients as the same across firms, but rather allows them to vary based on variables that control for inter-firm heterogeneity.

A potential way to control for this inter-firm heterogeneity would be to interact the heterogeneity terms with the performance measures in an OLS regression of compensation on performance. However, since the heterogeneity terms determine the compensation package,

¹To be precise, not all agency models predict that rewards will be an everywhere strictly increasing function of performance (although they cannot be an everywhere decreasing function). See Grossman and Hart (1983). On the other hand, non-increasing functions are, as a practical matter, unlikely: if there were a region over which rewards and performance were negatively correlated, then this would generate a perverse incentive for the agent to sabotage the performance measure in this region (e.g., reduce profits or destroy output).

²Garen (1994) makes a similar point.

which in turn determines performance, there is reason to believe that the right-hand side variables in such an OLS regression would be correlated with the error term, yielding biased and inconsistent coefficient estimates. Consequently, in Section 3, we estimate this alternative specification using a consistent estimator developed by Wong and Mason (1991). We find that controlling for inter-firm heterogeneity is important: The null hypothesis that these controls don't matter is readily rejected. We then proceed to calculate the coefficients on the performance measures for each firm. On average, they are positive, and both economically and statistically significant. For instance, taking as a measure of economic significance the change in the average CEO's compensation from a one-standard-deviation change in a performance measure, our method yields estimates of these changes that are about 2.8 times greater than estimates that do not fully account for heterogeneity..

In Section 4, we provide an alternative test of our hypothesis that variables that describe inter-firm heterogeneity help to determine firms' incentive schemes. We show that these variables are significant predictors of the inclusion or omission of various *contractual* features in the CEO's compensation package (e.g., whether it includes an incentive plan). Since these features, in turn, determine the sensitivity of pay to performance, we view this as corroborating our findings from Section 3.

Although restricting attention to firms in a single industry helps reduce inter-firm heterogeneity, there are, nonetheless, several potential drawbacks to focussing on the savings and loan industry *per se*. One is that previous studies of the pay-for-performance relationship have typically found the weakest results for regulated firms (Joskow *et al.*, 1993). Thus, looking at thrifts may impose, in some sense, a more severe test of the pay-for-performance relationship than may be true for unregulated industries. Thrifts are also highly leveraged, which gives rise to a number of confounding factors. First, a number of authors (see, e.g., Jensen, 1986) have suggested that leverage itself is an incentive scheme—the fear of losing their jobs if their firms go bankrupt spurs executives to do better. Second, regulatory oversight could substitute for direct incentive schemes, at least with regard to some executive behavior. Third, high leverage and deposit insurance conspire to turn the shareholders into owners of a put option: If the thrift is doing poorly, then they can “exercise” the option by forcing the deposit insurer to “buy” the thrift, but if it is doing well, they can let the option “expire”; that is, realize profits.³ Otherwise risk-neutral shareholders will, therefore, have risk-loving preferences; preferences that they will attempt to impart on their execu-

³Merton (1974, 1977) establishes the equivalence between the insurance guarantee and a common stock put option. Using the Black-Scholes pricing formula for a put and the assumption that insured deposits will be riskless, he derives an explicit formula for the cost of the guarantee per dollar of insured deposits. He then shows that the change in cost with respect to an increase in the deposit-to-asset ratio is positive; that is, the lower is tangible capital, the greater is the value of the put to the owners of the thrift.

tives through the compensation packages they offer.⁴ Consequently, an executive could be rewarded more for taking a lower-expected-profit-but-greater-risk course of action than if he took a higher-expected-profit-but-lower-risk course of action. Unless care is taken, this could lead one to find a negative relationship between pay and performance measured by profits. Fortunately, we are able to overcome these problems to a large extent. Using series on the dividend adjusted returns on stock prices, we obtain reasonable measures of the market riskiness of thrifts using CAPM beta.

We conclude in Section 5.

2 Reconsidering Principal-Agent Theory and the Estimation of the Pay-Performance Relationship

A possible reason for the small elasticities found in previous empirical tests of the pay-performance relationship may be that the models do not sufficiently account for inter-firm heterogeneity. To appreciate this criticism, consider the typical regression equation in this literature:

$$y_j = \mathbf{x}'_j \boldsymbol{\beta} + \mathbf{z}'_j \boldsymbol{\gamma} + \varepsilon_j, \quad (1)$$

where j indexes the observations, y is a measure of compensation (e.g., the log of salary), \mathbf{x} is a vector of performance measures (e.g., change in stock price, returns on assets, etc.), \mathbf{z} is a vector of other controls (e.g., industry, firm size, etc.), ε is the error term, and $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are the regression coefficients to be estimated. The test of the theory is, then, to see whether the coefficients on the performance measures are positive (i.e., $\boldsymbol{\beta} > \vec{0}$) and significant in both a statistical and economic sense.

Table 1 presents our estimations of (1).⁵ Our dependent variable is the natural log of the CEO's cash compensation (including salary, bonuses, and realized options) in the year *following* the year in which his or her performance is measured. The performance measures

⁴John and John (1993) develop a model in which shareholders try to convince would-be bondholders that they will *not* seek to maximize the value of their option (i.e., take on too much risk) by employing an executive under a contract that does not encourage risk taking. This model, however, relies on the assumption that executive contracts can be used as a commitment device by the shareholders because they are fully observable and non-renegotiable. As Katz (1991) persuasively argues, this assumption is unrealistic. Moreover, even if this were a credible model in general, it would not be applicable in our context: Deposit insurance makes would-be bondholders—that is, depositors—essentially indifferent to how risky a thrift's operations are.

⁵We use White standard errors to control for possible heteroskedasticity (see, e.g., Weber, 1986, for details).

Table 1: Fixed-Effects Estimates of the Compensation Performance Relationship

	Log(Cash Compensation) _{t+1}
Independent Variables	Coefficient Estimate (Standard Error)
Log of Assets	.2207** (.0841)
ROA (Rate of Return on Assets)	.0043** (.0018)
Percentage Change in Stock Price from Previous Year	.0008** (.0003)
R ²	.95
Test for firm-fixed effects, $F_{85,171}$	53.859**

*Significant at the 10% level

** Significant at the 5% level

*** Significant at the 1% level

(\mathbf{x}_j) are the rate of return on assets (ROA) and the annual percentage change in the stock price. The other control (\mathbf{z}_j) is the natural log of assets.⁶ To control for all other unspecified firm-level sources of heterogeneity we estimate and test for a fixed effects version of the model. We find that the fixed firm-level effects are jointly different from zero at the .01% level. We, therefore, report the results for the fixed effects model in Table 1. There are 86 thrifts in our sample and 430 thrift-years (the panel is balanced).

