CEO ownership and firm value:

Evidence from a structural estimation

Beau Page*

University of Rochester

November 2011

Abstract

I estimate a dynamic model in which a board of directors and a CEO interact to set the levels of CEO incentives and effort. The intent is to understand the effect that CEO incentives have on firm value. In the model incentive levels are the result of CEO risk aversion, the cost of CEO effort, the effect of CEO effort on firm value, the volatility of shocks to firm value, and the preference a board has over giving a CEO equity ownership. Model estimates show that CEO effort is an important component of firm value, and that CEO's exert a substantial amount of effort, on average 94.7% of the possible maximum. A one percentage point increase in CEO ownership is found to increase firm value 7.6 basis points, with the net (of the increase in CEO ownership) benefit to shareholders being 4.6 basis points.

Keywords: CEO Compensation, Dynamic Principal-Agent Model, Structural Estimation

* I wish to thank Toni Whited for her guidance and encouragement. I also thank Michael Raith, Boris Nikolov, Cliff Smith, Ed Owens, and Paul Nelson for their comments and criticisms, which have greatly improved this paper. Additional thanks to Chet Reeder and Ivan Ivanov for their feedback. I thank the participants in the Simon School AEC 510 Ph.D. student seminar for their comments and discussion.

From 1992 to 2007 the average CEO held 3.6% of the equity of his firm, worth $108 million 2000 dollars. CEOs acquired most of this ownership as a part of their compensation package intended to align their incentives with those of shareholders. This study seeks to quantify the gain shareholders receive from incentivizing CEOs. It also seeks to understand the economic mechanisms behind any such gains. Identifying these effects is difficult because CEO effort is unobservable, and the relation between CEO ownership levels and CEO effort is endogenous. To address these issues, I develop and estimate a dynamic model of interaction between a CEO and a firm's board of directors. The findings indicate that CEO effort makes an important contribution to firm value, which is sensitive to changes in the level of CEO ownership. On average, a 1% increase in CEO ownership increases shareholder value by 0.049% after taking into account the increase in CEO compensation. The model estimates imply that CEOs exert a substantial amount of productive effort, so that much larger ownership stakes are required to incentivizing them further.

Connecting the level of CEO ownership and firm value is hard because CEO efforts to increase or decrease firm value are unobservable and because incentive pay is set in equilibrium. The link from ownership to firm value goes through effort, and not measuring this effort leads to a downward bias of the estimate of the causal relation. For example, consider two firms, one with an lazy CEO and one with a hard-working CEO. With everything else held constant, because incentive pay is set in equilibrium, in the data we would observe the lazy CEO with high incentive pay, and the hard-working CEO with low incentive pay. However, because the two firms are otherwise identical, the two CEOs end up exerting the same amount of effort. In the data we would see identical values, and thus no relation between incentive pay and value, even though incentive pay is clearly working in this example.

By developing and estimating a structural model, I can identify the effect of ownership on firm value even though ownership is set in equilibrium and even though effort is unobservable. A CEO in the model must choose an effort level such that statistics from a simulation of the model are similar to those same statistics in the data. In a sense, effort can be viewed as the residual of the model. Once the estimation recovers the structural parameters that lead the CEO to exert the level of effort that matches the model to the data, I can test how the level of effort changes with changes in ownership, and how firm value changes with effort.

The model captures the intuition from the above example. The CEO chooses an effort level based on current ownership, the firm's value, and an expectation of future wealth. The CEO finds making effort costly, and his preferences are such that increasing the CEO's incentive pay has both income and substitution effects on his effort. (See also Edmans, Gabaix, and Landier (2009)). On one hand, with more ownership of the firm the substitution effect leads the CEO to substitute consumption for leisure—more work increases non-leisure consumption. On the other hand, the wealthier the CEO, the more he wants to increase his consumption of all goods, including leisure. He wants to enjoy his wealth, so he decreases effort. These competing effects temper the advantage of increasing CEO ownership, as wealthier CEOs can be harder to motivate. The final determinant of CEO effort is risk aversion. As the model is dynamic, the CEO's effort level is a function of his expected future wealth, which is risky. The CEO's risk aversion affects how he compares the benefits of more leisure now against the cost of lower expected future wealth.

The board of directors seeks to maximize shareholder value, given its perception of the costs of this ownership. Thus, it determines the CEO's level of ownership by equating the marginal benefit to firm value with the marginal cost, the main component of which is dilution of current shareholders' value. The board's optimization problem is complicated in two ways: the board cannot decrease the CEO's ownership, and has preferences over giving the CEO more ownership that go beyond the simple dilution effect. The board's inability to decrease ownership means there is a possible shadow cost to incentivizing a CEO. The board's preferences over granting equity can be positive, negative, or nonexistent, with intuitive effects on how much ownership the board grants the CEO.

I estimate the parameters of the model with simulated method of moments (SMM), using data on equity levels, firm size, the market-to-book ratio, and CEO tenure of public firms in the United States over the 1992 to 2007 period. SMM estimates of the parameter values are for the average firm in the sample. For this representative firm, the board of directors has a negative preference for giving the CEO ownership. This result implies the board sees hidden costs to granting equity aside from the value dilution. In addition, CEO effort has a large effect on firm value, as variation in CEO effort accounts for just under 10% of the variance of returns.

I use counterfactual analysis determine the sensitivity of firm values to an unexpected change in CEO ownership. The output of these exercises are upper and lower bounds on the elasticity of firm value with respect to ownership. A 1% increase in CEO ownership increases firm value between 0.008% and 0.090%, or \$0.560 million and \$6.274 million, respectively. A 1% decrease in ownership decreases firm value between 0.014% and 0.160%, or \$0.964 million and \$11.171 million, respectively. A best estimate of the actual effect is that an increase in ownership of 1% has a 0.079% effect on value. In well-governed firms, this potential increase is greatly diminished, as boards for these firms have chosen CEO ownership so that its net marginal benefits are near zero. In contrast, I find that in poorly-governed firms, CEOs hold too much equity. Increasing the ownership of the CEO of a poorly-governed firm leads to a decrease in firm value for two reasons. First, the average CEO of a poorly-governed firm is already exerting a large amount of productive effort, so there is little scope for increasing it.

Second, the CEO's wealth is so dependent on firm value, since the level of CEO ownership is so high in these firms, that the CEO's desire to increase leisure overcomes the desire to substitute higher wealth for leisure.

Much of the previous literature focuses on the absolute relation between CEO ownership and firm value. Morck, Shleifer, and Vishny (1988) document a non-monotonic relation between board ownership and firm value in the cross-section of firms. Similarly, McConnell and Servaes (1990) show a hump-shaped relation between Tobin's Q and the level of insider ownership, in which value is at its maximum when insider ownership is 37.6% (based on their 1986 sample). Jensen and Murphy (1990) argue that the pay of most CEOs in their sample is not sufficiently sensitive to changes in shareholder wealth, and that an increase in ownership would increase firm value. Results in Himmelberg, Hubbard, and Palia (1999) and Palia (2001) imply that, once the equilibrium nature of the incentive process is taken into account, there is no gain to increasing ownership. They find that heterogeneity among firms leads them to different value-maximizing CEO ownership levels, meaning the cross-sectional relation between ownership and firm value has no implication for the right level of CEO ownership in any specific firm. While this line of reasoning is valid, their results are susceptible to a form of the Lucas (1976) critique: the estimated regression parameters are not structural, so they cannot be used to evaluate changes in policy. In contrast, I estimate the structural parameters and change the board's equilibrium policy, which allows me to show the sensitivity of firm value to changes in ownership.

The papers closest to mine are Margiotta and Miller (2000) and Coles, Lemmon, and Meschke (forthcoming). Margiotta and Miller (2000) show that low incentive levels are sufficient to prevent CEOs from incurring large losses through shirking. Their objective is to show that low levels of CEO ownership are reasonable in the light of Jensen and Murphy's

(1990) concerns. Coles et al. (forthcoming) estimate a structural model based on Holmstrom and Milgrom (1987) to show that the hump-shaped relation from McConnell and Servaes (1990) is the outcome of firms with different productivity maximizing their value. In contrast, the current study's focus is on the sensitivity of firm value to changes in ownership, not the cross-sectional relation between value and ownership.

Other papers that estimate structural models of CEO compensation are Gayle and Miller (2009a,b), and Taylor (2010b). Gayle and Miller (2009a) estimate a moral hazard model to show that growth in firm size explains trends in executive compensation, but they do not directly address incentive compensation. Gayle and Miller (2009b) compare estimates of a moral hazard model with those of a hybrid moral hazard and private information model. They show that the hybrid model fits the data better, arguing that private information is an important determinant of CEO compensation. Taylor (2010b) focues on the sensitivity of total compensation to new information about the CEO. Dittman and Maug (2007) calibrate a static structural model to find the best mix of straight equity, stock options, and cash compensation that keeps CEO effort at the same level while decreasing the cost to the firm. They find that most CEOs should hold more straight equity, hold no stock options, and receive lower salaries.

The remainder of the paper is organized as follows. I develop the model in Section I. Section II describes the data set. The results are discussed in Section III. Section IV gives the results of counterfactual experiments. Section V reports the results from estimating the model on some subsamples. Robustness checks of assumptions used in the model are discussed in Section VI. Section VII concludes.

I Model

In this section I develop a model of the setting of CEO equity incentives. The model describes the interaction between the firm's board of directors and its CEO. The firm's CEO chooses his level of effort based on his incentives, which come primarily from the amount of firm ownership he has. The CEO faces a dynamic and finite horizon problem, in that he is hired and then works for the firm up to T years, but no longer. His desire to work changes over the course of his career as he faces the horizon problem when approaching retirement. The board seeks to maximize the future value of the firm (future value, as current value is already determined) by giving the CEO equity compensation. This compensation is designed to increase the CEO's effort, but it also decreases shareholder value by the size of the grant. In addition, the board has a preference for increasing the CEO's ownership, which leads it to give the CEO more or less ownership than naïve shareholder maximization requires.

I.A Timing and economic environment

The model is set in discrete time, with each period corresponding to a year. At the beginning of period t, the CEO receives a signal of the firm's potential end-of-period market-to-book ratio $\widehat{M_t}$. He then chooses his level of effort for this period a_t, which is in the compact set $[0, 1]$. The firm realizes its market-to-book ratio $M_t = a_t\widehat{M_t}$, and market value M_tB_t. The CEO then receives his income for the period from his salary, dividends (both from outside investments and the firm), and any sales of his current ownership. At this point the CEO may retire, a possibility discussed in more detail below. Once the CEO's employment is resolved, the board of directors chooses the level of ownership the CEO will have in the next period e_{t+1}. At this point period t ends and period $t + 1$ begins.

The two decisions in the model are the CEO's effort choice and the board's CEO ownership choice. The objective is to understand the effect that these decisions have on firm value. The firm's value has two parts: firm assets B and market-to-book ratio M. As the firm's investment decisions are not the focus of this paper, I follow Taylor (2010a) and hold firm assets constant through time. As such, I remove the time subscript from B going forward. The firm's market-to-book ratio follows an AR(1) process in logs:

$$\ln M_{t+1} = \alpha + \ln a_{t+1} + \rho \ln M_t + \varepsilon_{t+1}, \tag{1}$$

where α is a constant that determines the level of M, ρ is the autocorrelation coefficient, and ε_{t+1} is a normally-distributed random error with mean 0 and standard deviation σ. Equation (1) shows the effect of CEO effort; it impacts the firm's return between periods t and $t+1$. This captures the essence of effort in the model; CEO effort is the sum total of actions the CEO takes to increase firm value. When the CEO does not take the available actions, or takes other actions that do not increase firm value, then his effort is lower; he shirks. Shirking affects the firm by decreasing shareholders' returns.

The board of directors chooses the CEO's contract at the end of each period, after the CEO has exerted effort, while the CEO makes decisions during each period. The CEO's decisions are fully informed; he knows how much of the firm he owns and the firm's value. The board makes its decision with only partial information; it does not know how much effort the CEO will exert for a given level of ownership, because the CEO makes that decision in the next period after he sees $\widehat{M}_{t+1}$. This timing makes sense intuitively.[1] A firm's board meets only a few times each year, while CEO compensation is usually set once a year. The

[1] It also makes sense theoretically. Edmans and Gabaix (Forthcoming) argue that, in a similar, though more general, model, this timing convention improves tractibility of solving the model. In addition, they show that this timing convention provides similar results to a continuous time model in which shocks and decisions are all simultaneous.

CEO, on the other hand, is more involved in the daily running of the firm. If the firm receives a positive shock to value, the CEO can quickly make decisions that take advantage of this information.

I.B The CEO's decision

The CEO is a forward-looking agent who maximizes his expected lifetime utility through his level of effort. His lifetime utility is the discounted sum of individual period utilities for each year of his tenure as CEO and the utility he receives from consuming his wealth in retirement. There is a finite number of years, denoted T, that a CEO can be employed; upon reaching a tenure of T years, the CEO retires and receives his retirement utility. In every period t before the CEO reaches maximum tenure, there is a probability $\chi(\mathbf{x}_t, a_t)$ that he separates from the firm, at which point he receives his retirement utility. The separation probability is a function of the model's state variables, which are denoted by the vector $\mathbf{x}_t$. For the CEO, the state variables are the firm's size B, the market-to-book signal $\widehat{M}_t$, his ownership e_t, and the length of his tenure with the firm, τ.

Each period the CEO consumes his income, which is the sum of his salary $f(B)$, the dividend paid on his outside wealth w, his share of the firm's payout de_tM_tB, and revenue from selling a portion of his firm ownership $\delta(\mathbf{x}_t, a_t)\, e_tM_tB$. His outside wealth, W, is constant through time, and the dividend he receives from outside wealth equals his total wealth times the discount rate r: $w = rW$. The firm's payout is a constant fraction d of firm value, so the size of the payout fluctuates with changes in firm value. The proportion of shares the CEO sells, $\delta(\mathbf{x}_t, a_t)$, is a function of the state variables. Here salary is all cash payments the firm gives the CEO and is a solely a function of firm size. I ignore cash-based incentive compensation, such as bonuses and changes in salary. Ignoring bonuses is not detrimental to the model for three reasons. First, cash compensation is only partially driven by incentives. It reflects changes in beliefs about the CEO's ability [as in Taylor (2010a)], and/or changes in the CEO's negotiating power (e.g., if he receives an outside job offer). Second, the bulk of the CEO's incentives come from changes in the value of the CEO's stock holdings, not bonuses [see Hall and Liebman (1998) and Core, Guay and Verrechia (2003)]. And third, some CEO bonuses are based on factors that are questionably related to increasing firm value, for example M&A activity [see Grinstein and Hribar (2004)].[2] A CEO's incentives are determined primarily by his ownership, so the model focuses on ownership alone.

The CEO receives disutility from exerting effort. The cost of effort enters multiplicatively. When thinking about not exerting effort (i.e. shirking), as the consumption of a good, then a multiplicative cost of effort makes more intuitive sense than an additive cost of effort. The basic argument, as given in Edmans et al. (2009), is, assuming there are only two goods, a consumption good and shirking, an increase in a CEO's wealth should lead to greater consumption of both good types. For example, if shirking is thought of as leisure (not exerting effort), then greater wealth will have both a substitution effect (the CEO wants to divert consumption of leisure to consumption of the good) and an income effect (because leisure has become more valuable since the CEO's labor endowment has increased). An additive model does not have an income effect, so as wealth increases, shirking goes to zero. In this manner, shirking should be treated like any other normal good, in that it should be sensitive to changes in CEO wealth. Getting this wealth effect requires a multiplicative utility function.

The CEO has a constant relative risk aversion utility function. His single-period utility

[2] Murphy (1999, p. 2506): "Although virtually all bonus plans provide incentives to increase company profits, ...[their makeup] suggest[s] a plethora of additional incentives, most conflicting with stated company objectives."

comes from his consumption, c_t, and the cost he pays for making effort, a_t. The single-period utility for the CEO is:

$$u\left(e_t, \widehat{M}_t, B, a_t\right) = \frac{\left[c_t\left(1-a_t\right)^{\eta}\right]^{1-\gamma}}{1-\gamma}. \tag{2}$$

Here the coefficient of relative risk aversion is γ, and the cost of effort parameter is η. The cost of effort parameter controls the curvature of the CEO's effort cost, or, in other words, how painful it is to the CEO to increase his effort. If $\eta = 0$, then the CEO has no disutility of effort. If $\eta = 1$, then the disutility of effort is linear with a slope of -1; increasing effort is equally costly for any effort level. For $0 < \eta < 1$, the disutility of effort increases with the effort the CEO exerts; it is less painful to increase effort when effort is low, and more painful when effort is high. Just how quickly this pain enters into the CEO's utility is governed by the curvature of the effort cost function.

At retirement the CEO uses his existing wealth to purchase an annuity that pays off over his remaining expected lifetime. If he is forced to retire because he reaches the maximum tenure, T, then his post-retirement length is N periods. If he separates from the firm before he reaches the maximum tenure, his post-retirement length is $N^R = N+T-\tau$. His retirement utility is then:

$$U^R\left(\mathbf{x}_t, a_t\right) = \frac{1-\left(1+r\right)^{-N^R}}{r} * \frac{c_R^{1-\gamma}}{1-\gamma}, \tag{3}$$

where $c_R = \frac{1}{N^R}\left[W + \left(1-\delta_{t-1}\right)e_{t-1}M_{t-1}B\right]$ is the amount the CEO consumes in each period of retirement.

The CEO solves a finite-horizon problem when deciding how much effort to expend each period. His problem is:

$$U(\mathbf{x}_t) = \max_{a_t} u(\mathbf{x}_t, a_t) + \begin{cases} [1-\chi(\mathbf{x}_t, a_t)]\,\beta E[U(\mathbf{x}_{t+1})\,|\mathbf{x}_t, a_t] + \chi(\mathbf{x}_t, a_t)\,\beta U^R(\mathbf{x}_t, a_t), & \tau < T \\ \beta U^R(\mathbf{x}_t, a_t), & \tau = T \end{cases}. \tag{4}$$

I.C The board's decision

The solution to equation (4) depends on the CEO's expectations of his future ownership percentage of the firm, which is determined by the board of directors. As the focus of this paper is on the CEO's effort choice, I keep the model of the board's decision simple. The board seeks to maximize firm value by incentivizing the CEO to exert productive effort. The cost of incentivizing the CEO is the value of the new shares given to him. The board faces the constraint that it cannot take ownership away from the CEO; it can only add to it or leave it the same. In addition, the board has a preference for giving equity to the CEO, which affects how much ownership it grants. This preference could be positive, negative, or neutral. If it is positive (negative), then the board increases the CEO's ownership by more (less) than a simple, static model of shareholder-value maximization implies. I assume the board is risk neutral. The board of director's maximization problem is:

$$V(\mathbf{y}_t) = \max_{E(a_{t+1}), e_{t+1}} \beta E M_{t+1} B - \lambda\left[e_{t+1} - (1-\delta(\mathbf{x}_t, a_t))\,e_t\right] M_t B - f(B), \tag{5}$$

$$\text{subject to } e_{t+1} \geq (1-\delta(\mathbf{x}_t, a_t))\,e_t$$

$$E(a_{t+1}) = \arg\max_{a_{t+1}} E[U(\mathbf{x}_{t+1})]$$

where $\mathbf{y}_t$ is a vector of state variables for the board, and λ is the preference parameter for equity. For $\lambda > 1$, the board has negative preference for giving the CEO ownership (i.e. the board pays a personal cost when increasing the CEO's ownership). For $\lambda = 1$, the preference is neutral. The state variables for the board, contained in vector $\mathbf{y}_t$, are the firm's size B, the market-to-book ratio M_t, CEO ownership e_t, and the length of the CEO's tenure with the firm τ. The board does not see the CEO's signal of firm value $\widehat{M}_t$, only the realized market-to-book ratio M_t.

The board makes its compensation decision in the period before the CEO chooses the effort level, so the incentive compatibility constraint is in reference to the expected level of CEO effort. Equation (5) does not have a participation constraint for the CEO. My focus is on the incentives a CEO receives from his ownership, not on the level of CEO compensation or its structure. For simplicity I assume the CEO's cash compensation is exogenously determined and the participation constraint is not binding.

I.D Solving the model

Assuming CEOs are risk and effort averse, $\eta \in (0, 1)$ and $\gamma > 0$, each single-period utility for the CEO is strictly concave in effort no matter the state. With each single-period utility concave, the sum of all of them is also strictly concave. Thus the CEO's problem has a unique optimum for any given state, which can be found using backwards induction. However, this solution is not analytically tractable. I cannot show that the board's problem has a unique solution, as it depends on the expected value of the derivative of the CEO's policy function with respect to ownership, which is not necessarily monotonic. To avoid this issue, I follow Edmans et al. (2009) in assuming that, if there are multiple solutions to the board's problem, the board chooses the one that grants the CEO the least equity.[3]

The solution to the board's problem is easy to characterize. The board solves a static problem. Although a_{t+1} is a function of e_{t+1}, I do not explicitly include that dependence. The board chooses the CEO's next period equity holdings $e_{t+1} = \max\{(1 - \delta(\mathbf{x}_t, a_t)) e_t, e^*\}$, where e^* solves the first-order condition:

$$\beta EM_{t+1,e_{t+1}} = \lambda M_t. \tag{6}$$

Here $EM_{t+1,e_{t+1}}$ is the first derivative of the expected firm market-to-book ratio one period in the future with respect to CEO ownership. Equation (6) states that the board chooses the level of equity, which sets the marginal increase in expected discounted firm value equal to the marginal cost of giving that equity.

The CEO's effort for a given level of ownership varies with his time until retirement. When $\tau = T$, the CEO knows he will retire at the end of the period. In this case he makes an simple comparison of the costs of exerting effort against the benefit of increased current income and greater wealth to draw from in retirement, using his effort choice to equate the marginal cost and benefit. He chooses effort such that:

$$-\frac{u_a\left(e_t, \widehat{M}_t, B, a_t\right)}{\beta} = U_a^R\left(\mathbf{x}_t, a_t\right). \tag{7}$$

Looking one period in the past, when $\tau = T - 1$, the first-order condition is:

[3] At reasonable parameter values, the number of states possibly affected by multiple equilibria is so small as to have no appreciable effect. The results are similar using the assumption that the board chooses the solution that gives the CEO the most equity.

$$-\frac{u_a\left(e_t,\widehat{M}_t,B,a_t\right)}{\beta} = [1-\chi(\mathbf{x}_t,a_t)]\,E\left[u_a\left(e_{t+1},\widehat{M}_{t+1},B,a_{t+1}\right)+\beta U_a^R(\mathbf{x}_{t+1},a_{t+1})\,|\mathbf{x}_t,a_t\right] \tag{8}$$
$$-\chi_a(\mathbf{x}_t,a_t)\,E\left[u\left(e_{t+1},\widehat{M}_{t+1},B,a_{t+1}\right)+\beta U^R(\mathbf{x}_{t+1},a_{t+1})\,|\mathbf{x}_t,a_t\right]$$
$$+\chi(\mathbf{x}_t,a_t)\,U_a^R(\mathbf{x}_t,a_t)+\chi_a(\mathbf{x}_t,a_t)\,U^R(\mathbf{x}_t,a_t).$$

The simple result of equation (7) becomes more complex in equation (8). The CEO must balance the cost of increasing effort now against the effects it has on the probability he separates from the firm at the end of the current period, his expected utility next period, his choice next period, and his final wealth (whether he separates from the firm in the current period or retires in the next). His choice has significant dynamic consequences. For a CEO early in his tenure, these dynamic factors weigh heavily on his decision-making process, as there are as many as nineteen future periods affected by his effort. They also make the solution intractable.

In order to simulate firms, I solve the model numerically using backwards induction. I first solve the CEO's problem for his last period of employment, then the board's decision in the period before that, followed by the CEO's decision one period before retirement, then the previous period's board decision, and so on.

I.E Comparative statics

To illustrate how the model works, I solve it using the following parameter values: discount rate $r = 0.02\%$, CEO risk aversion $\gamma = 2.50$, CEO effort cost $\eta = 0.35$, board equity preference $\lambda = 1.54$, maximum effort return $\alpha = 0.04$, market-to-book persistence $\rho = 0.90$, and volatility of market-to-book shocks $\sigma = 0.40$. These values result in average firm values and CEO ownership levels that are similar to the data. I assume that CEOs are employed for a maximum of $T = 20$ years, and a CEO who works for the entire 20 years has a retirement that lasts $N_R = 20$ years. I follow Hall and Murphy (2002) in giving each CEO \$5 million of outside wealth (W), worth \$0.1 million of outside yearly income (w). Firms pay out $d = 1\%$ of their firm value to shareholders each year. I use estimates from the data, given in the Estimation section below, for the CEO salary, CEO ownership sales, and separation probability functions.

Panel A in Figure 1 shows the evolution of the natural log of the CEOs effort,[4] while Panel B shows the evolution of CEO percentage ownership across tenure. Effort is lowest at the beginning of a CEO's career and increases until the fifth year before plateauing. The board increases ownership for most of the CEO's tenure to counteract the horizon problem that comes from mandatory retirement and the increasing probability of early CEO separation. Dai, Jin, and Zhang (2010) find this same pattern: CEO ownership gradually increases over a CEO's tenure, along with an increase in firm value correlated with the ownership increase. Panel C of Figure 1 compares effort for low, middle, and high market-to-book ratios, while Panel D shows the same for average CEO equity. When firm value is high, the CEO makes the most effort, putting forth the least effort when firm value is low. The board partially counteracts this lower effort by giving the highest levels of equity to the CEO when firm value is low, and the lowest levels when firm value is high. Panel E of Figure 1 shows that CEO effort increases in firm size, while Panel F shows that the percentage of the firm owned by the CEO decreases with firm size. Effort increases with firm size because the value of CEO ownership is larger for a big firm than a small one. In this simulation, the average dollar value of CEO ownership is \$12.67 million in small firms, and \$486.61 million in large

[4] I use the natural log of effort because it is the percent impact that effort has on firm value, and is thus easier to interpret.

firms, despite the respective ownership percentages being 5.0% and 2.26%, respectively.

Figure 1 shows the results of a simulation based on one set of parameters. Figure 2 reports comparative statics from varying model parameters. While the persistence ρ and volatility σ of market-to-book do have an effect on effort and ownership decisions, their main effect is on market-to-book itself, so I focus on the other four parameters. As α is the intercept term for the evolution of market-to=book, increasing it increases average firm value. Panels A and B in Figure 2 show that increasing firm value has the expected effect: CEO effort increases and ownership decreases. CEOs work harder when their shares are more valuable, and the board does not have to grant the CEO as much because of the higher value of the shares.

Panels C through H in Figure 2 show the effect on CEO effort and CEO ownership from changing CEO risk aversion γ, CEO effort cost η, and board efficiency λ. These plots indicate that each parameter has a unique effect on ownership and firm value, which is useful in identifying them empirically. Average CEO effort is non-monotonic as risk aversion increases; it decreases for values of risk aversion below 2.50, and increases for risk aversion above 2.50. Consistent with previous theory [e.g., Holmstrom and Milgrom (1987)], CEO ownership decreases in the level of risk aversion. Increasing the cost of effort decreases effort. The board reacts to lower effort by increasing incentives, so CEO ownership increases with the cost of effort.[5] As λ increases, decreasing the board's preference for giving equity, CEO ownership decreases, which leads to lower CEO effort as well. The cause of the non-monotonicity in CEO effort as risk aversion increases provides a way to identify it separately from effort cost and board efficiency. Given the value of risk aversion, the cost of effort and board efficiency have the same effect on CEO ownership, but the opposite effect on effort

[5] CEO ownership increasing in the cost of effort is the consequence of assuming the participation constraint is not binding and a multiplicative cost of effort. Without a constraint on the level of CEO income, the board is free to increase incentives to counteract the effect of higher cost of effort. The board does so because it is offsetting both a substitution effect and an income effect.

level. This tension allows for them to be separated empirically.