The performance measures are positive and significant in predicting salary, which is consistent with theory and earlier studies (Ciscel and Carroll, 1980; Barro and Barro, 1990). What is disputable is whether these coefficients are *economically* significant. In particular, a one-standard-deviation change in ROA shifts the average CEO's salary by \$4478 and a similar change in the stock-price performance measure shifts it by \$6548.

The validity of the test embodied by (1) relies critically on the maintained hypothesis that the coefficients are identical across the observations; or, since different observations typically refer to different firms, that they are identical across firms. We have strong doubts about this hypothesis.

⁶See the Data Appendix for details on these variables.

Firstly, we know from principal-agent theory that the sensitivity of pay to performance depends on the particulars of the agency problem. For instance, compare two firms, A and B. Suppose that A's profits (a measure of firm performance) are a noisy signal of its CEO's performance (e.g., A's demand is exceptionally volatile), while B's profits are a more informative signal of its CEO's performance (e.g., its demand is stable). As a consequence, the relationship between CEO pay and firm performance in firm A will almost surely differ from the relationship in firm B; that is, the two firms' β s will differ.⁷ Alternatively, suppose that a large part of the CEOs' efforts are spent controlling costs, so that the more effort a CEO expends, the lower his or her firm's marginal costs will be in expectation. Suppose, too, that A produces more output than B. Consequently, the benefit of CEO effort is greater in A than in B. Correspondingly, A will induce more CEO effort than B, with the result that A's CEO faces a more sensitive pay schedule than will B's CEO; that is, A's β s will exceed B's β s (see Hermalin, 1992). It is straightforward to think of more reasons why two firms might adopt different compensation schemes.

Secondly, even if we were to make the heroic assumption that all the firms in the sample faced identical situations, this would *not* be sufficient to ensure that they use the same incentive contracts. Hermalin (1994) shows that equilibria exist in which otherwise identical firms employ their CEOs under different contracts. To illustrate the intuition for this, consider two Cournot duopolists, A and B. As in the previous paragraph, assume that CEO effort reduces marginal costs in expectation. Suppose that B anticipates that A will give its CEO strong—i.e., highly sensitive—incentives to expend effort on reducing costs. In determining its best response, B would reason as follows: “A is likely to have low marginal costs, so, *ceteris paribus*, it can be expected to have more output. Since we are Cournot competitors, this means I will want to produce less output. Consequently, marginal-cost reduction is less beneficial to me than to A, which means I want to give my CEO weaker—i.e., less sensitive—incentives.” In other words, weak incentives for B's CEO is B's best response to strong incentives for A's CEO. The same argument shows that the converse is also true. Consequently, there exists an equilibrium in which the two firms employ their CEOs under different—but optimal—incentive schemes.

In terms of equation (1), the consequences of heterogeneity in compensation schemes are

⁷Which will have the more sensitive relationship is theoretically indeterminate. On the one hand, because the signal is less informative in firm A, the cost of inducing a given level of effort from A's CEO will be greater than inducing the same level of effort from B's CEO. Consequently, A will induce less effort from its CEO than will B (assuming comparable benefits from their CEOs' efforts). This difference in effort will tend to make A's CEO's pay less sensitive to firm performance. On the other hand, the greater cost faced by A arises because A must use a more sensitive, hence riskier, pay schedule to induce a given level of effort. This effect will tend to make A's CEO's pay more sensitive to firm performance. For a more complete discussion of this issue see Hermalin (1992) or Hermalin and Katz (1996).

clear. First, the estimated coefficients will be a weighted average of the true coefficients. This could seriously bias the coefficients downward. This would be true if the low-incentive firms were the more numerous in the sample. It would also be true if the variance of the performance measures were greater for the low-incentive firms.⁸ Since, as we suggested above, greater variance in the performance measure could lead to less sensitive incentives, this is not an unreasonable concern.

Moreover, even if all the firms in the sample used highly sensitive compensation schemes, the regression results could still suggest otherwise. From Holmstrom (1979), we know that incentive schemes should give the most weight to the most informative performance measures; which these are, however, could vary from firm to firm for reasons already discussed. Consequently, although all firms heavily weight a performance measure, they do not weight the same measure. It is possible, therefore, to have a situation in which, say, half the firms heavily weight one of two performance measures but not the other, while the other half heavily weight the second measure but not the first. Hence, the estimated coefficients on these measures are both low; inaccurately suggesting that the firms in the sample are making only minor use of incentive schemes. Indeed, the downward bias could be considerable. Recall that if a measure is less informative (in the Blackwell sense), then it is riskier (see Hermalin and Katz, 1996). This means that the firms that *don't* weight a performance measure are more likely to have a high variance for that measure. Their coefficient on this measure will, therefore, receive greater weight than will the coefficient of the firms who do weight this measure (see footnote 8).

Finally, the standard model (1) assumes that the \mathbf{z}_j enter the performance-compensation relation as separate regressors rather than as determinants of firm-specific β 's. A preferable strategy would be a specification that enabled us to test this hypothesis to see whether the "controls" in fact differ from other exogenous heterogeneity factors that determine the incentive contracts chosen by specific firms.

⁸Suppose there were two types of firms *equally* represented in the sample. Suppose, for convenience, that the true relationship was

$$y = x\beta_i + \varepsilon,$$

where i indexes the type of firm and that $\mathbb{E}\{x\} = 0$ for both types. Let σ_i^2 denote the variance of x for type i . Suppose $\beta_1 < \beta_2$ and $\sigma_1^2 \gg \sigma_2^2$. Then, since the estimated coefficient is

$$\hat{\beta} \approx \frac{\sigma_1^2\beta_1 + \beta_2\sigma_2^2}{\sigma_1^2 + \sigma_2^2},$$

it follows that the estimated coefficient is much closer to β_1 than to β_2 . It is readily shown that this intuition extends to a more general framework such as equation (1).