The non-monotonicity of the relation between CEO effort and risk aversion comes from the different ways that low and high risk-averse CEOs react to incentive changes. Panels I and J of Figure 2 plot average CEO effort and CEO ownership across CEO tenure for $\gamma = \{1.50, 2.50, 3.50\}$. When risk aversion is low, $\gamma = 1.50$, effort begins low then increases across the CEO's tenure. His effort is high at the end of his career because, having low risk aversion, he is willing to hold more ownership. Early in his career he shirks more because he is willing to accept the higher probability of dismissal from the firm having a low value. The CEO knows that he will work hard late in his career, so he does not work as hard early on because he is maximizing his lifetime utility. A CEO with higher risk aversion is not willing to accept the higher probability of dismissal from having a low firm value, and does not particularly want to increase the volatility of his income by holding a lot of shares of the firm. These two implications of high risk aversion lead his effort to start high, and then decrease throughout his tenure as his ownership does not grow like they do for a CEO with lower risk aversion. These effects are then exacerbated by the effect of early CEO separations. The probability of the CEO leaving the firm is higher for low value firms, so the CEOs remaining in the sample the longest also work for more valuable firms. More valuable equity increases the volatility of expected utility, which has a negative effect on very risk-averse CEOs exerting effort. The non-monotonicity of effort with risk aversion is thus caused by high early effort for high levels of risk aversion pulling the average up, and high late effort for low risk aversion pulling the average effort up – low average effort occurs when risk aversion is such that effort is fairly constant throughout the CEO's tenure.

II Data

The sample covers publicly-traded U.S. firms over the 1992 to 2007 period. Data on CEO ownership, compensation, and tenure are from Execucomp, firm data from the Annual Compustat file, and return data from CRSP.

Execucomp does not consistently label the CEO for each firm, so I use the date when an executive became CEO to determine who is the CEO of each firm-year. From Execucomp I extract the number of shares held by the CEO at year end, new and existing options holdings, and cash compensation (calculated as all compensation not in the form of equity: *total_curr*, *noneq_incent*, *othcomp*, and *allothtot*). To determine the level of CEO ownership, I combine shares held with the share-equivalence of stock options, found using the method of Core and Guay (1999) [described even more fully in Edmans et al. (2009)]. One difference between my method and that used by Edmans et al. (2009) is that I compute stock volatility, and thus the upper and lower bounds on expected stock volatility, using the entire CRSP database, not the Execucomp subsample. Beginning in 2006, more complete data on CEO option holdings is available, which allows me to calculate option deltas without the assumptions on existing holdings used by Core and Guay (1999). I calculate the amount a CEO sells in a period in the same way as Core and Guay (2010). CEO sales in a given year are the value of end-of-year equity holdings minus the value of new equity grants minus what beginning-of-year equity holdings would be worth at end-of-the-year prices if the CEO's ownership did not change. This value can be positive (CEO buys more equity with his own money) or negative (CEO sells equity). When this variable is positive, I set the value to zero, as only the effect on a CEO's incentives from selling equity is of interest in this study. From Compustat I extract the value of firm assets, market value, and distributions.

I remove firms with missing values, negative or zero assets or market values, and negative equity ownership. I adjust firm size for inflation using the CPI (dollar amounts are in year 2000 dollars). I define market-to-book as firm market value over firm assets, which differs from its traditional definition. The value of CEO stock holdings is based on the equity market value of the firm, not the value of the firm as a whole, which makes my definition more appropriate here. I winsorize market-to-book at the 1% level. Because of some large outliers, I trim the percentage of CEO ownership at the 1% level. Some of the moments I use to estimate the model use the level of CEO ownership in the previous year. I do not have prior holdings information for CEOs in their first year in office, so those observations are removed. This leaves a sample of 17,244 firm-CEO years. Descriptive statistics are in Table 1.

III Estimation

This section describes the identification and estimation of the model. I estimate the parameters of primary interest using SMM estimation. Other parameters I estimate outside of SMM or set to a reasonable value. I set the discount rate r to 2.00%, which is an approximation of the risk-free interest rate. Outside CEO wealth is set to \$5 million, as in Hall and Murphy (2002), which means CEOs receive periodic income w of \$0.100 million each period. The average dividend yield, d, is 1.00%. OLS estimates from regressing the log of yearly cash compensation (salary and other cash pay) on the log of firm assets yield the salary function:

$$f(B) = 0.076B^{0.384}. \tag{9}$$

A logit model for the probability of a CEO leaving his position in a given year, with the probability being a function of firm size, market-to-book, and CEO tenure gives the probability a CEO separates from his firm in a given year:

$$P(\chi_t = 1) = [1 + \exp\left(2.110 + 0.184 \ln B + 0.012 (\ln B)^2 + 0.232 \ln M_t - 0.060 (\ln M_t)^2 + 0.033\tau_t - 0.001\tau_t^2\right)]^{-1}. \tag{10}$$

The proportion of firm equity a CEO sells each period, δ, comes from a Tobit model, as the proportion of shares sold cannot be less than 0. This proportion is a function of the firm's market-to-book ratio, the level of CEO ownership entering the period, and CEO tenure [both first and second order, because of curvature in the relation, as demonstrated in Core and Guay (2010)]. Firm size is not a regressor because it does not have an effect once these other variables are included in the regression model. Each period the CEO sells a proportion of his ownership equal to:

$$\delta(M, e, \tau) = \max\left\{0, 0.1719 + 0.0319 \ln M + 0.0424 \ln e + 0.0031\tau - 0.0001\tau^2\right\}. \tag{11}$$

To identify the remaining six parameters, I use the observed relation between CEO equity ownership levels and firm and CEO characteristics. I describe the state space and my solution method in the appendix.

III.A Identification

I estimate the parameter vector $\theta = (\rho, \sigma, \alpha, \gamma, \eta, \lambda)'$ using SMM. Identification of the first three of these parameters comes from their empirical counterparts. I estimate the following AR(1) regression:

$$\ln MarketToBook_{i,t+1} = \psi_0 + \psi_1 \ln MarketToBook_{it} + \zeta_{i,t+1}. \tag{12}$$

Let υ^2 be the variance of $\zeta_{i,t+1}$. Recall from equation (1) that the log of market-to-book follows an AR(1) process in the model. Thus ρ is identified by ψ_1, and σ is identified by υ^2. The mean of the logarithm of market-to-book identifies α, as there is a one-to-one relationship between the intercept of an AR(1) model and the unconditional mean of the variable.

The remaining three parameters have to do with the CEO's effort and the board's incentive decisions. I estimate these for an average CEO, so it is how effort and ownership vary and covary for each CEO, not across CEOs, that provides the necessary identification. Thus I use within-firm moments. Using within-firm variation controls for any differences that come about from differences in CEO ability across firms, although it does not control for within-firm changes in ability or changes in beliefs about ability. The key to identification is in differentiating between risk aversion, cost of effort, and board preference empirically. The comparative statics section above provides guidance for this differentiation.

Changes in risk aversion change the slope of effort with respect to CEO tenure, while changes in the cost of effort and board preference have no effect on this effort slope. While CEO effort is unobservable, it affects changes in a firm's valuation, so changes in a firm's log market-to-book ratio ($\Delta^{\ln M} = \ln M_t - \ln M_{t-1}$) capture the effects of effort. Risk aversion determines the covariance of effort and CEO tenure $\text{cov}\left(\Delta^{\ln M}, \tau\right)$, so this covariance is the first moment used to identify risk aversion. How this moment changes for different values of γ, η, and λ is shown in Figure 3 panel A. As risk aversion increases, $\text{cov}\left(\Delta^{\ln M}, \tau\right)$ decreases strongly. The moment is basically unaffected by changes in λ, and it increases slightly with η, although the slope of this line is small compared to the risk aversion plot. Matching this moment identifies risk aversion, as only risk aversion causes enough variation to set the simulated model moment equal to the data moment.

Additional identification of risk aversion comes from the covariance of CEO ownership and firm value. As risk aversion increases, the spread of CEO ownership between high and low market-to-book firms narrows. The next two moments come from the covariance of CEO ownership and the firm's market-to-book ratio, both in logs. A CEO's effort can either increase or decrease during his tenure depending on the value of γ, so this covariance is computed for early- and late-career CEOs. Panel B of Figure 3 plots this covariance for CEO tenure less than or equal to ten years $\text{cov}\left(\ln e, \ln M | \tau \leq 10\right)$, and Panel C plots the same moment for CEO tenure over ten years $\text{cov}\left(\ln e, \ln M | \tau > 10\right)$. This moment increases strongly with risk aversion, while its movement with the cost of effort and board preference is not monotonic. In the data, both of these moments are positive. Figure 3 show that the simulated model only produces positive values for these moments if risk aversion is high enough. For no value of the cost of effort is this moment positive, and changes in board preference have little effect on the simulated moment. The plots in Panel C tell a similar story to the one Panel B tells, except there is a strong decrease in this covariance for a range of values of the cost of effort. This decrease in the moment for values of the cost of effort greater than 0.30 suggests the cost of effort is in the bottom half of its range. Matching these three moments – $\text{cov}\left(\Delta^{\ln M}, \tau\right)$, $\text{cov}\left(\ln e, \ln M | \tau \leq 10\right)$, and $\text{cov}\left(\ln e, \ln M | \tau > 10\right)$ – identifies γ separately from η and λ.

Taking the identification of CEO risk aversion as given, I can separate the cost of effort from board preference. Increases in the cost of effort and board preference have the opposite effect on the within variance of CEO ownership, $V\left(e\right)$, as is shown in Panel D of Figure 3. The relation of effort to ownership is decreasing in both of these parameters, but for different reasons. For a given level of ownership, a CEO chooses a lower effort level as the cost of effort increases. As the board preference parameter increases (decreasing the board's willingness to increase the CEO's ownership), a CEO with a given level of ownership expects less future ownership, lowering his current desire to work. Figure 3 Panel E plots how the moment cov $\left(\Delta^{\ln M}, \ln e_{-1}\right)$, changes with increases in η and λ, which captures the effect of the covariance of effort and incentives. Note that the moment uses the previous period's CEO ownership, since the incentives that matters for the current period's effort choice are determined in the previous period. The moment is weakly increasing in the cost of effort, as well as in the board preference parameter.[6] Changes in these two parameters drive $V\left(e\right)$ in two directions, while driving cov $\left(\Delta^{\ln M}, \ln e_{-1}\right)$ in the same direction, which provides the tension necessary to separate the cost of effort from board ownership preference.

For additional identification, I use both the covariance of changes in the logarithm of market-to-book and changes in the logarithm of CEO ownership for early and late tenures: cov $\left(\Delta^{\ln M}, \Delta^{\ln e}|\tau \leq 10\right)$ and cov $\left(\Delta^{\ln M}, \Delta^{\ln e}|\tau > 10\right)$. These two moments are informative about risk aversion, the cost of effort, and board efficiency, as they deal with how sensitive changes to ownership are to changes in firm value. The last moment I use in estimating the model is the mean of CEO percentage ownership, $E\left(e\right)$. The basic research question is how much incentive ownership gives CEOs, so it is important to match the average level of incentives. Mean CEO ownership is the only moment aside from the AR(1) regression that is calculated from the entire panel, and not as a within estimate.

[6] Note as well that both of these moments, $V\left(e\right)$ and $cov\left(\Delta^{\ln M}, \ln e_{-1}\right)$ are decreasing in γ. This highlights that they are useful in separating η and λ from each other.

III.B Estimation results

Panel B of Table 2 contains the estimated parameters. The point estimate of the autocorrelation of log market-to-book (ρ) is similar to the data moment. The standard deviation of the shocks to log market-to-book ($\sigma = 0.3933$) is smaller than the data moment ($\upsilon = 0.4215$). The volatility of shocks in the AR(1) regression is larger than the actual shock because it includes variation from changes in CEO effort as well as true exogenous shocks. Changes in CEO effort make up 1.03% of overall volatility, and its covariance with shocks to market-to-book make up another 8.66%, which brings the net proportion of market-to-book volatility from changes in CEO effort to 9.69%.

The estimate for α is 0.0342, which represents the maximum potential increase in return a CEO expect. There is no heterogeneity for this parameter in the model, so this point estimate captures the average effectiveness of CEOs in the sample, along with other factors that lead to the observed level of market-to-book. By itself it is difficult to interpret, as no CEO in the model puts forth full effort (i.e., $a = 1$). As α affects changes in a firm's market-to-book, looking at expected changes in firm value – returns, basically – helps in understanding the estimate. Ignoring cash distributions (which I assume are constant in percentage terms), the firm's expected return in period t is:

$$E \ln M_{t+1} - \ln M_t = \alpha + E \ln a_{t+1} - (1 - \rho) \ln M_t. \tag{13}$$

The mean log of CEO effort in the simulated model is -0.0548. Combining average effort with α, the portion of the return affected by the CEO contributes a mean annual return of $\alpha + E \ln a = -2.06\%$. On average, the CEO costs the firm about 2% of return per year through shirking. The model-predicted average market-to-book is lower than market-to-book in the data, which may imply a downward bias in the estimate.

The average CEO in the sample has a coefficient of relative risk aversion (γ) of 3.0647. For comparison, previous studies give a wide range of reasonable values for risk aversion, from 0 to 55 [see Table VII in Bliss and Panigirtzoglou (2004) for a partial list]. A more recent study, Chetty (2006), shows that risk aversion less than 2 fits some established facts about U.S. labor supply. Although my estimate of risk aversion is greater than 2, there is no a priori reason to believe that CEOs have the same level of risk aversion as the general labor supply. Previous studies on CEO compensation [e.g., Hall and Murphy (2002), Fernandes, Ferreira, Matos, and Murphy (2010)] assume CEO risk aversion is either 2 or 3, which is in-line with my estimate.

The estimate of effort cost parameter, η, is 0.3785. One way of understanding this estimate is to think about the elasticity of the cost of effort with respect to effort, which is given by $-\eta a/(1-a)$. Average CEO effort in the simulation is $a = 0.9474$, which coincides with a cost elasticity of just under 7 – if the CEO puts forth enough effort to increase the firm's return by 1% from its current level, his personal utility cost increases by almost 7%. So, to induce a CEO to exert more effort requires a benefit large enough to overcome this substantial disutility. Note that the maximum effort a CEO could exert would be $a = 1.0000$, but no CEO in the model would exert maximum effort because doing so would make the single-period utility negative infinity. CEOs already, on average, exert a substantial amount of effort, which means there is not a lot of scope for increasing it.

The last parameter (λ) captures the board's preference for giving ownership to the CEO. The estimate is 1.5375, implying that board members act as if giving equity compensation to the CEO is more expensive to them than it is to shareholders. There are two unmodelled contributors to this estimate. The first is the dynamic shadow price for granting equity. For simplicity, the board's optimization problem is static in the model. It is the board's job to maximize firm value, and the basics of this are captured in the static model. In some situations the board would optimally lower a CEO's ownership, but are unable to take away previously granted equity. There is a possible future cost to granting current ownership that is not captured by the static model, but would be in a dynamic model, which biases λ upward. Second, board members may feel they pay a personal cost when granting CEOs more ownership, such as the possibility of negative press when overall CEO compensation is high. This cost could also reflect the presence of other costs of incentivizing the CEO, such as the possibility of earnings' management. Whatever the underlying source, CEO ownership is less than a naïve model of shareholder maximization predicts.

III.C Model fit

Panel A of Table 2 compares the data moments to those estimated by the model. A formal test of overidentification, whether or not the model is able to match all of the moments, fails. The failure of the model to match all of the moments is unsurprising given the sample size, and the fact that some strong assumptions go into making the model estimable. The moments the model has the hardest time matching are the variance of CEO ownership, average log market-to-book, and the early and late covariance of CEO ownership and market-to-book. We learn more from looking at why the model is unable to match these individual moments.

The most statistically significant difference between simulated and data moments is for the within-firm variance of CEO percentage ownership. There are two main sources of variance for equity in the model: changes in firm value and CEO tenure. Obviously there are other real world shocks that affect the variation in CEO ownership. The two most obvious are shocks to the CEO's outside wealth, which is held constant in the model, and changes in beliefs about the CEO's ability [as in Taylor (2010b)]. One other possible reason for this inability to match ownership volatility is that the model assumes that all CEO ownership is held as straight stock, while in the data a significant portion of CEO equity is held as options. The value of options is generally more sensitive to changes in firm value than is the value of straight stock, which means that the stock equivalent value of options is more variable than that of straight stock. Including shocks to outside wealth, learning about CEO ability, or stock options separate from equity in the model would require adding at least one state variable, making estimation difficult. What is important here is that the model does match average CEO ownership. Matching average CEO ownership matters because the research question is on the level of CEO incentives.

The model's estimate of mean log market-to-book is significantly lower than for the data. The reason the model does not match this moment is connected to not being able to match the variance of CEO ownership. The way to increase average market-to-book is to increase α, the CEO impact parameter. However, the variance of CEO ownership is sharply decreasing in CEO impact, since increasing α both increases effort (which has an upper bound) and decreases CEO equity (because of higher firm values). Setting $\alpha = 0.0425$, which enables to model to match this moment, leads the variance of CEO ownership to fall to 0.0002. As above, it is the lack of variation in CEO ownership that causes the model to fail.

The other two moments the model is unable to match are the covariances of the log of CEO ownership and log market-to-book for early and late career CEOs. These two moments are used to identify risk aversion. Making a small change like setting $\gamma = 3.15$ is enough to match the early covariance, which is where the match is most statistically different. This change decreases the variance of CEO ownership to 0.0003. Such a small change in the risk aversion parameter is necessary to match this moment that the failure to match it is not concerning; the effect on identifying risk aversion is small. Again, it appears that the lack of enough sources of variance for CEO ownership keeps the model from matching the data. Since it is the level of incentives that matter, this inability to match ownership volatility should not largely affect the interpretation of the model's results.

Table 3 compares the estimated model against the data. Panel a compares CEO ownership in both percentages and in dollar terms, while Panel b compares firm valuation. The estimated model data does a reasonable job in replicating the actual data. Mean ownership is one of the moments I match, so it fits well, but the model estimates median ownership that is too high (2.80% vs. 1.57%). Median ownership is too high because boards for large firms grant too much CEO ownership in the model. For the model to simulate large firms that give the correct amount of ownership requires large firms (and their CEOs) to differ in some systematic way from smaller firms. For example, CEOs of large firms may have more outside wealth, may have different risk aversion [as in Baker and Hall (2004)], or may have less volatile firm returns. I examine the possible differences between large and small firms when I estimate the model using subsamples. The other differences between the simulated data and the real data are either connected to this problem with large firms, or the low median market-to-book ratio, which is the result of the model not matching the mean log market-to-book ratio.

IV Sensitivity of Effort to Incentives

While knowing the risk aversion and cost of effort for CEOs is interesting, we want to understand the effect of changing CEO ownership on firm value. If boards of directors determine CEO ownership by equating its marginal costs and benefits, then knowing the marginal benefit is the same as knowing the marginal cost. The explicit cost is the dollar amount of additional ownership given to the CEO, and anything more is the result of some unseen friction.

To place an upper and a lower bound on the marginal benefit of CEO ownership, I perform two related counterfactual experiments. In each experiment I alter the board's policy function to give CEOs more or less ownership, then allow CEOs to change their effort levels based on the new board policy. I do this by multiplying the estimated board policy function by $x\%$, with x values from 50 to 200. In every state a firm may experience, CEO ownership is $x\%$ of the estimated policy. I then update the CEO's policy function using backward induction as before. The difference between the two experiments lies in the timing of the change in policy. In the first experiment I unexpectedly change the policy during the simulation to observe effort and firm value before and after the change. In the second experiment I change the policy before the simulation, so that the new policy has always been in effect. In a sense, the second experiment is the ultimate outcome of the first, as the first would evolve into the second if the simulation is allowed to run for enough periods after the change.

Both CEO effort and firm value are sensitive to unexpected changes in CEO ownership, as Panels A and B of Figure 4 show, with the sensitivity being asymmetric for increases versus decreases. Panel A of Table 4 reports the elasticities of effort and firm value to changes in ownership. A rise in ownership of 1% increases the average log of CEO effort by 0.128% one period after the change (the "initial effect"), while a 1% decline decreases effort by 0.454%. These lead to an increase (decrease) in average firm value of 0.008% (0.014%). The change in policy is permanent, which means a CEO continues to work more (less) than the baseline in every period after the change. As a result, firm values continue to increase (decrease) with respect to the baseline. For example, in the tenth year after the change, firm values are 0.052% higher (0.101% lower) than the baseline. In dollar terms, a 1% increase in average CEO ownership is worth $1.688 million to the CEO when the board makes the change.[7] One year after the change, average firm value is $0.560 million higher than the baseline; ten years after it is $3.613 million higher.

The problem with interpreting the first experiment lies in the fact that the model does not have rational expectations on the part of investors in computing firm value. Investors are forward looking, so firm value one year after a change in compensation policy should impound the future effect of increased CEO effort, not only current effort. The second experiment mitigates this concern by showing the opposite extreme: the effect of this change in the limit (Panel B of Figure 4). Table 4 reports the elasticities. The sensitivities here are that an increase (decrease) in ownership of 1% increases (decreases) CEO effort by 0.151% (0.509%), and increases (decreases) firm value by 0.090% (0.160%). Average firm value is $6.274 million higher than the baseline for a 1% increase in ownership policy, while the CEO's ownership is worth $1.971 million more. Because this experiment is similar to impounding all of the increased future effort's effects in the current stock price, it is an upper bound for the effect of increasing CEO ownership.

The range of the elasticities is fairly large. To find a likely best estimate of the elasticity, I compute the present value of making the change by repeating the first experiment and letting it run forward until the differences from the baseline stabilize (i.e., it turns into the second experiment). Using a discount rate of 2% (the same used in estimating the model), I discount all of the differences back to the policy change and sum them. I also do this for the CEO's

[7] The majority of the change in ownership occurs at the time of the policy change. Given an increase (decrease) in ownership, 10 years after the change the CEO's ownership is worth $1.794 million more ($1.845 million less) than the baseline.

ownership. The 1% increase in the board's policy function leads to a present value change in the value of the firm of \$5.267 million, and a present value change in CEO ownership of \$1.875 million. The difference between the increase in firm value and the increase in CEO ownership is \$3.392 million. For a decrease in ownership, the decrease in firm value is \$9.606 million, the decrease in CEO ownership of \$1.960 million, and the difference between the two is \$7.647 million.

The counterfactual experiments give an upper and a lower bound on the sensitivity of CEO effort and firm value to changes in CEO ownership. Increasing CEO ownership by 1% leads to an increase in average firm value between 0.8 basis points and 9.0 basis points. Based on a best estimate for the average firm, the marginal benefit is \$5.267 million, with an explicit cost of dilution of \$1.875 million. The difference of \$3.392 million is the possible gain from increasing CEO ownership, which means it is also an estimate of the hidden costs to the board associated with increasing ownership.

V Subsample Estimation

The sample is split into subsamples based on four characteristics: size, options usage in CEO compensation, governance quality, and time.

V.A Large vs. small firms subsample

Since the estimated model is unable to reproduce the low level of CEO ownership for large firms, there must be systematic differences between large and small firms for the model to fail to match both in one estimation. To understand these differences, I estimate the model for the largest and smallest thirds of the sample, where size is measured by book assets (see

Panel A of Table 5).

Contrary to Baker and Hall (2004), CEOs of large firms in the sample are substantially less risk averse than the average CEO ($\gamma = 1.8469$ for large firms vs. $\gamma = 3.0647$ for the full sample); they also face a larger cost of effort ($\eta = 0.4239$) than the average CEO ($\eta = 0.3785$). This cost is reasonable, as large firms generally have more stable returns, so increasing their return by 1% is more difficult than increasing the return of a smaller firm by the same amount. The average market-to-book ratio is lower for large firms than for the full sample, which leads to a smaller maximum benefit from CEO effort α.[8] The estimate of board preference λ is 1.7637, which is larger than the estimate from the full sample; boards of large firms have a lower preference for giving the CEO ownership than the average board. This lower preference is the primary reason the model estimated for the whole sample is unable to match ownership for large firms.

I find the lower and upper bounds for increases in firm value from a 1% increase in CEO ownership of 0.004% and 0.053%. The bounds are smaller in magnitude than the bounds for the full sample, which reflects the greater cost of effort faced by CEOs of large firms. In dollar terms, a 1% increase in average ownership for these CEOs is worth \$7.616 million. The increase in average firm value is between \$1.789 million and \$26.421 million. Computing the present value of the increase in firm value minus the increase in dilution gives an increase of 0.032% (\$15.877 million) for a 1% increase in ownership.

The estimates for the small firm sample tell a different story. Firm valuation is less autocorrelated, more volatile, and of a higher level than the it is for the average firm. Despite these differences, CEOs of small firms are not very different from average CEOs. Their risk

[8] This does not necessarily imply that CEOs of large firms have lower ability than CEOs of smaller firms. In addition to CEO ability, α also captures any other reason for the level of firm valuation. That is why it is the maximum benefit from CEO effort, instead of a direct measure of CEO ability.

aversion is 3.2506, which is not much larger than the 3.0647 estimated from the full sample. Their cost of effort is 0.3498, which is slightly less than that of the full sample (0.3785). The largest difference is that the boards of small firms have a stronger preference for granting equity ($\lambda = 1.1689$ compared to $\lambda = 1.5375$ for the full sample). A 1% increase in ownership for small firms leads to an increase in average firm value between 0.014% and 0.090%. The upper bound is the same as the upper bound for the full sample, but the lower bound is almost double that of the full sample (and 3.5 times larger than for large firms). In dollar terms, a 1% increase in ownership for small firms is \$0.306 million, and the bounds for the increase in firm value are \$0.090 million and \$0.584 million. The relatively low ratio of the upper bound to the dollar increase in ownership reflects the less negative board preference. The present value of a1% change is an increase of 0.020% (\$0.126 million).

Large firms and small firms have different needs, which leads them to choose different types of CEOs. This heterogeneity is the cause of the model estimated on the full sample to struggle fitting some aspects of the data. In comparing large and small firms, counterfactual analysis shows that the value of small firms is more sensitive to changes in CEO ownership. Also, there appears to be an opportunity to increase the value of large firms by increasing CEO ownership, unless there is reason to believe that large firms face substantial hidden costs.

V.B More stock options vs. fewer stock options subsample

The largest statistical difference between the estimated moments and the data moments is the within variance of CEO ownership. CEO ownership in the estimated model does not vary enough. One of the possible reasons for this failure is that a significant number of CEOs receive stock options instead of straight equity as incentives. Stock options are more volatile than straight equity, so CEOs with a high proportion of stock options are likely to have more volatile ownership stakes. To test this idea, I split the sample into two groups: high stock options firms and low stock options firms. High (low) stock options firms are those whose CEOs have, on average during the sample period, more (less) than 50% of their ownership in stock options. These groups are exclusive, as there are relatively few observations for which the ownership of the CEO of a high stock option firm was made up of less than 50% options, and vice-versa.

This intuition that holding mostly stock options implies more volatile ownership does not match the data. In fact, the within volatility of CEO ownership is almost five times lower for CEOs whose ownership is comprised mostly of stock options (0.0001) than for CEOs with more straight equity than options (0.0006). This difference is because there are systematic differences between firms that grant a large amount of options and those that prefer to grant straight equity. Panel B of Table 5 shows the parameter estimates for these two subsamples. The two groups face roughly the same valuation environment: the autocorrelation and volatility of firm value hardly differ, and the maximum benefit from CEO effort is not that different. They differ in that CEOs of firms that use stock options heavily are more risk averse and face a much smaller cost of effort. As a result, the average level of CEO ownership is 1.87% for firms that heavily use stock options, and 6.27% for firms that do not. This difference is huge considering the difference in average firm size and firm value is not large between the two groups.

The results on board preference (λ) is mixed. Low options-using firms have a positive preference for giving their CEOs ownership; they give more than a naïve shareholder maximization strategy implies. The sensitivity of firm value for low options-using firms is between 0.006% and 0.067%. High options users have the same negative preference as average firms, although this result may not be as strong if options create stronger incentives than straight equity. The sensitivity of firm value for these firms is between 0.0003% and 0.024%.