These criticisms suggest an alternative specification to equation (1):

$$\mathbf{y}_j = \mathbf{X}_j \boldsymbol{\beta}_j + \boldsymbol{\varepsilon}_j, \quad (2)$$

where

$$\boldsymbol{\varepsilon}_j \sim N(\vec{\mathbf{0}}, \sigma_j^2 \mathbf{I}_{n_j})$$

is the error variance associated with $\boldsymbol{\varepsilon}_j$, $\boldsymbol{\beta}_j$ is a $K \times 1$ vector of coefficients hypothesized to vary randomly across firms (but to be stationary within a firm), n_j is the number of observations (years of data) on the j th firm, \mathbf{I}_n is an $n \times n$ identity matrix, \mathbf{y}_j is the $n_j \times 1$ vector of observations on the compensation variable, and \mathbf{X}_j is the $n_j \times K$ vector of observations on the performance measures. The coefficients on the performance measures are themselves functions of exogenous heterogeneity variables. Let \mathbf{Z}_j be the $K \times L$ block diagonal matrix of heterogeneity variables,

$$\mathbf{Z}_j = \begin{pmatrix} \tilde{\mathbf{z}}_j' & \vec{\mathbf{0}} & \vec{\mathbf{0}} \\ \vec{\mathbf{0}} & \ddots & \vec{\mathbf{0}} \\ \vec{\mathbf{0}} & \vec{\mathbf{0}} & \tilde{\mathbf{z}}_j' \end{pmatrix},$$

where $\tilde{\mathbf{z}}_j$ is a l -element vector of variables hypothesized to determine a firm's incentive contracts and $L = K \cdot l$. Then the following set of equations determines $\boldsymbol{\beta}_j$:

$$\boldsymbol{\beta}_j = \mathbf{Z}_j \boldsymbol{\delta} + \boldsymbol{\omega}_j, \quad (3)$$

and

$$\boldsymbol{\omega}_j \sim N(\vec{\mathbf{0}}, \boldsymbol{\Omega}_j),$$

for firms $j = 1, \dots, J$ independently. $\boldsymbol{\delta}$ is a common $L \times 1$ vector of heterogeneity coefficients and $\boldsymbol{\Omega}_j$ is the firm-specific $K \times K$ disturbance covariance matrix for the heterogeneity terms.

One strategy for estimating (2) and (3) would be to follow Garen (1994), who estimated (2) and (3) starting with Jensen and Murphy (1990)'s firm-specific ordinary least squares estimates of

$$y_{tj} = \mathbf{x}'_{tj} \boldsymbol{\beta}_j + \varepsilon_{tj}, \quad (4)$$

where t indexes time. The estimates, $\hat{\boldsymbol{\beta}}_j$, were then used as dependent variables in a second stage ordinary least squares estimate of (3) for each firm. This strategy will not produce the best linear unbiased estimate of either $\boldsymbol{\beta}_j$ or $\boldsymbol{\delta}$ (Amemiya, 1978; Hsiao, 1986; Laird *et al.*,

1987), because it ignores the variance component structure, $\text{Var}(y_{tj})$, which is a function of both ε_{tj} and ω_j . Since our interest is in both the β_j and δ parameters and population inference, equation (2) is more correctly viewed as a random-coefficients model in which the regression coefficients are assumed to be the dependent variables of another regression (3). Also, the use of an estimated coefficient in (3) precludes estimating it using OLS: The resulting error term is heteroskedastic, so care must be taken to insure consistent estimates of the standard errors. In particular, if ω_j and ε_{tj} are correlated, then a systems estimator, such as we employ, will be a more efficient estimator than a two-stage procedure.

Combining (2) and (3) yields a random-coefficients model (general hierarchical model) for our panel of savings and loans:

$$\mathbf{y}_j = \mathbf{X}_j \mathbf{Z}_j \boldsymbol{\delta} + (\boldsymbol{\varepsilon}_j + \mathbf{X}_j \boldsymbol{\omega}_j). \quad (5)$$

Equation (5) allows for firm specific heterogeneity in the determination of all aspects of the contract structure. The relationship between the exogenous heterogeneity factors and the performance coefficients can be tested using the parameter estimates $\boldsymbol{\delta}$. The model also provides estimates and standard errors for the performance coefficients themselves, β_j , and allows for direct tests of the performance-compensation relation for thrift CEO's. Note because \mathbf{X}_j contains a column of 1's, this specification encompasses the specification in which the \mathbf{z} -variables are additively separable from the other \mathbf{x} -variables (i.e., specification (1)).

Observe that we can't estimate (5) by ordinary least squares (equivalently, maximum likelihood) because the combined error term $(\boldsymbol{\varepsilon}_j + \mathbf{X}_j \boldsymbol{\omega}_j)$ is correlated with the independent variables, unless $\mathbb{E}\{\mathbf{X}_j \boldsymbol{\omega}_j\} = \vec{0}$. This, however, would be an unreasonable assumption: Performance could be correlated with the error term in (3) because a positive (negative) shock to β_j should induce a positive (negative) shock to performance through its impact on incentives.⁹ The consequence of this correlation is to bias downward the estimates of the variance-covariance matrix of ε_{tj} and ω_j , Γ , obtained from (unrestricted) maximum likelihood estimation. We can avoid this bias by using a restricted maximum likelihood (REML) estimator that accounts for the loss in degrees of freedom from estimating $\boldsymbol{\delta}$ (Laird and Ware, 1982; Wong and Mason, 1991). Specifically, the likelihood function is maximized with respect to Γ only, holding $\boldsymbol{\delta}$ constant at an initial estimate, $\boldsymbol{\delta}^*$, which yields an estimate Γ^* . Following Wong and Mason, we then estimate β_j and $\boldsymbol{\delta}$ by their posterior means and assess the uncertainty of this estimation by their respective posterior variances and covariances with Γ set at Γ^* . Having obtained a new estimate, $\boldsymbol{\delta}^*$, we then iterate the REML step;

⁹A caveat is that if one views the performance measures as profits, and thus as having the cost of incentives subtracted from them, then an error in setting compensation (i.e., $|\omega_j| \neq 0$) means a reduction in profits. However, unless it is just the magnitude of the error that determines the reduction in profits (i.e., $-|\omega_j|$ and $|\omega_j|$ result in precisely the *same* reduction in profits), we would still expect some correlation between \mathbf{X}_j and ω_j , although now the sign of that correlation is less obvious.

and then iterate the estimation-of- β_j -and- δ step. This iterative process is accomplished by applying an expectations-maximization (EM) algorithm (Dempster *et al.*, 1981; McLachlan and Krishnan, 1997) until convergence.¹⁰ In this application, we assume that all thrifts' incentive contracts respond to the same set of heterogeneity variables (our estimation strategy would, however, allow us to relax this assumption).