The use of stock options does not appear to be an important factor in the model's inability to match the within variance of CEO ownership. In fact, the firms that use stock options the most have low volatility of CEO ownership. The two most likely causes for this inability are volatility of CEO outside wealth and changes in CEO ability or the perception of CEO ability.

V.C Corporate governance subsamples

Boards of directors have a negative preference for granting CEOs ownership (i.e., boards see granting ownership as more costly than naïve shareholder maximization predicts). This negative preference reflects the cost of granting ownership, which is made up of real costs to the firm (e.g., the shadow price of granting ownership, or concerns about the CEO manipulating public disclosures,[9] etc.) and personal costs to the board members (e.g., negative press from high CEO compensation). To understand the source of the board preference estimate, I split the sample according to two measures of corporate governance: the level of institutional ownership, and the level of blockholder ownership. Institutional owners are more likely to be active in communicating their preferences to the board of directors, so high levels of institutional ownership is a sign of better governance. Shleifer and Vishny (1986) argue that high levels of blockholder ownership make corporate takeovers and proxy contests easier. With a lower takeover cost, the threat of losing control of the company disciplines managers and boards. For both subsamples, I find the average level of institutional or blockholder ownership for each firm inside the sample. I then separate the samples into low and high

[9]For an example, see Bergstresser and Philippon (2006), who show the higher CEO incentives are correlated with greater discretionary accruals, which may be a sign of greater earnings management.

groups, where low (high) indicates the bottom (top) one-third of the distribution for the firm averages. I do this separation by firm averages because of the need to compute within-firm volatility and covariances.

High institutional ownership firms have lower mean book assets than have low institutional ownership firms, \$2.8 billion to \$6.2 billion, but larger median book assets, \$1.3 billion to \$0.8 billion. High institutional ownership firms are more valuable, with an average market-to-book of 1.73 compared to 1.65. Seeing that average CEO ownership is lower for high institutional ownership firms is not surprising, given their larger (median) size and valuations. Panel C of Table 5 gives the parameter estimates for these two subsamples. Firms with low institutional ownership have more volatile valuations than high institutional ownership firms, and their CEOs are more risk averse. High institutional ownership firms have significantly lower costs of effort. Lower costs mean the board of directors at a high institutional ownership firms can give less ownership incentives to the CEO, compared to how much is necessary at a comparable low institutional ownership firm. As to how boards decide on the level of ownership, firms with high institutional ownership have a similar negative preference as the whole sample. Low institutional ownership firms have a strong positive preference towards giving ownership to the CEO.

The average level of ownership is much higher for low institutional ownership firms, 5.2% of outstanding shares versus 2.5%, and they are biased towards over-incentivizing the CEO, which should increase effort. However, the average return from CEO effort for low institutional ownership firms, measured as the natural logarithm of CEO effort, is -2.33%, versus -1.85% for high institutional ownership firms. CEOs of high institutional ownership firms add an average of 50 basis points more value to their firm each year than CEOs of low institutional ownership firms, even though CEOs of low institutional ownership firms hold substantially more ownership.

The range of elasticities for high institutional ownership firms is 0.004 to 0.026 (\$0.202 million to \$1.343 million), while for low institutional ownership firms it is 0.004 to 0.037 (\$0.510 million to \$4.950 million). A 1% increase in ownership cost \$0.893 million for a high institutional ownership firm, and \$5.491 million for a low institutional ownership firm. For well-governed firms, the dollar value of a 1% change in ownership is large compared to the upper end of the increase in firm value range, implying that hidden costs for increasing CEO ownership are low for well-governed firms. For poorly governed firms, a 1% increase in ownership is worth \$5.491 million, which is larger than the upper bound for the increase in firm value. Firms with low institutional ownership appear to give their CEOs too much ownership, as the marginal cost exceeds even the most aggressive estimate of the marginal benefit.

Panel D of Table 5 shows the parameter estimates for the low and high blockholder ownership firms. High blockholder ownership firms have more volatile value, and more risk-averse CEOs than do low blockholder ownership firms. Similar to high institutional ownership firms, high blockholder ownership firms have CEOs with a much lower cost of effort than the CEOs of low blockholder ownership firms. Unlike with high versus low institutional ownership firms, high blockholder ownership firms have a more strongly positive preference for giving the CEO ownership than do low blockholder ownership firms, although this difference is only marginally statistically significant because of the sample size.

Similar to the results from the institutional investor sample split, CEOs of firms with high blockholder ownership put forth more effort (-1.42%) than CEOs with low blockholder ownership (-2.32%). Also similar is the sensitivity of firm value to changes in ownership for high blockholder ownership firms and high institutional ownership firms. The elasticity range is 0.003 to 0.027 ($0.125 million to $1.132 million), for a $0.769 million change in ownership. Low blockholder ownership firms have negative elasticities (-0.014 to -0.077); in other words, increasing average CEO ownership decreases average firm value. This decrease is the result of the wealth effect in the model. The CEOs of the largest firms in this subsample hold such a large amount of wealth in their firm ownership that the wealth effect overtakes the substitution effect; increasing their wealth leads them to want more leisure. As with the low institutional ownership firms, low blockholder ownership firms appear to give their CEOs more ownership than is necessary.

The results from splitting the sample along governance lines indicate that well-governed firms do a good job of incentivizing their CEOs, which seems to lead them to have practically no hidden costs to increasing ownership. This implies that the hidden costs in the full sample are at least partly the result of agency problems at the board level. Poorly-governed firms give their CEOs more ownership than they should, on average. Even ignoring hidden costs, the explicit marginal cost of ownership appears to be larger than even the most aggressive estimate of its marginal benefit. For low blockholder ownership firms, there is no marginal benefit because of the level of current ownership.

V.D Time period subsamples

Hall and Liebman (1998) show that the average amount of ownership held by CEOs increased greatly throughout the 1980s and early 1990s. In this study, mean CEO ownership is 3.8% of outstanding shares ($76 million) in the early sample (1993-1996), 4.0% ($158 million) during the tech boom (1997-2001), and 3.1% ($108 million) in the late sample (2002-2007). CEOs early in the sample hold a larger percentage of their firms on average than late in the sample, which is due to the early sample containing smaller firms.

Panel E of Table 5 contains the estimates of the model's parameters for these three subsamples. Early period firms have less volatile valuations, less risk-averse CEOs, and lower cost of effort for their CEOs. Firms during and after the tech boom are quite similar. For the last two sample periods, the only two parameters that are statistically different from each other are the volatility of market-to-book and board preference.

Board preference becomes more negative throughout the sample period. On one hand, increasing λ during the sample period may imply that boards are becoming less efficient in giving stock ownership to CEOs – that granting ownership is becoming more costly. On the other hand, any systematic changes between time periods that are not included in the model will show up in the estimate of λ. There is a similar pattern in the firm value elasticity counterfactuals. The upper and lower bounds on the elasticities increase, from (0.006, 0.059) in the early sample to (0.011, 0.111) in the late sample. The marginal benefit to giving ownership is higher in the late sample, implying that the marginal costs must also be larger. The increased cost shows up after translating the elasticities into dollar terms. In the early sample, a 1% increase in ownership is worth \$2.562 million, and the bounds on the increase in firm value are \$0.648 million and \$6.570 million. In the late sample, a 1% increase in ownership is worth \$2.835 million, with bounds on the increase in firm value of \$1.932 million and \$19.719 million. That the increase in ownership is so close to the lower bound for the value increase points to there being substantial hidden costs to increasing firm value in the 2002-2007 period.

There are three reasons to be cautious in suggesting that boards face greater costs to giving ownership later in the sample. First, the sample in the earliest sub-period is smaller; it has an average of 930.5 firms per year, while the two later periods have 1,118.6 and 1,231.2 firms per year. The firms collected by Execucomp in the early 1990s may not be similar enough to firms during and after the tech boom to make any comparisons. Second, if CEO outside wealth is larger during the tech boom, then CEOs at that time also require greater incentives. Another way of saying this is that the level of the market dropped by a large amount during the recession of 2001, and did not recover to late 1990s levels for a number of years. Thus CEOs late in the sample period may be poorer than CEOs during the tech boom, and not need as much ownership. I do not have time-varying outside wealth embedded in the model or in this particular estimation, which may bias the estimate of board preference. Third, parameter estimates come from within-firm variation. By cutting the sample into sub-periods, the within estimates necessarily rely on fewer firm-year observations. This shorter by firm time series increases the variance of the sampling distribution of the moments used to estimate the model. In short, subsample moment estimates have a higher probability of not being similar to population moments.

The results from estimating the model for different sub-periods are interesting. The firms in the earliest sub-period appear to be quite different from later firms, either because they were changed by the technology revolution of the 1990s, or they were replaced by new firms during the IPO booms of the mid and late 1990s. There is evidence for an increase in the cost of giving CEOs ownership through time.

VI Robustness Checks

The results of the estimation depend on the model's assumptions. To make the model estimable, some parameters have been set to reasonable values. This section describes how the results are affected by changes in these assumed parameter values. Specifically, different values for the discount factor, maximum tenure length, retirement length, and CEO outside wealth are considered, with the results in Table 6.

VI.A Discount factor

For the main results, the discount factor is assumed to be 0.98, corresponding to a discount rate of 2%. To see the effects of higher and lower discount rates, I re-estimate the model assuming the discount factor is either 0.95 or 0.995. None of the estimated model parameters are greatly affected by using a different discount factor. The largest change is the CEO effectiveness parameter, α, which is significantly higher for the lower discount factor. As all of the other parameters are basically the same, a larger value of α implies CEOs exert less effort, which makes sense as with a lower discount factor a CEO values future wealth less. This effect is not large, however, as with a discount factor of 0.95 average effort is 0.9444, compared to 0.9474 for a discount factor of 0.98.

VI.B Maximum CEO tenure

To see the effect of different assumed maximum CEO tenure lengths, I estimate the model using tenures of 15 and 25 years. As tenure increases, the estimates of risk aversion and effort cost both increase, while the effectiveness parameter decreases. Even though effort is more costly under the assumption of 25 years maximum tenure, CEOs exert more effort under that assumption: 0.9524 compared to 0.9446 for a 15 year maximum tenure. This makes intuitive sense, as a CEO with a longer expected career has less of a horizon problem near the beginning of his career, and so he works harder. None of the parameter estimates are that different from the base estimation.

VI.C Retirement length

For retirement length, I estimate the model using retirement lengths of 15 and 25 years. There is significant movement in risk aversion (decreasing in retirement length) and the cost of effort (increasing). In addition, the estimate of the board's preference for giving the CEO ownership decreases as retirement length increases, which means the estimate of the board's cost from giving ownership decreases as the assumed retirement length increases. Since a CEO will receive longer enjoyment of wealth with a longer retirement period, his effort is greater for the longer assumed retirement lenght (0.9488 vs. 0.9407).

VI.D Outside wealth

There are two different assumptions about wealth to test. First is the level of wealth, which is assumed to be $5 million in the base estimation. Second is the distribution of wealth across CEOs, which is unlikely to be flat in the data. Since it is likely that CEOs of larger firms are also wealthier, I let CEO wealth increase with firm size.

First I estimate the model using outside wealth of $3 million and $7 million. The estimate of risk aversion strongly increases in assumed level of outside wealth, and the cost of effort decreases. With the assumption of $7 million, the board's preference for giving the CEO ownership is significantly higher than the base estimation; to match the data, the model requires boards to see giving ownership as very costly if outside CEO wealth is high. Interestingly, the estimate of CEO effort is lower for both robustness checks than for the base estimation: 0.9413 for the $3 million assumption and 0.9435 for the $7 million assumption.

In the model CEO salary is purely a function of firm size, so using a similar assumption to one use in Baker and Hall (2004), I assume CEO outside wealth is a multiple of CEO salary. To keep the results of estimating the model under an assumption of CEO wealth increasing in firm size comparable to the main estimation, I use multiples that produce average and median outside wealth equal to \$5 million. These multiples are, respectively, 3.55 and 4.56. The estimate of risk aversion is smaller under the assumption that wealth increases with firm size (2.1007 and 2.3896 compared to 3.0647), and estimates of the cost of effort are larger (0.4924 and 0.4104 compared to 0.3785). In addition, the estimates of the board's preference for giving ownership are closer to one, meaning if outside wealth increases in firm size as assumed here, then boards do not see granting CEO ownership as costly as the base estimation implies. Average CEO effort is slightly smaller for increasing CEO wealth with a mean of \$5 million (0.9424), but larger for wealth with a median of \$5 million (0.9513).

VII Conclusion

In this paper I develop and estimate a structural model of a CEO's incentive and effort levels. Using the estimates of the underlying structural parameters, I find the sensitivity of firm value to changes in CEO ownership. An increase in CEO ownership of 1% increases average firm value between 0.008% and 0.090%, with a best estimate of 0.076%. The size of this increase implies that there may be some value to increasing average CEO ownership. As CEOs already exert a large amount of effort, a significant increase in ownership is necessary to increase firm value by any significant amount. The evidence that boards do not increase ownership indicates there may be hidden costs to increasing ownership. Hidden costs appear to be much smaller for firms with high institutional ownership and high blockholder ownership, suggesting that these hidden costs are lower in the presence of good governance.

References

Baker, G., Hall, B., 2004. CEO incentives and firm size. Journal of Labor Economics 22, 767-798.

Bergstresser, D., Philippon, T., 2006. CEO incentives and earnings management. Journal of Financial Economics 80, 511-529.

Bliss, R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. Journal of Finance 59, 407-446.

Chetty, R., 2006. A new method of estimating risk aversion. American Economic Review 96, 1821-1834.

Coles, J., Lemmon, M., Meschke, F., Forthcoming. Structural models and endogeneity in corporate finance: The link between managerial ownership and corporate performance. Journal of Financial Economics.

Conyon, M., Core, J., Guay, W., 2011. Are U.S. CEOs paid more than U.K. CEOs? Inferences from risk-adjusted pay. Review of Financial Studies 24, 402-438.

Core, J., Guay, W., 1999. The use of equity grants to manage optimal equity incentive levels. Journal of Accounting & Economics 28, 151-184.

Core, J., Guay, W., 2010. Is CEO pay too high and are incentives too low: A wealth-based contracting framework. Academy of Management Perspectives 24, 5-19.

Core, J., Guay, W., Verrechia, R., 2003. Price versus non-price performance measures in optimal CEO compensation contracts. Accounting Review 78, 957-981.

Dai, Z., Jin, L., Zhang, W., 2010. CEO incentive dynamics and their effect on firm value. Unpublished working paper. Available at SSRN: http://ssrn.com/abstract=1763016.

Demsetz, H., Villalonga, B., 2001. Ownership structure and corporate performance. Journal of Corporate Finance 7, 209-233.

Dittman, I., Maug, E., 2007. Lower salaries and no options? On the optimal structure of executive pay. Journal of Finance 62, 303-343.

Edmans, A., Gabaix, X., Forthcoming. Tractability in incentive contracting. Review of Financial Studies.

Edmans, A., Gabaix, X., Landier, A., 2009. A multiplicative model of optimal CEO incentives in market equilibrium. Review of Financial Studies 22, 4881-4917.

Erickson, T., Whited, T., 2000. Measurement error and the relationship between investment and q. Journal of Political Economy 108, 1027-1057.

Fernandes

Gabaix, X., Landier, A., 2008. Why has CEO pay increased so much? Quarterly Journal of Economics 123, 49-100.

Gayle, G., Miller, R., 2009a. Has moral hazard become a more important factor in managerial compensation? American Economic Review 99, 1740-1769.

Gayle, G., Miller, R., 2009b. Identifying and testing models of managerial compensation. Unpublished working paper. Carnegie Mellon University.

Grinstein, Y., Hribar, P., 2004. CEO compensation and incentives: Evidence from M&A bonuses. Journal of Financial Economics 73, 119-143.

Hall, B., Liebman, J., 1998. Are CEOs really paid like bureaucrats? Quarterly Journal of Economics 113, 653-691.

Hall, B., Murphy, K., 2003. The trouble with stock options. Journal of Economic Perspectives 17, 49-70.

Himmelberg, C., Hubbard, R. G., Palia, D., 1999. Understanding the determinants of managerial ownership and the link between ownership and performance. Journal of Financial Economics 53, 353-384.

Holmstrom, B., Milgrom, P., 1987. Aggregation and linearity in the provision of intertemporal incentives. Econometrica 55, 303-328.

Jensen, M., Murphy, K., 1990. Performance pay and top management incentives. Journal of Political Economy 98, 225-265.

Lucas, R., 1976. Econometric policy evaluation: A critique. In: Brunner, K. and Meltzer, A. (Eds.), The Phillips Curve and Labor Markets, Carnegie-Rochester Conference Series on Public Policy, Volume 1. pp. 19-46.

Margiotta, M., Miller, R., 2000. Managerial compensation and the cost of moral hazard. International Economic Review 41, 669-719.

McConnell, J. J., Servaes, H., 1990. Additional evidence on equity ownership and corporate value. Journal of Financial Economics 27, 595-612.

Morck, R., Shleifer, A., Vishny, R. W., 1988. Management ownership and market valuation. Journal of Financial Economics 20, 293-315.

Murphy, K., 1999. Executive compensation, in Handbook for Labor Economics, ed. O. Ashenfelter and D. Card. Elsevier, Amsterdam.

Palia, D., 2001. The endogeneity of managerial compensation in firm valuation: A solution. Review of Financial Studies 14, 735-764.

Taylor, L., 2010a. Why are CEOs rarely fired? Evidence from structural estimation. Journal of Finance 65, 2051-2087.

Taylor, L., 2010b. CEO wage dynamics: Evidence from a learning model. Unpublished working paper. University of Pennsylvania.

Appendix

The Appendix describes the state space used in solving the model, along with the weighting matrix used for estimation and inference and the computational method.

I solve the model numerically. Therefore, I can only use a discrete state space to estimate the model, despite most of the variables being continuous. I solve the model using 18 points of support for firm size A, 24 points of support for market-to-book M, and 64 points of support for both equity e and effort a. CEO tenures in themodel last no longer than 20 years ($T = 20$), and a CEO who does not quit before this point expects to live 20 years after retirement (N). I solve the model using backward induction, then perform 10 simulations of 33,000 firms over 200 years, keeping firm value and CEO ownership data for the last 16 years of the simulation since that is the length of the sample (1992-2007). I then compute moments for each simulation using these 594,000 firm-year observations, and average across simulations to get the simulated moments for a given set of parameters θ.

The solution parameter vector $\widehat{\theta}$ to SMM estimation solves:

$$\widehat{\theta} = \arg\min_{\theta} \left[\widehat{J}_N - \frac{1}{S} \sum_{s=1}^{S} \widehat{j}_n^s (\theta) \right]' \widehat{\Omega}_N \left[\widehat{J}_N - \frac{1}{S} \sum_{s=1}^{S} \widehat{j}_n^s (\theta) \right], \tag{14}$$

where $\widehat{J}_N$ are the moments calculated from the data, $\widehat{j}_n^s (\theta)$ are the same moments calculated from simulation s of the model given parameters θ, and $\widehat{\Omega}_N$ is an arbitrary positive definite matrix that converges in probability to a deterministic positive definite matrix Ω. There are two obvious choices for this weight matrix $\widehat{\Omega}_N$: an identity matrix and the optimal weighting matrix. The moments are estimated with differing levels of precision, and their point estimates vary in magnitude. Since an identity matrix would place equal weight on each moment, those with the largest absolute value would implicitly receive more consideration by the optimization, since they would have the biggest impact on the objective function. Instead I use the inverse of the variance-covariance matrix of the moments, adjusted for clustering at the firm level. This is the optimal weighting matrix, which I find using the Erickson and Whited (2000) influence-function approach.

The state space extends in four dimensions for CEOs (five for the board of directors, since CEO effort is a state variable for it), which complicates estimation because of the curse of dimensionality – the CEO (board) state space has 552,960 (35,389,440) points of support. To get around this problem I use GPU computation, which allows parallelization in solving the model. For each tenure state, I can solve the model for the remainder of the state space at the same time, as none of the states rely on each other once tenure is considered. This allows me to solve the model efficiently enough for estimation.

Table 1: Descriptive Statistics

This table displays the descriptive statistics of the variables used in estimating the model. The sample is based on a merge of Compustat and Execucomp for the years 1993 to 2007 at an annual frequency. All dollar amounts are in year 2000 dollars. Assets is the total asset value of the firm. Market-to-Book is the market value of equity divided by firm assets. CEO ownership is the value of a CEO's straight equity and the equity equivalent value of stock options.

	Mean	S. D.	25%	50%	75%	N
Assets ($000,000s)	11,422.20	56,475.00	484.55	1,442.28	5,194.95	17,244
Market Value ($000,000s)	6,595.26	21,582.00	515.39	1,352.82	4,341.59	17,244
Market-to-Book	1.424	1.717	0.503	0.957	1.721	17,244
CEO Ownership (%)	3.575%	5.627%	0.664%	1.574%	3.616%	17,244
CEO Ownership ($000,000s)	115.74	1,036.00	9.156	23.43	62.03	17,244

Table 2: Simulated Moments Estimation

Estimates of the moments and parameters of the model based on a sample of firms from a merger of Compustat and Execucomp for the years 1993 to 2007. Estimation is done by SMM, which matches moments computed using a simulated sample of firms with the same moments calculated from the data. Panel A shows the estimates for the moments both from the data and the simulation, along with t-statistics for the difference between the two. Mean ownership, $E(e)$, is the average percentage ownership held by the CEO in the sample. Variance of ownership, $V(e)$, is the within-firm variance of CEO ownership. Each of the covariance moments is computed using a within estimator as well. For these covariances, $\ln e$ is the natural logarithm of current CEO ownership, $\ln e_{-1}$ is the natural logarithm of previous period's CEO ownership, $\Delta^{\ln e}$ is the change in the natural logarithm of CEO ownership, $\ln M$ is the natural logarithm of market to book, $\Delta^{\ln M}$ is the change in the natural logarithm of market to book, τ is current CEO tenure, early refers to CEO tenures less than 10 years, and late refers to CEO tenures above 10 years. Average natural logarithm of market to book, $E(\ln M)$, is the full sample mean of the ratio of firm market capitalization over assets. Serial correlation, ψ_1, and residual variance, ν^2, come from the regression:

$$\ln M_{t+1} = \psi_0 + \psi_1 \ln M_t + \zeta_{t+1},$$

where ν^2 is the variance of ζ_{t+1}. Panel B contains the estimates and standard errors (in parentheses) of the model's parameters: CEO risk aversion γ, the cost of CEO effort η, the measure of board efficiency λ, the maximum return to CEO effort α, the autocorrelation of firm value ρ, and the standard deviation of shocks to firm value σ. Panel B also includes a test of the overidentification restrictions of the model, χ^2 along with its p-value.

Panel A. Moments

	Actual Moments	Simulated Moments	T-statistics
$E(e)$	0.0357	0.0358	0.1541
$V(e)$ (x10)	0.0059	0.0042	-9.3900
Early cov$(e, \ln M)$	0.0211	0.0171	-5.1674
Late cov$(e, \ln M)$	0.0083	-0.0016	-2.4744
cov$(\Delta^{\ln M}, \tau)$	-0.0307	-0.0333	-0.2120
cov$(\Delta^{\ln M}, e_{-1})$	0.0032	0.0019	-1.9025
Early cov$(\Delta^{\ln M}, \Delta^{\ln e})$	0.0095	0.0096	0.4074
Late cov$(\Delta^{\ln M}, \Delta^{\ln e})$	0.0012	0.0012	-0.0593
$E(\ln M)$	-0.1034	-0.1974	-5.0357
ψ_1	0.9062	0.9127	1.5628
υ^2	0.1777	0.1713	-1.6965

Panel B. Parameters

γ	η	λ	α	ρ	σ	χ^2
3.0647	0.3785	1.5375	0.0342	0.8980	0.3933	36.2686
(0.0233)	(0.0033)	(0.0082)	(0.0055)	(0.0031)	(0.0070)	($<$0.0001)

Table 3: Actual versus Simulated Data

This table compares the actual mean and median values of CEO ownership (in percentage and dollar terms) and firm value (in market-to-book and market capitalization terms) for the sample and for a simulation of the estimated model.

Panel A. CEO Ownership

	Ownership (%)		Ownership ($ millions)	
	Mean	Median	Mean	Median
Actual	3.575	1.574	115.740	23.428
Simulation	3.576	2.779	177.414	21.792

Panel B. Firm Value

	Market-to-Book		Market Capitalization ($ millions)	
	Mean	Median	Mean	Median
Actual	1.424	0.957	6,959.262	1,352.823
Simulation	1.363	0.837	6,987.756	943.924

Table 4: Effort and Value sensitivities

This table reports the results of the counterfactual experiment in which CEO ownership is increased/decreased by 1% in all states of the world. The effect on CEO effort and firm value are given for three different time frames. The initial change is the change in effort in the year following the change in ownership policy, and the effect the change in effort directly has on firm value. The permanent change is the ultimate resulting increase in average effort and firm value that comes from simulating the model forward after the policy change until it reaches a new steady state. The present value is the discounted value of the increase/decrease in firm value for all the periods between the period the change occurs and the period in which the simulation reaches a steady state.

	Initial Change		Permanent Change		Present Value	
	Decrease	Increase	Decrease	Increase	Decrease	Increase
CEO Effort	-0.454	0.128	-0.509	0.151	na	na
Firm Value (%)	-0.014	0.008	-0.160	0.090	-0.139	0.076
Firm Value ($)	-0.560	0.964	-11.171	6.274	-9.606	5.267

Table 5: Parameter Estimation in Subsamples

Estimates of the model parameters and standard errors (in parentheses) based on subsamples of firms from a merger of Compustat and Execucomp for the years 1993 to 2007. Estimation is done by SMM, which matches moments computed using a simulated sample of firms with the same moments calculated from the data. The parameters are CEO risk aversion γ, the cost of CEO effort η, the measure of board efficiency λ, the maximum return to CEO effort α, the autocorrelation of firm value ρ, and the standard deviation of shocks to firm value σ. This table also includes the estimate of sensitivity to changes in CEO ownership from a counterfactual increase in CEO ownership by 1%. Panel A compares firms separated into large and small sizes, where small and large are the top and bottom 33% of the asset distribution. Panel B compares high and low options using firms based on whether the average percentage of their CEOs' ownership is above or below 50%. Panel C shows the parameter estimates for high and low institutional investor ownership firms, and panel D shows the estimates for high and low blockholder ownership firms. For these two subsamples high and low are for the top and bottom 33% of their respective distributions. Panel E shows the estimates for different sub-periods of the sample.