3 Estimating the Pay-Performance Relationship While Accounting for Heterogeneity

In this section, we discuss our estimation of (5). In the previous sections, we identified three sources of heterogeneity: thrift size, the informativeness (riskiness) of the performance measures, and the value of the shareholders' put option.

In our estimation of (5), we use the natural log of the thrift's assets as a measure of its size. Although other measures exist (e.g., deposits), we view assets as preferable (a high degree of multicollinearity among the potential size measures prevents us from using more than one). One reason is that the thrifts themselves often measure their size in terms of their assets, so assets could accord with how the thrifts themselves think of size when deciding on compensation. Furthermore, an analysis of 165 proxy statements for a related sample of thrifts revealed that a significant minority of them partially base the CEO's compensation on assets, which suggests a second motive for controlling for assets besides their role in the determination of β . In any case, the strong correlation among the various size measures lessens the importance of choosing "the right one."

As we discussed above, measuring the risk thrifts face is difficult. Here, we estimate the CAPM beta using the last sixty periods of monthly returns. Summary statistics for the CAPM beta estimates are in the Appendix. Although the CAPM beta estimates are somewhat problematic given the high proportions of non-trading debt in thrift portfolios, we felt the advantages of a market-based measure was superior to any other risk measure we could construct.

Estimating the value of the shareholders' put option is impossible with the data available. To gain some insight about the importance of the put option to the shareholders and management we use two measures: one, the officers' ownership share and, two, the thrift's tangible capital.

The idea behind this first measure is that the more the officers own of the firm, the more aligned their interests will be with the shareholders'. This could make them more likely to

¹⁰Estimation was carried out using the program *GENMOD* supplied by the Population Studies Center of the University of Michigan.

pursue strategies that exploit the put option. On the other hand, unless their ownership is part of a highly diversified portfolio, they will also be risk averse; if their stock ownership represents a significant proportion of their wealth, their risk aversion could lead them to minimize the probability that the thrift is put to the deposit insurer. Even though the overall effect of the officers' share ownership is ambiguous, it is unlikely that the two effects cancel each other out; hence, we need to control for share ownership.

The second measure of the put option's importance reflects how close the option is to being "in the money." The less tangible capital the thrift has, the closer it is to insolvency, which raises the value of the put option. Consequently, we expect low-tangible-capital thrifts to be more likely to pursue strategies that exploit the put option than high-tangible-capital thrifts.

A final source of heterogeneity, which we have not discussed, is the variation in ability across thrifts' management. That is, managers differ in their competence, commitment, creativity, and other attributes. This heterogeneity in ability will lead to heterogeneity in compensation for two reasons: first, more able managers are more valuable; and, hence, are likely to be paid more. Second, differences in ability will translate into differences in the optimal incentive schemes they are given.¹¹

To control for these differences in ability, we use a modified version of the nonparametric efficiency measure employed by Hermalin and Wallace (1994). Specifically, we say that thrift j is inefficient relative to thrift m —alternatively, thrift m *dominates* thrift j —if thrift j 's revenues are less than thrift m 's, but thrift j 's costs are not less than thrift m 's. The idea behind this test is that thrift j could do better by emulating thrift m , so thrift j cannot be operating efficiently.¹² Following Hermalin and Wallace (1994), we say that thrift j is inefficient if it is dominated by 2.5% or more of the thrifts that could potentially dominate it (i.e., thrifts with greater revenue). On the assumption that efficiency and managerial ability should be correlated, we use this inefficiency measure as a proxy for managerial ability.

As before, the performance measures (\mathbf{X}_j) are the rate of returns on assets, the annual change in the stock price, and a constant. The determinants of the compensation scheme (\mathbf{z}_j) are a constant, the relative market risk of the firm's strategies, CAPM beta, our production efficiency measure (1 = efficient), the percent of tangible capital, the share of stock ownership

¹¹See Hermalin and Weisbach (1997) for a model in which differences in ability across the managers of different firms lead to different governance practices across these firms.

¹²Assuming that the thrifts face the same input and output prices. It is in this regard that our test differs from Hermalin and Wallace's test (which, itself, is based on tests developed in Varian, 1984), as their test does not require this assumption. Unfortunately, we do not have the data to use their test. Although not ideal, it turns out that Hermalin and Wallace find that our test and theirs are extremely closely correlated (from Hermalin and Wallace, 1994, p. 365, the correlation coefficient is .96), so the cost we incur by using a less robust test is likely to be minimal.

of the officers, and total assets.

In Table 2, we consider the heterogeneity variables as determinants of the coefficients on the performance measures rate of return on assets (ROA) and the percentage change in stock price. A likelihood-ratio test reveals that the heterogeneity variables are jointly significant at better than the 1% level.

Table 2: Determinants of the Compensation
Performance Relationship

Independent Variable	Dependent Variable = Log(Cash Compensation) _{t+1}
	Coefficient Estimate (Standard Error)
Intercept	10.052*** (.276)
CAPM Beta	.159* (.089)
Production Efficiency	.062 (.632)
Tangible Capital	.047** (.021)
Share of Stock Owned by Officers	-0.003 (.343)
Log of Assets	.389*** (.041)
ROA	.137* (.071)
CAPM Beta × ROA	.038 (.023)
Production Efficiency × ROA	-0.025 (.026)
Tangible Capital × ROA	-0.011** (.005)
Share of Stock Owned by Officers × ROA	.004** (.002)
Log of Assets × ROA	-0.005 (.011)

Table 2: Determinants of the Compensation
Performance Relationship (continued)

Independent Variable	Dependent Variable = Log(Cash Compensation) _{t+1}
	Coefficient Estimate (Standard Error)
Percentage Change in Stock Price	.003 (.003)
CAPM Beta × Percentage Change in Stock Price	.0015* (.0008)
Production Efficiency × Percentage Change in Stock Price	.001 (.001)
Tangible Capital × Percentage Change in Stock Price	−0.0005** (.0002)
Share of Stock Owned by Officers × Percentage Change in Stock Price	-3.6×10^{-5} (9.0×10^{-5})
Log of Assets × Percentage Change in Stock Price	8.0×10^{-5} (.0004)

* Significant at the 10% level

** Significant at the 5% level

*** Significant at the 1% level

We find, as expected, that larger firms (greater assets) pay their CEOs significantly more, *ceteris paribus*. Since both the dependent variable and assets are measured in logs, we can interpret the coefficient as an elasticity: A one-percent increase in assets results in a 0.39% increase in salary. Thrifts with more tangible capital pay more, *ceteris paribus*. Finally, the greater is CAPM beta, the greater is CEO salary, *ceteris paribus*. To the extent that a higher CAPM beta translates into riskier incentives, this last result is consistent with a need for a *non*-contingent compensation for making the CEO bear risk through his incentives. The *direct* effects of the other heterogeneity controls are not statistically significant.