γ	η	λ	α	ρ	σ	Sensitivity
Panel A. Firm Size						
Large Firms						
1.8469	0.4239	1.7637	0.0149	0.9462	0.3387	0.0455
(0.0295)	(0.0058)	(0.1501)	(0.0081)	(0.0095)	(0.0163)	
Small Firms						
3.2506	0.3498	1.1689	0.0951	0.8126	0.4761	0.0708
(0.0042)	(0.0096)	(0.0339)	(0.0054)	(0.0066)	(0.0049)	
Panel B. Use of Stock Options						
High Stock Options Use						
3.3357	0.2061	1.2659	0.0120	0.9050	0.4182	0.0209
(0.0120)	(0.0014)	(0.0163)	(0.0121)	(0.0063)	(0.0070)	
Low Stock Options Use						
2.1832	0.3646	0.8795	0.0202	0.9018	0.4028	0.0591
(0.0325)	(0.0076)	(0.0058)	(0.0101)	(0.0036)	(0.0019)	
Panel C. Institutional Ownership						
High Institutional Ownership						
2.0010	0.1746	1.3343	0.0528	0.8560	0.4019	0.0240
(0.0249)	(0.0032)	(0.0415)	(0.0044)	(0.0040)	(0.0062)	
Low Institutional Ownership						
2.6023	0.2399	0.6674	0.0459	0.8651	0.4769	0.0324
(0.0113)	(0.0055)	(0.0032)	(0.0076)	(0.0075)	(0.0077)	
Panel D. Blockholder Ownership						
High Blockholder Ownership						
3.4794	0.1874	0.6978	0.0164	0.8970	0.4528	0.0139
(0.0514)	(0.0042)	(0.0761)	(0.0128)	(0.0082)	(0.0058)	
Low Blockholder Ownership						
2.4359	0.2322	0.7850	0.0349	0.9570	0.3755	-0.0661
(0.0093)	(0.0047)	(0.0316)	(0.0127)	(0.0073)	(0.0056)	
Panel E. Time Period						
Early: 1993 - 1996						
1.9886	0.3370	1.5077	0.0300	0.9044	0.3225	0.0516
(0.0123)	(0.0046)	(0.0110)	(0.0056)	(0.0019)	(0.0032)	
Tech Boom: 1997 - 2001						
2.2890	0.4293	1.6359	0.0554	0.8877	0.4728	0.0833
(0.0390)	(0.0063)	(0.0529)	(0.0086)	(0.0222)	(0.0063)	
Late: 2002 - 2007						
2.2935	0.4288	2.0366	0.0471	0.8931	0.3731	0.0974
(0.0079)	(0.0034)	(0.0269)	(0.0049)	(0.0024)	(0.0028)	

Table 6: Robustness Checks

Estimates of the model parameters and standard errors (in parentheses) using different assumptions for the model. The parameters are CEO risk aversion γ, the cost of CEO effort η, the measure of board efficiency λ, the maximum return to CEO effort α, the autocorrelation of firm value ρ, and the standard deviation of shocks to firm value σ. The test statistic and p-value for the test of the overidentification restrictions of the model, χ^2 are in the final column. Panel A reports the estimates for the model if for two different values of the discount factor. Panel B reports the results from using either 15 or 25 years as the maximum CEO tenure. Panel C reports the estimates using either 15 or 25 years for retirement. Panel D reports parameter estimates for different values of outside wealth. Panel E reports estimates using the assumption that CEO outside wealth increases with firm size. Specifically for this last panel, 3.55x refers to outside wealth being 3.55 times the CEO's salary, which gives a mean value of outside wealth equal to \$5 million, and 4.56x refers to outside wealth being 4.56 times CEO salary, for a median value equal to \$5 million.

γ	η	λ	α	ρ	σ	χ^2
Panel A. Discount Factor						
$\beta = 0.995$						
2.9423	0.3757	1.5190	0.0359	0.9015	0.3971	35.1014
(0.0090)	(0.0075)	(0.0106)	(0.0027)	(0.0028)	(0.0022)	(<0.0001)
$\beta = 0.95$						
2.9517	0.3638	1.4849	0.0454	0.8983	0.3985	65.6320
(0.0310)	(0.0040)	(0.0546)	(0.0051)	(0.0046)	(0.0142)	(<0.0001)
Panel B. CEO Tenure						
15 Years						
2.9274	0.3634	1.5609	0.0413	0.9046	0.3978	62.4285
(0.0250)	(0.0070)	(0.0247)	(0.0020)	(0.0062)	(0.0052)	(<0.0001)
25 years						
3.0765	0.3839	1.4461	0.0332	0.8975	0.3949	33.2293
(0.0349)	(0.0061)	(0.0103)	(0.0032)	(0.0033)	(0.0128)	(<0.0001)
Panel C. Retirement Length						
15 Years						
3.2466	0.3477	1.8232	0.0468	0.8977	0.3994	36.9249
(0.0297)	(0.0054)	(0.0174)	(0.0050)	(0.0028)	(0.0018)	(<0.0001)
25 Years						
2.7668	0.4242	1.3466	0.0370	0.9027	0.3944	73.0343
(0.0201)	(0.0094)	(0.0354)	(0.0035)	(0.0104)	(0.0234)	(<0.0001)
Panel D. Outside Wealth Level						
$W =$ \$3 million						
2.1770	0.4737	1.4232	0.0449	0.8971	0.4009	61.4158
(0.0110)	(0.0071)	(0.0249)	(0.0136)	(0.0056)	(0.0073)	(<0.0001)
$W =$ \$7 million						
3.4566	0.3452	1.9164	0.0469	0.8975	0.3948	42.3320
(0.1207)	(0.0077)	(0.0173)	(0.0041)	(0.0047)	(0.0048)	(<0.0001)
Panel E. Outside Wealth Increasing with Firm Value						
3.55x Salary						
2.1007	0.4924	1.2277	0.0418	0.9071	0.3940	108.0956
(0.0240)	(0.0051)	(0.0070)	(0.0050)	(0.0057)	(0.0069)	(<0.0001)
4.56x Salary						
2.3896	0.4104	1.2215	0.0380	0.9050	0.3934	111.4713
(0.0141)	(0.0127)	(0.0139)	(0.0118)	(0.0092)	(0.0070)	(<0.0001)

Figure 1: Effort and Ownership in the Model

Figure 1 shows how the average natural logarithm of CEO effort and average CEO ownership covary with CEO tenure, firm valuation, and firm book size for the base parameter values given in the paper.

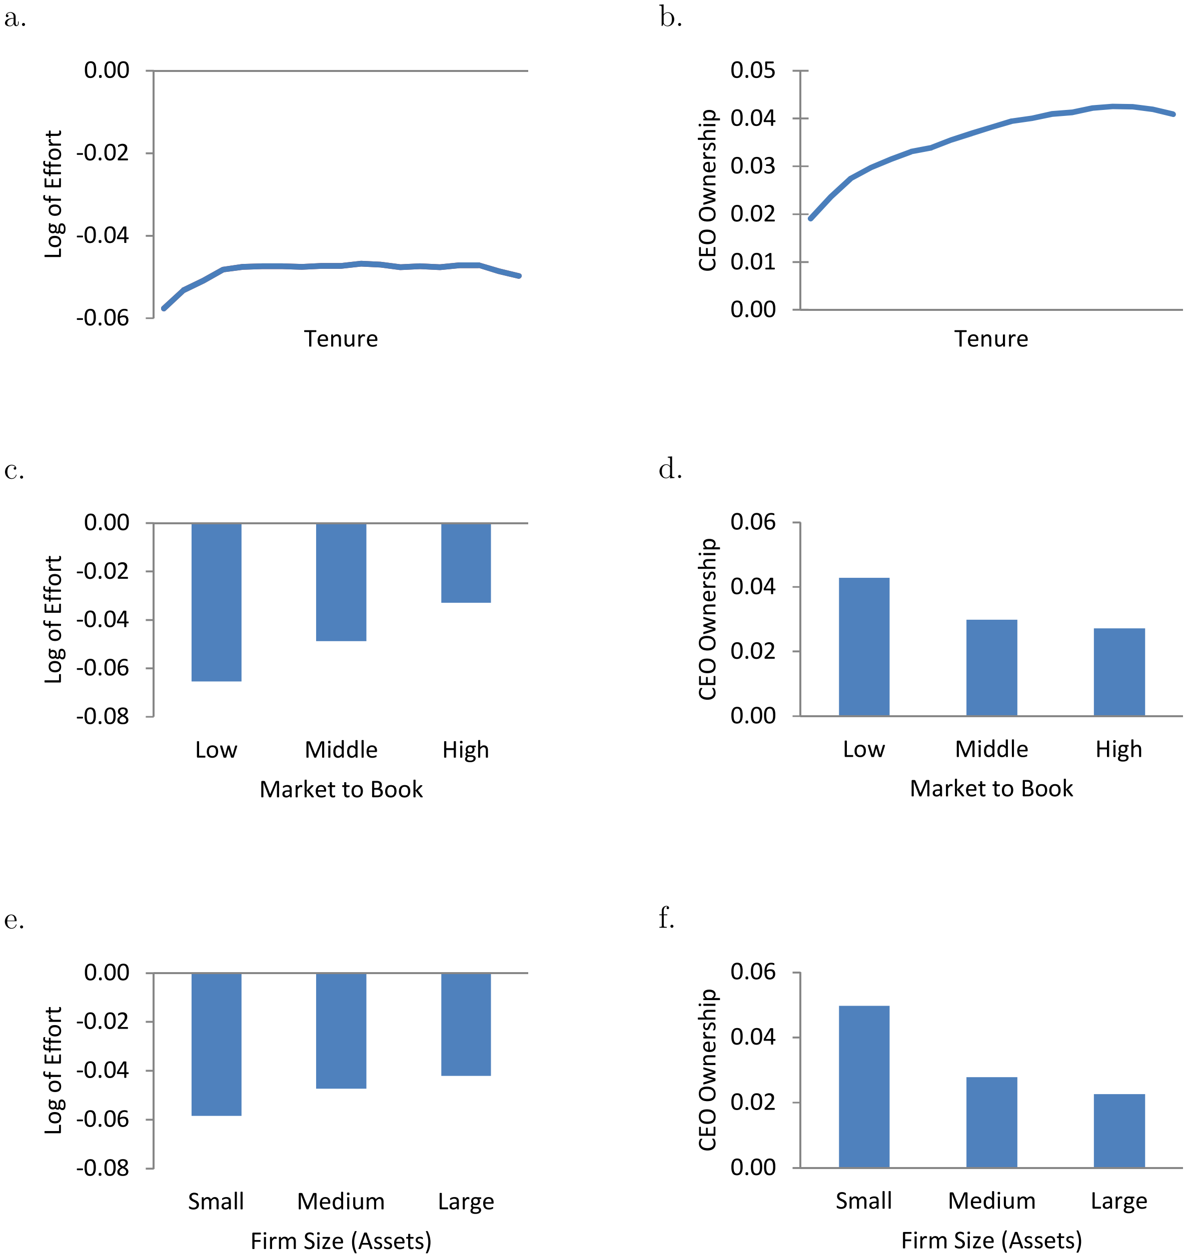

Figure 2: Comparative Statics

Figure 2 depicts the relation between the maximum return to CEO effort α, CEO risk aversion γ, the cost of CEO effort η, and board efficiency λ and i) the mean natural logarithm of CEO effort, and ii) mean CEO percentage ownership.

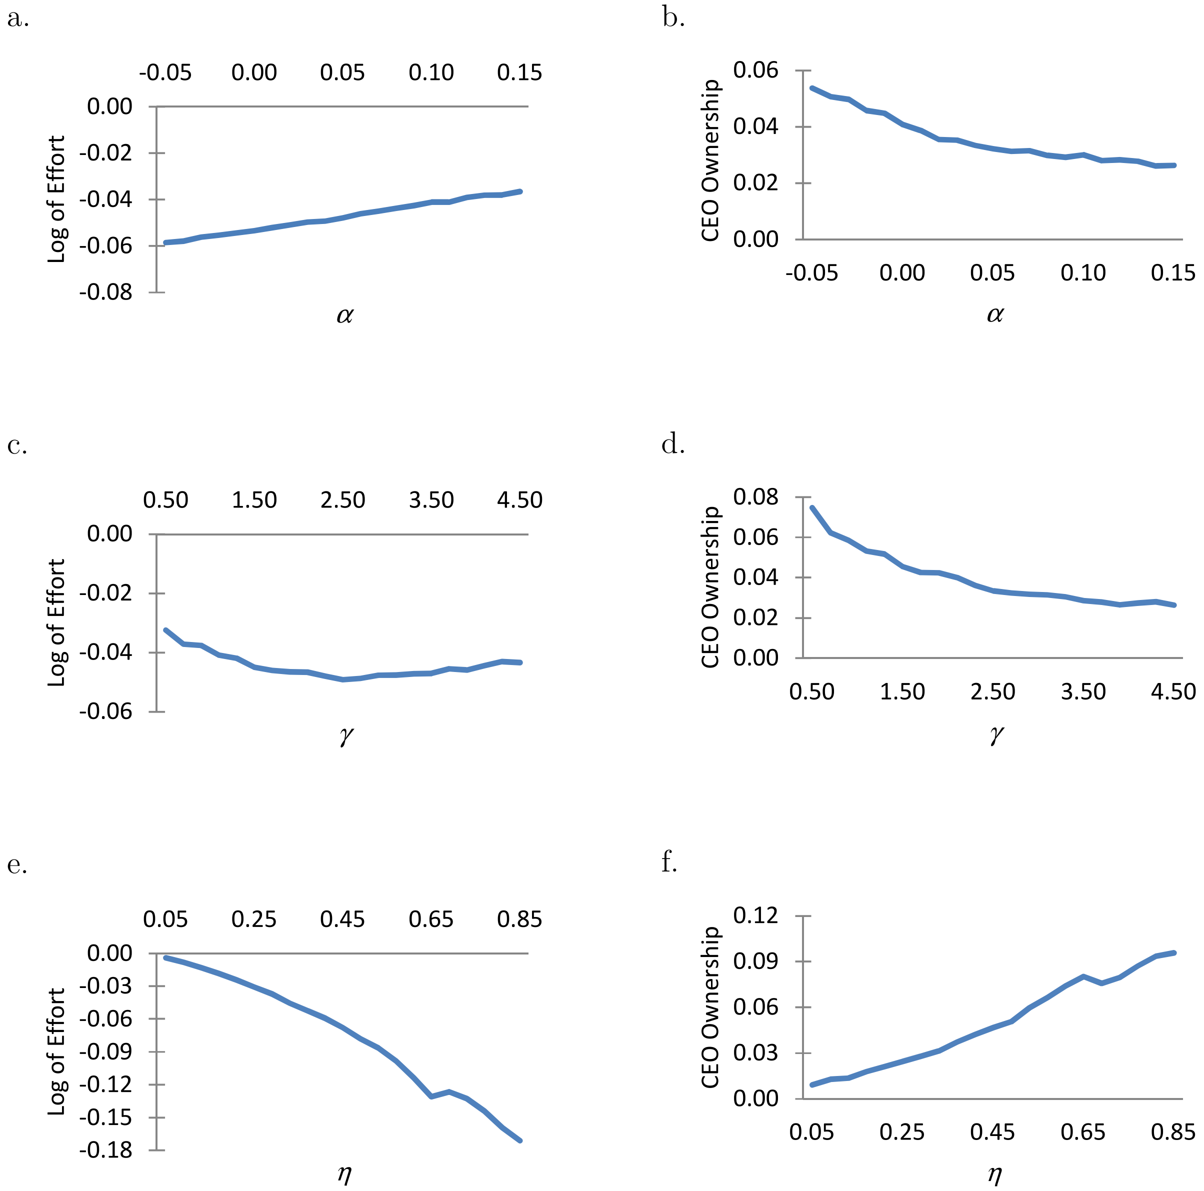

Figure 2 (continued)

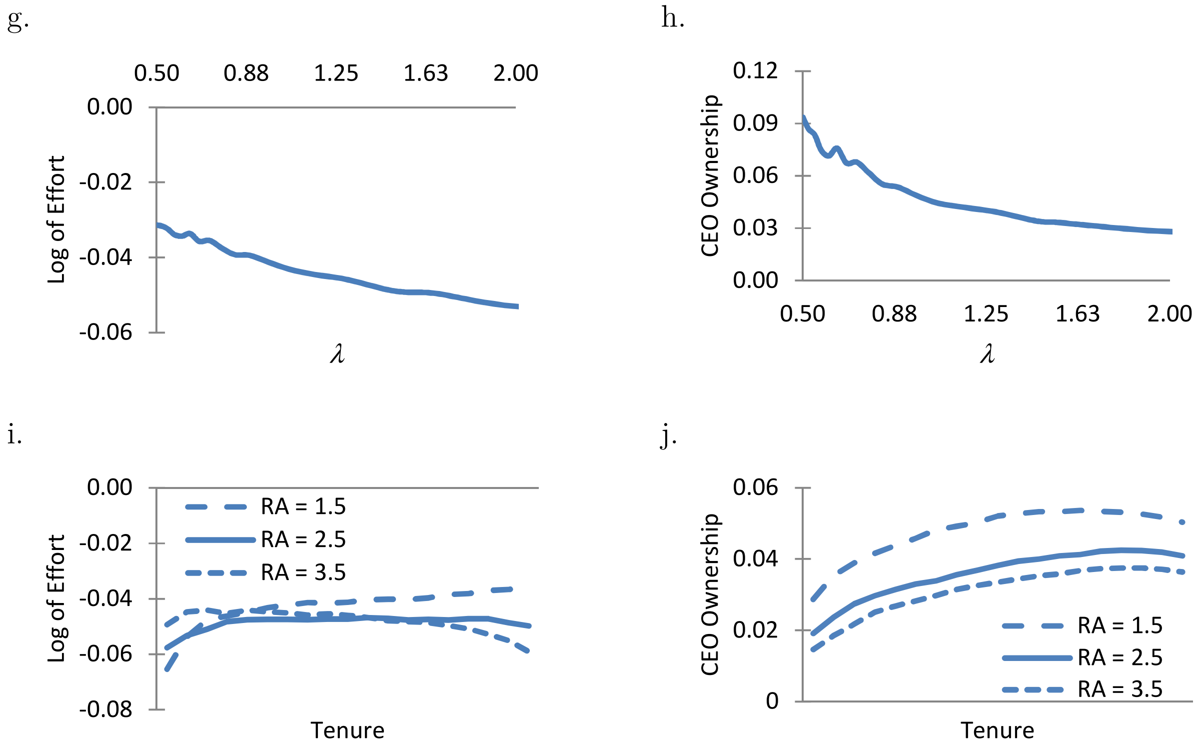

Figure 3: Identification Strategy

Figure 3 plots how the moments I use in estimating the model vary with changes in CEO risk aversion γ, the cost of CEO effort η, and board efficiency λ. Panels a, b, and c show how the covariance of changes in the natural logarithm of market to book and CEO tenure, and the covariance of the log of CEO ownership and the log of market to book for early and late career CEOs identifies the CEO risk aversion parameter. Panels d and e show how the within-firm variance of CEO ownership and the covariance of changes in the log of market to book ratio and the log of previous period's CEO ownership can be used to separately identify the cost of CEO effort and board efficiency.

a.

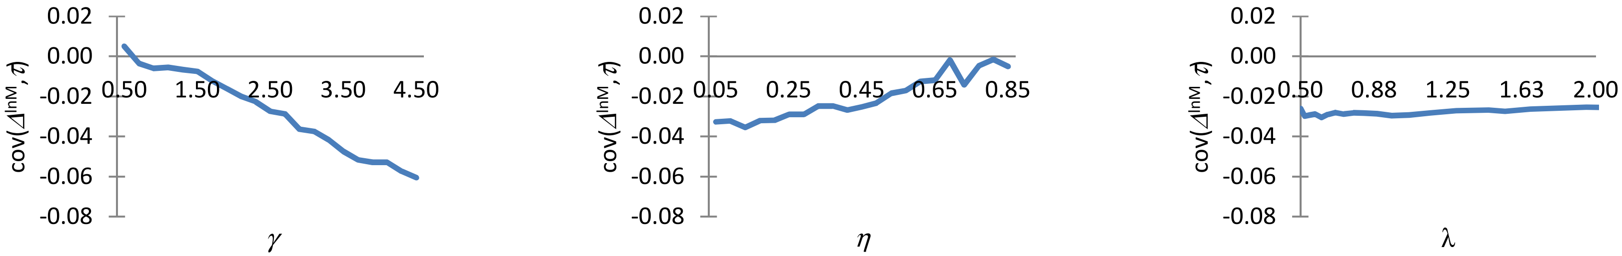

b.

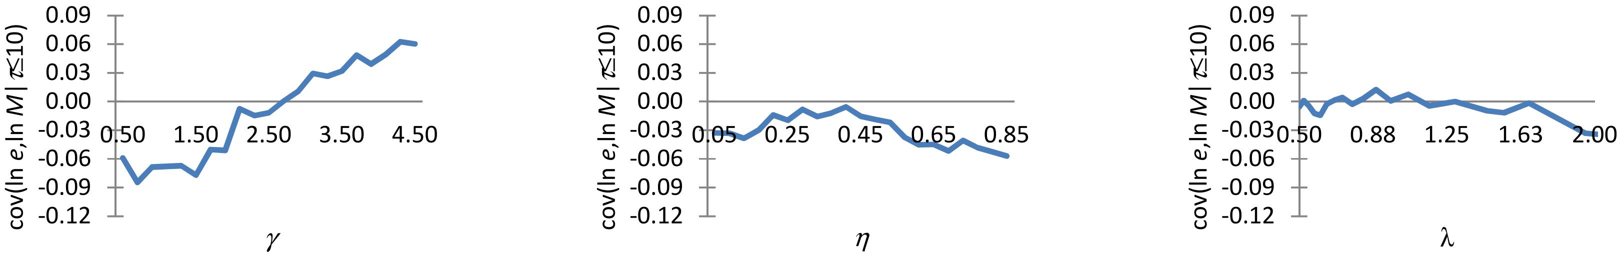

c.

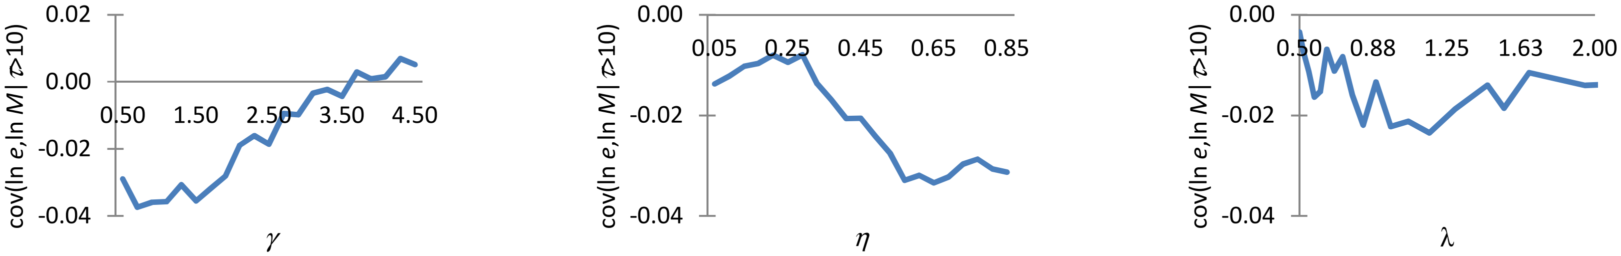
cov(ln e,ln M | t>10)
0.02
0.00
-0.02
-0.04
0.50
1.50
2.50
3.50
4.50
γ
0.05
0.25
0.45
0.65
0.85
η
0.50
0.88
1.25
1.63
2.00
λ

d.

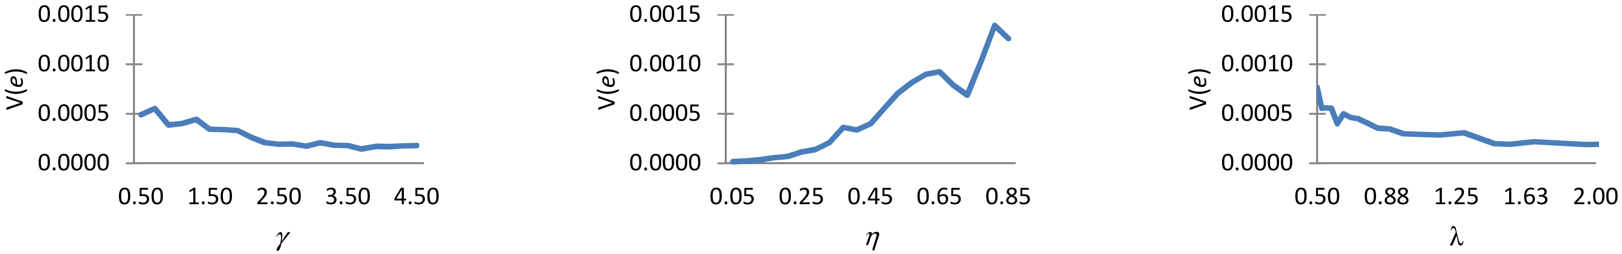
V(e)
0.0015
0.0010
0.0005
0.0000
0.50
1.50
2.50
3.50
4.50
γ
0.05
0.25
0.45
0.65
0.85
η
0.50
0.88
1.25
1.63
2.00
λ

e.

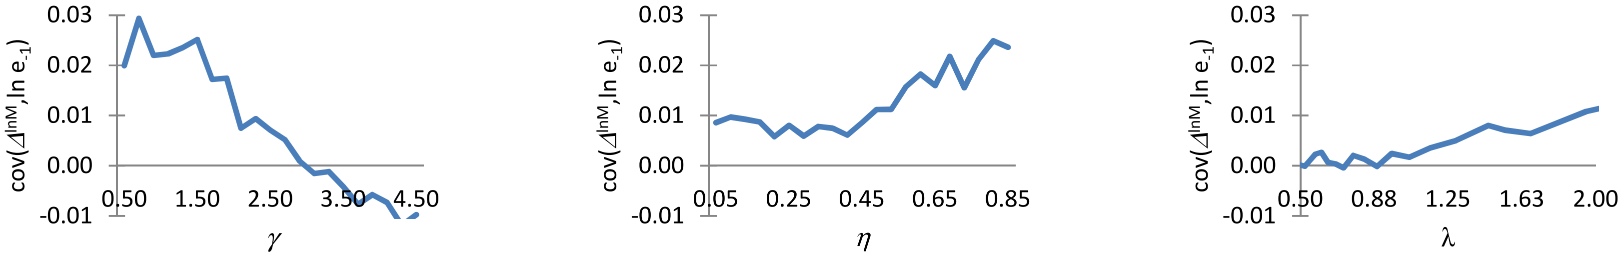
cov(Δ^lnM,ln e_-1)
0.03
0.02
0.01
0.00
-0.01
0.50
1.50
2.50
3.50
4.50
γ
0.05
0.25
0.45
0.65
0.85
η
0.50
0.88
1.25
1.63
2.00
λ

Figure 4: Effect of Ownership on Firm Value

Figure 4 shows the effect of changing CEO ownership on i. CEO effort and ii. firm value. The range for both figures runs from 0.5 to 2.0, where 0.5 (2.0) is setting CEO ownership at half (twice) the level at which it is currently set.

a.

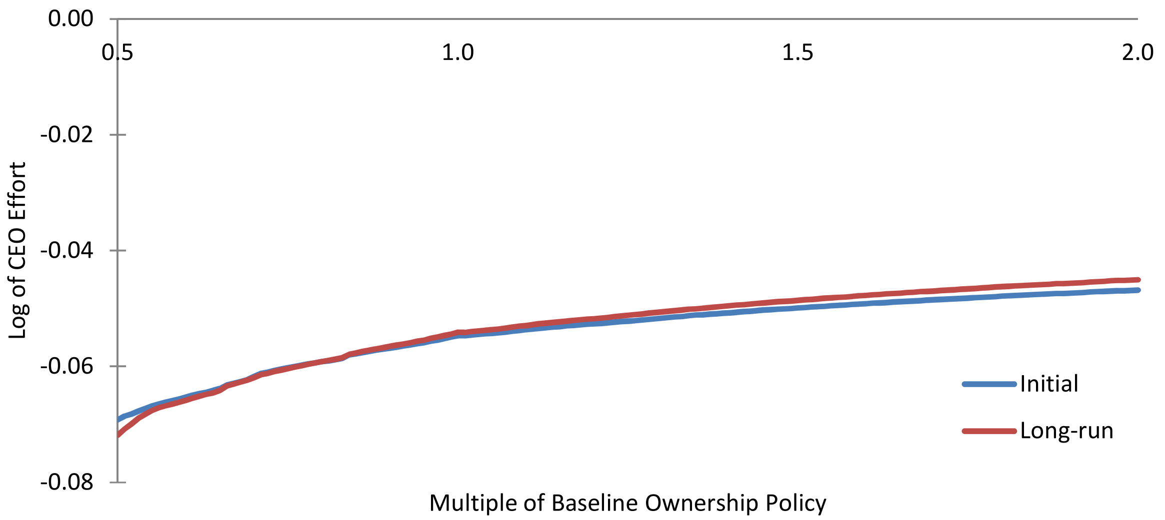

b.