The two heterogeneity controls that have a statistically significant impact on $\beta_{j,ROA}$ (the marginal effect of ROA on CEO salary for the j th firm) are tangible capital and the share of stock owned by officers. The former has a negative effect, while the latter has a positive effect. The effect of tangible capital is consistent with the view that firms with low tangible capital are “in the money” with respect to their put option. Hence, shareholders will want their managers to gamble more and they provide them incentives to do so by increasing $\beta_{j,ROA}$. The positive impact of stock ownership is more puzzling, since the opposite would seem more consistent with theory: Since stock itself provides incentives, the firm would seem to have less need for incentive pay. On the other hand, our stock-ownership measure, the share of stock owned by officers, is only an indirect measure of the incentive effects of stock ownership.¹³ It is, however, a direct measure of managerial control. As Hermalin (1992) showed, increasing the control (bargaining power) of the agent can lead, in theory, to *more* sensitive incentive pay.

The two heterogeneity controls that have a statistically significant impact on $\beta_{j,STKCHG}$ (the marginal effect of the percentage change in the stock price) in the salary regression are the CAPM beta and tangible capital. The positive effect of CAPM beta is consistent with the view that the greater the noise, the less informative the signal, so the more powerful the incentive tied to that signal needs to be to motivate managers. The negative effect on tangible capital is consistent with the “put-option theory” advanced earlier.

Although the heterogeneity factors are independently of interest, our focus is on estimating the (overall) coefficients of the performance measures (i.e., the β_j). Our finding that the heterogeneity factors are statistically and economically important in the determination of the random-effects performance measures supports our contention that different thrifts employ different compensation schemes and that these differences across firms are significant. Importantly, these differences are associated with the relationship between the performance measures and the ownership structure and the riskiness of the firms.

¹³Ideally, we would like to know how sensitive the CEO’s *total* well-being is to changes in the stock price. This would require knowing not only the derivative of his wealth with respect to a dollar change in the stock price, but also knowing his total wealth, since the impact on his utility depends on his current wealth.

Figure 1 provides histograms (the back rows labeled REML) for the firm-specific random-effects coefficients, $\beta_{j,ROA}$ and $\beta_{j,STKCHG}$. The mean of $\beta_{j,ROA}$ across all firms is .0227 and the mean for $\beta_{j,STKCHG}$ is .0026. As shown by the histograms, the coefficient distributions are slightly more skewed and have fatter tails than would be expected in normal distributions, however, they are generally symmetric. More importantly, treating the estimated performance coefficients as random variables, we easily reject a null hypothesis that the true population performance incentives are not positive:

$$\begin{aligned}\Pr \{ \beta_{j,ROA} \leq 0 \} &= 0.03 \\ \Pr \{ \beta_{j,STKCHG} \leq 0 \} &= 0.0\end{aligned}$$

Alternatively, 53 out of 86 $\beta_{j,ROA}$ are positive. Assuming, as a null hypothesis, that the $\beta_{j,ROA}$ were equally likely to be positive as negative, the probability of getting 53 or more positives from a sample of 86 is only .011. Similarly, 84 out of 86 $\beta_{j,STKCHG}$ are positive. The probability of 84 or more positives from a sample of 86 under a null that they were equally likely to be positive as negative is essentially zero. Hence for all these reasons, we can conclude that there is a statistically significant and positive mean effect on thrift CEO compensation for both the accounting-based performance measure, ROA, and the market-based performance measure, change in annual stock price.

Additionally, our results indicate that the performance measures, ROA and stock price change, have economically significant effects on thrift CEO compensation levels: For the average CEO, the effect of a one-standard-deviation change in ROA on salary is \$12,490 (compared to \$4478 from estimating (1) by OLS). The effect of a one-standard-deviation change in the stock-price performance measure on salary is \$18,245 (compared to \$6548 from estimating (1) by OLS). In short, by accounting for heterogeneity, the *economic* significance of the pay-for-performance relationship is much greater than suggested by regressions that ignore this heterogeneity.

Finally, recall that our reason for using REML was our concern that OLS (*unrestricted* maximum likelihood) coefficients would be biased. To see whether that concern was warranted, we estimated (5) using fixed-effects OLS. We then calculated β_j for each firm. The results are graphed in Figure 1 (the front rows labeled “Fixed Effects”). As is clear, the results from this method vary markedly from our REML results. To confirm this visual examination, we also calculated Kolmogorov-Smirnov tests of the null hypothesis that the two methods yielded the same distribution performance coefficients. This null hypothesis was rejected at better than the .0001 level in both cases.¹⁴ It is worth noting firstly that the OLS estimates understate the heterogeneity in the incentive coefficients and, secondly,

¹⁴For the impact of ROA on salary, the Kolmogorov-Smirnov Z statistic was 3.431. For the impact of change in stock price on salary, it was 5.109.

that they suggest a lower mean effect than does our random-coefficients model. This second point indicates that ignoring heterogeneity among firms can lead to results that understate the economic significance of the pay-for-performance relationship.

4 Directly Testing the Impact of Heterogeneity on Compensation Schemes

In the previous section, we showed that if we accounted for heterogeneity in compensation schemes, we would obtain more reasonable estimates of the sensitivity of pay to performance. Our second approach, reported in this section, is to see whether our heterogeneity controls can predict the terms of the compensation contracts directly.

From the *SNL Executive Compensation Review for Thrift Institutions* volumes, we know whether or not the CEO's compensation package contained each of the following six features: (i) an employment contract, (ii) a golden parachute, (iii) an incentive plan, (iv) a restricted stock option plan, (v) an unrestricted stock option plan, and (vi) a supplemental pension plan. If our hypothesis that heterogeneity leads to differences in compensation schemes is correct, then we would expect our measures of heterogeneity to predict whether or not these features were included in the CEO's compensation package.

To carry out this test, we assumed that the probability the package has the h th feature is $\Phi(\tilde{\mathbf{z}}_j' \boldsymbol{\xi}_h)$, where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal, $\tilde{\mathbf{z}}_j$ is the same vector of heterogeneity controls and other controls used in Table 2, and $\boldsymbol{\xi}_h$ is the coefficient vector to be estimated. To estimate these vectors, we estimated the following system by nonlinear seemingly unrelated regression (NLSUR):

$$\begin{pmatrix} f_j^1 \\ \vdots \\ f_j^6 \end{pmatrix} = \begin{pmatrix} \Phi(\tilde{\mathbf{z}}_j' \boldsymbol{\xi}_1) \\ \vdots \\ \Phi(\tilde{\mathbf{z}}_j' \boldsymbol{\xi}_6) \end{pmatrix} + \begin{pmatrix} \zeta^1 \\ \vdots \\ \zeta^6 \end{pmatrix}, \quad (6)$$

where $f_j^h = 1$ if the thrift j 's package has feature h and $f_j^h = 0$ if it doesn't and where ζ^h is the error term in that regression.

Our estimates of (6) are given in Table 3. The most important question is whether the heterogeneity controls have explanatory power *vis-à-vis* the features of the compensation package. A Wald test yields a χ_{24}^2 statistic of 43.57, which is significant at the 5% level. This is consistent with our hypothesis that these heterogeneity controls predict the features of the compensation package, and, as such, are important to understanding the relationship between pay and performance.

Table 3: Nonlinear SUR Estimates of Factors Affecting CEO
Employment Benefits
(Dependent Variable = CEO Receives Benefit)

Independent Variables	Employment Contract	Golden Parachute	Incentive Plan	Restricted Stock Option Plan	Stock Option Plan	Supplemental Pension Plan
Intercept	3.142*** (.698)	1.396*** (.434)	-0.013 (.402)	-2.003*** (.732)	3.183** (1.60)	-20.715*** (3.93)
Log of Assets	-0.109*** (.044)	-0.030 (.027)	.055** (.026)	.054 (.043)	-0.059 (.106)	1.45*** (.276)
CAPM Beta	.069 (.112)	-0.032 (.013)	-0.017 (.098)	.432*** (.129)	.501*** (.166)	.217 (.204)
Production Efficiency	-0.188 (.178)	-0.254 (.158)	-0.022 (.145)	.114 (.189)	-0.716* (.382)	-0.583** (.257)
Tangible Capital Ratio	-0.031 (.022)	.014 (.022)	-0.057*** (.020)	-0.056 (.036)	-0.114*** (.027)	-0.078** (.038)
Share of Stock Owned by Officers	-0.074*** (.013)	-0.049*** (.013)	-0.016 (.011)	-0.0057 (.011)	.212** (.097)	-0.056** (.019)
Wald Test	$\chi^2_{24} = 43.57^{**}$					

* Significant at the 10% level

** Significant at the 5% level

*** Significant at the 1% level

We argued above that larger thrifts would value managerial effort more than smaller thrifts and would, thus, make more use of incentive plans. This is borne out in the data: The more assets a thrift has, the greater the likelihood it includes an incentive plan. Large asset thrifts are also more likely to use supplemental pension plans. Larger asset thrifts are statistically significantly less willing to use employment contracts which provide written guarantees that the employer will provide specific types of benefits in exchange for continued employment. Larger thrifts, therefore, appear to prefer strategies that do not entrench the CEO by providing high levels of insurance against dismissal.

The more of the firm the officers own and, thus, the more voting power they have, the less they need to worry about change-in-control issues. This, too, is borne out in the data: The probability of insurance against dismissal (i.e., employment contracts and golden parachutes) decreases with the share of the firm owned by the officers.

The more sensitive its stock price is to movements in the market (i.e., the greater its CAPM beta), the more likely the officers of a thrift would desire to participate in positive changes in the firms stock prices through stock option plans. This is consistent with what we find in Table 3.

It is somewhat disturbing that thrifts with lower tangible capital ratios are more likely to use stock option plans suggesting “bet-the-bank” incentives may underlie the choice of stock option plans for thrift CEOs. That is, shareholders could be attempting to induce their managers to take actions that exploit their put option.

Supplemental pension plans are more likely to be used by inefficient firms and when the officers have a small stake in equity ownership. Because the CEOs in these firms do not have a large ownership stake in long-cycle risk-management strategies, paying benefits through retirement packages may serve to better align the incentives of CEO to an equally long-run horizon. To the extent, however, that these firms are also inefficient producers of their lines-of-business (see Hermalin and Wallace, 1994, for evidence), this suggests that this type of compensation benefit does not provide appropriate incentives for the efficient management of shorter horizon production costs. That the use of supplemental pension plans decreases with capitalization is somewhat puzzling; one possibility is that struggling firms cannot afford cash compensation and, so, prefer deferred compensation.

Incentive plans are designed to provide specific shorter horizon performance related benefits. As shown in Table 3, these plans are chosen by less well capitalized firms. Obviously, an inadequately capitalized firm faces questions about its long-run prospects, which could necessitate the use of shorter horizon, usually one year, pay-off strategies through incentive contracts. This result is again consistent with principal-agent predictions.

Admittedly, many of the coefficients in Table 3 are not significant. Part of the problem may lie with the lack of variation in some of the features, such as employment contracts and stock option plans (see Table A.1). Another part of the problem may lie with the crudeness of

the dependent variables. In particular, it would be better to know *how* sensitive the incentive plan is to performance rather than simply whether it exists or not. Despite such deficiencies, however, Table 3 does provide additional support for our claim that heterogeneity is an important determinant of the structure of CEO compensation in thrifts.

5 Conclusion

In this paper, we presented a theoretical reason why the “standard” specification used in testing the pay-performance relationship may be inadequate and may tend to find “small” effects of performance on executive pay: This specification treats all firms in the sample as if they used the same compensation scheme, when, in fact, theory predicts that they will use different schemes to reflect heterogeneity on a number of dimensions, including firm size, managerial ability, and how informative firm performance is about managerial performance. We proceeded to develop an alternative specification that does not assume that all firms use the same compensation scheme. Rather our specification allows each firm to employ a different compensation scheme.

We tested this specification on data from the savings and loan (thrift) industry. Although these data have many advantages, they also have some shortcomings. Among their shortcomings is that it is difficult to measure risk (which we require as a proxy for informativeness). We “solved” this problem by using stock returns series to construct an estimate of each thrift’s CAPM beta. Finally, the thrift industry is a regulated industry and previous empirical research suggests that the pay-performance relationship is weakest in such industries. The fact that we find a strong pay-performance relationship even in this case is, in our view, strongly supportive of our alternative specification for tests of the principal-agent prediction.

We found that we could reject the null hypothesis that the standard specification was as good as our specification: Our controls for heterogeneity were jointly significant. Using the results from this regression, we were able to calculate the relationship between pay and performance for each firm individually. These calculated values revealed a more pronounced positive relationship between pay and performance than did the standard specification. Various tests indicate that these calculated values are significant.

In addition, our estimation strategy allows us to identify a number of important ways that heterogeneity factors lead to incentive structures in the thrift industry. Our finding that tangible capital has a statistically significant, negative impact on $\beta_{j,ROA}$ and $\beta_{j,STKCHG}$ suggests that the “put option” arising from deposit insurance acts as an important determinant of incentive schemes. Further, direct evidence for incentive strategies that seek to induce managers to “bet-the-bank” is our finding that poorly capitalized firms are more likely to use stock-option plans and short-term incentive plans. Another ominous result is that managerial

ability, as measured by our efficiency measure, appears to have no statistically significant effect on the use of either ROA or stock-price changes as incentives.

Many other hypothesized effects were also found: For instance, the riskiness of the firm, as measured by its CAPM beta, was shown to have a statistically significant and positive impact on $\beta_{j,STKCHG}$. This result is consistent with the view that the less informative a given signal, the more powerful the incentive tied to the signal must be. Similarly, CAPM is also shown to be an important determinant of whether firms use stock-option and restricted stock-option plans. Although the share of stock held by the firm's officers leads to somewhat ambiguous results for $\beta_{j,ROA}$, it leads to a greater value for $\beta_{j,STKCHG}$. In the overall compensation package, firms with less stock held by officers (thus, in which officers were less protected by their voting power *vis-à-vis* change-in-control issues), tended to use employment contracts and golden parachutes more.

Work remains to be done, however. First, it would be nice to replicate these results using other data sets; particularly, as the thrift industry has many special features.

Second, alternative measures of compensation and CEO well-being could be considered. For instance, it is known that poor performance weakly predicts the CEO being fired (see, e.g., Warner *et al.*, 1988). Moreover, it is also known that this relationship varies across firms with respect to other factors, such as the composition of the board of directors (see Weisbach, 1988). What other heterogeneity factors affect this relationship?

Third, how do these factors affect the weights that these two incentives (risk of dismissal and performance pay) receive; that is, what might cause some firms to rely on one incentive more than the other?

Finally, many of the heterogeneity factors are, in the long run, endogenous. As firms and industries adjust and these heterogeneity factors change, do we tend to see convergence or divergence among the firms with respect to their incentive schemes? That is, given a longer panel than we have available, could one say anything about the evolution of incentive schemes over time?

Despite these open questions and the usual shortcomings that plague any empirical work, we nevertheless feel confident that our results support our theoretical arguments and establish the importance of controlling for inter-firm heterogeneity in incentive schemes when estimating the pay-performance relationship.

Data Appendix

The data used in this study were obtained from a variety of sources. The SNL Security data was compiled from two publications: The *SNL Quarterly Thrift Digest* (1988 - 1992) and the *SNL Executive Compensation Review for Thrift Institutions* (1988 - 1994). These

Table A.1: Summary Statistics ($N = 86$)

Variables	Mean	Standard Deviation
SNL Securities		
Cash Compensation of CEO (1988 Dollars)	\$ 255,400	\$ 244,828
Total Assets (1988 Dollars)	\$ 29,826,807	\$ 69,260,715
Employment Contracts	84.8%	35.8%
Golden Parachutes	72.5%	44.6%
Incentive Plans	56.7%	49.7%
Restricted Stock Option Plans	15.6%	36.3%
Stock Option Plans	94.0%	23.0%
Supplemental Pension Plan	22.8%	42.0%
Tangible Capital (Equity \div Assets)	.062	.038
Share of Stock Owned by Officers	4.63%	6.17%
ROA (Rate of Return on Assets)	0.36%	0.83%
Annual Change in Stock Price	5.20%	34.7%
Other Sources		
Production Efficiency (1 = efficient; 0 = inefficient)	22.5%	13.0%
CAPM Beta	.733	.713

documents are published by SNL Securities, Charlottesville, VA. The data contained in the SNL Executive Compensation Review for Thrift Institutions is collected from proxy statements for all publicly traded thrift institutions. The publications provide the most comprehensive coverage for publicly traded thrifts in the country. Due to bankruptcies and ownership patterns, thrifts from the states of Arizona and Texas are under represented in the SNL industry surveys. We included all institutions with complete data on the variables used in the study for the period 1988 to 1992.

The summary measures for the variables used in the study are reported in Table A.1. CEO compensation is computed as salary, bonus, deferred compensation, and other forms of cash-equivalent compensation received by the CEO in a given year. All payments or gains realized from long-term incentive plans, where identifiable in the proxy statement are also included. Severance pay, matching contributions, and signing bonuses are excluded. The percentage change in stock price variable measures the percentage change in the thrift's stock price as of the last close in a year compared to the last close in the previous year. The share of stock ownership measure is the proportion of stock held by officers of the thrift.

The other benefits are:

- **Employment contract:** An agreement between the employer and the CEO that provides a written guarantee that the employer will provide defined types of compensation and benefits to the CEO during a specified period in return for continued employment.
- **Golden Parachute:** A contract or agreement that guarantees payment of defined compensation and benefits to executives in the event of an acquisition, merger, divestiture or other change in control. The parachute may be a provision of, or separate from the CEO's employment contract.
- **Incentive Plan:** A plan designed to retain the CEO and reward achievement of specific performance objectives during a specified period of time. Short-term incentive awards, or annual bonuses, are typically earned over one year, and may be performance based or discretionary. Payment of the awards may be in cash, stock, or a combination of the two.
- **Restricted Stock Option Plans:** A plan that provides for issuance to the CEO of stock which have transfer and sale limitations. Generally, the stock—received at nominal or no cost—vests over a period of time or at one time in the future, and is contingent upon continued employment and, in some cases, performance. The stock must be subject to substantial risk of forfeiture to avoid taxation as current compensation.
- **Stock Option Plan:** A plan that grants executives the opportunity over a period of years to purchase shares of common stock, generally at the trading price of the shares on the date of grant. The options under such a plan may be nonqualified or qualified to receive special tax treatment (incentive stock options).
- **Supplemental Pension Plan:** A nonqualified plan that provides for retirement benefits, usually to supplement qualified plan benefits. Supplemental plans can take two forms: ERISA excess plans which restore benefits that qualified plans cannot pay, and SERPs which go beyond replacing qualified plan benefits due to Section 415 limits by changing the plan formula to include, for example, bonuses or other incentive pay.

Tangible Capital Ratio is measured as total shareholder equity minus goodwill and intangible assets expressed as a percent of total assets. The information for the share of stock held by officers was obtained from the most recent proxy statement for each firm.

The efficiency measure was computed based on the methodology developed in Hermalin and Wallace (1994). We use a zero/one version of their measure which equals one (the firm is efficient) if the firm is dominated in production by less than 2.5% of firms that are larger

than it. The data for constructing this measure was obtained from the Federal Home Loan Bank's quarterly thrift financial reports.

The CAPM beta's were computed for each firm using the preceding 60 months of monthly dividend adjusted returns. The proxy for the market returns was the S&P 500 and all values were risk adjusted by monthly returns on the 3-Month T-bills. As shown, in Table A.1 the mean level of CAPM beta for firms in the sample was .733, however, as is obvious given the standard deviation of this measure there was considerable variance in thrift betas. The returns series were obtained from the Center for Research in Security Prices (CRSP) files and the 3-month T-Bills data was obtained from CITIBASE.

References

- Amemiya, T. (1978). A note on a random coefficients model. *International Economic Review*, **19**(3), 793–796. October.
- Barro, J. R. and Barro, R. J. (1990). Pay, performance, and turnover of bank CEOs. *Journal of Labor Economics*, **8**, 448–481.
- Carroll, T. M. and Ciscel, D. H. (1982). The effects of regulation on executive compensation. *The Review of Economics and Statistics*, **64**, 505–509.
- Ciscel, D. H. and Carroll, T. M. (1980). The determinants of executive salaries: An econometric survey. *The Review of Economics and Statistics*, **62**, 7–13.
- Dempster, A. P., Rubin, D. B., and Tsutakawa, R. K. (1981). Estimation in covariance components models. *Journal of the American Statistical Association*, **76**(374), 341–353. June.
- Garen, J. E. (1994). Executive compensation and principal-agent theory. *Journal of Political Economy*, **102**(6), 1175–1199.
- Gibbons, R. (1997). Incentives and careers in organizations. In D. Kreps and K. F. Wallis, editors, *Advances in Economic Theory and Econometrics, 7th World Congress, Volume II*. Cambridge University Press.
- Grossman, S. J. and Hart, O. D. (1983). An analysis of the principal-agent problem. *Econometrica*, **51**, 7–45.
- Haubrich, J. G. (1994). Risk aversion, performance pay, and the principal-agent problem. *Journal of Political Economy*, **102**, 258–276.

- Hermalin, B. E. (1992). The effects of competition on executive behavior. *The RAND Journal of Economics*, **23**, 350–365.
- Hermalin, B. E. (1994). Heterogeneity in organizational form: Why otherwise identical firms choose different incentives for their managers. *The RAND Journal of Economics*, **25**, 518–537.
- Hermalin, B. E. and Katz, M. L. (1996). Corporate diversification and agency. Working Paper, University of California at Berkeley, <http://haas.berkeley.edu/~eap/WP.html>.
- Hermalin, B. E. and Wallace, N. E. (1994). The determinants of efficiency and solvency in savings and loans. *The RAND Journal of Economics*, **25**, 361–381.
- Hermalin, B. E. and Weisbach, M. S. (1997). Endogenously chosen boards of directors and their monitoring of the CEO. *American Economic Review*. forthcoming.
- Holmstrom, B. (1979). Moral hazard and observability. *The Bell Journal of Economics*, **10**, 74–91.
- Hsiao, C. (1986). *Analysis of Panel Data*. Cambridge University Press, Cambridge, England.
- Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review*, **76**, 323–329.
- Jensen, M. C. and Murphy, K. J. (1990). Performance pay and top-management incentives. *Journal of Political Economy*, **98**, 225–264.
- John, K. and John, T. A. (1993). Top-management compensation and capital structure. *Journal of Finance*, **48**, 949–974.
- Joskow, P., Rose, N., and Shepard, A. (1993). Regulatory constraints on CEO compensation. In *Brookings Papers on Economic Analysis, Microeconomics*, Washington, D.C. The Brookings Institution.
- Katz, M. L. (1991). Game-playing agents: Unobservable contracts as precommitments. *The RAND Journal of Economics*, **22**, 307–328.
- Laird, N. M. and Ware, J. H. (1982). Random-effects models for longitudinal data. *Biometrics*, **38**, 963–974. December.
- Laird, N. M., Lange, N., and Stram, D. (1987). Maximum likelihood computations with repeated measures: Application of the EM algorithm. *Journal of the American Statistical Association*, **82**, 97–105.

- McLachlan, G. J. and Krishnan, T. (1997). *The EM Algorithm and Extensions*. Wiley-Interscience, New York.
- Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, **29**, 449–470.
- Merton, R. C. (1977). An analytic derivation of the cost of deposit insurance and loan guarantees. *Journal of Banking and Finance*, **1**, 3–11.
- Varian, H. (1984). The nonparametric approach to production analysis. *Econometrica*, **52**, 579–599.
- Warner, J., Watts, R., and Wruck, K. (1988). Stock prices and top-management changes. *Journal of Financial Economics*, **20**, 461–492.
- Weber, N. C. (1986). The jackknife and heteroskedasticity: Constant variance estimation for regression models. *Economics Letters*, **20**, 161–163.
- Weisbach, M. (1988). Outside directors and CEO turnover. *Journal of Financial Economics*, **20**, 431–460.
- Wong, G. Y. and Mason, W. M. (1991). Contextually specific effects and other generalizations of the hierarchical linear model for comparative analysis. *Journal of the American Statistical Association*, **86**(414), 487–503. June.