Agency Conflicts and Cash: Estimates from a Structural Model

Boris Nikolov
University of Rochester

Toni M. Whited*
University of Rochester

First Draft: November 18, 2009
Revised, November 25, 2011

*Corresponding author: Toni Whited, William E. Simon Graduate School of Business Administration, University of Rochester, Rochester, NY 14627. (585)275-3916. toni.whited@simon.rochester.edu. We are grateful for helpful comments from Hui Chen, Laurent Frésard, Arthur Korteweg, Laura Liu, Erwan Morellec, Beau Page, Yuri Tserlukevich, and seminar participants at Lingnan University, City University of Hong Kong, Chinese University of Hong Kong, HKUST, Harvard Business School, Boston University, University of Oklahoma, DePaul University, University of Oregon, Oxford, University of New South Wales, McGill University, University of Texas, Drexel University, Carnegie Mellon University, and University of Washington.

Electronic copy available at: http://ssrn.com/abstract=1508556

Agency Conflicts and Cash: Estimates from a Structural Model

Abstract

We estimate a dynamic model of firm investment and cash accumulation to ascertain whether agency problems affect corporate cash policy. We model three specific mechanisms that misalign managerial and shareholder incentives: limited managerial ownership of the firm, compensation based on firm size, and managerial perquisite consumption. Our estimates indicate that agency issues related to perquisites are more important for explaining corporate cash balances but that agency issues related to firm size are more important for firm value. We find that firms with lower blockholder and institutional ownership have higher managerial perquisite consumption. We also find that lower managerial ownership is a key factor in the secular upward trend in cash holding.

Electronic copy available at: http://ssrn.com/abstract=1508556

Do manager-shareholder conflicts distort corporate cash holding decisions, and do these distortions in turn affect shareholder value? This question is economically interesting in light of the buildup of high levels of corporate cash that both preceded and followed the recent financial crisis and recession. This question is also challenging because cash accumulation is a dynamic decision: managers' accumulation of liquid assets today only makes sense if the managers intend to spend them in the future. These dynamics complicate the standard intuition that managers tend to waste large cash flows, because cash flows are not cash stocks. One must therefore also study the incentives to accumulate cash stocks in the first place. Further complicating this question is the issue that different types of agency problems might have different effects on cash accumulation.

To address these questions, we develop and estimate a model of firm investment and cash accumulation. We start with a standard dynamic model featuring a manager who makes decisions about both investment and financing this investment, via current profits, externally raised funds, or accumulated cash balances. Interest taxation implies that cash is costly to hold. However, cash also has value because outside financing is more expensive and because firms have unanticipated funding needs associated with investment opportunities or profit shortfalls. This trade-off generates a precautionary motive for holding cash that naturally depends strongly on investment dynamics.

Layered on top of this basic trade-off is the manager's compensation package, which consists of an equity share, as well as a cash bonus that is related to current profits, and thus to firm size. The equity share aligns the manager's interests with those of shareholders, but the profit-sharing moves him away from the objective of maximizing stockholder value by causing him to view capital as more productive than it actually is. Finally, the manager has a limited ability to divert both profits and liquid assets from their optimal uses within the firm. This managerial resource diversion (perquisite consumption, self-dealing, transfer

Electronic copy available at: http://ssrn.com/abstract=1508556

pricing, or outright stealing) has a strong positive effect on cash accumulation, relative to capital, because in the model the manager can only divert liquid assets and the cash flows produced by capital, but not capital itself.

Our estimation results show that cash holdings depend on both the private benefits managers obtain from resource diversion and on the dependence of managers' compensation on firm size. However, we find that the latter has a larger effect on shareholder value because it distorts not only cash policy, but also real investment policy. We find that the elasticity of firm value with respect to resource diversion is approximately -0.3, whereas the elasticity of firm value with respect to the compensation scheme is near -1.

Our estimation produces several further results. First, standard neoclassical models of the type in Riddick and Whited (2009) do not fit the data well because they only capture the precautionary motive for holding cash. Second, modeling agency issues goes a long way toward reconciling the model with the data. In particular, a model without agency produces average ratios of cash to assets that are only half as large as the 0.13 ratio in the data. Adding agency issues allows the model to match this moment almost exactly. Third, we find substantial cross-sectional heterogeneity in the severity of the two agency issues that we model. We examine firms with poor governance as measured by blockholder and institutional ownership, finding that poor governance leads to more perquisite consumption, more cash accumulation, and greater loss of shareholder value. In contrast, we find the opposite when we measure governance by the presence of antitakeover provisions. This result occurs because firms with many antitakeover provisions differ from their well-governed counterparts in many dimensions such as size, industry and age, all of which affect cash holding. Fourth, small firms hold much more cash than large firms, and agency plays a small role in this difference. In sum, looking at the cross sectional heterogeneity in our sample of firms reveals that governance is but one of many important factors that affect cash holding, and most of these other factors, such as the serial correlation of productivity shocks, or the degree of returns to scale, are inherently unobservable. Using a structural model allows us to see not only whether, but also how all of these different forces affect cash.

The role of agency conflicts in shaping firms' incentives to accumulate and use liquid assets has been of interest in corporate finance at least since Jensen (1986), who points out that managers' decisions about the use of internal funds is central to the conflict between managers and shareholders. Although many empirical researchers have studied the effects of agency conflicts on cash holding, this topic remains of interest because no single prominent conclusion has emerged from these exercises. For example, Mikkelson and Partch (2003) find no difference in the governance structures of firms that hold different amounts of cash. Opler, Pinkowitz, Stulz, and Williamson (1999) find that cash accumulation has little impact on investment, and Bates, Kahle, and Stulz (2009) find that governance has played almost no role in the recent buildup of corporate cash. In contrast, Harford (1999) and Harford, Mansi, and Maxwell (2008) find that firms with more antitakeover provisions hold less cash and tend to do value-destroying acquisitions, and Dittmar and Mahrt-Smith (2007) find that poor governance lowers the value of cash significantly.

To try to sort out the conflicting results in the literature, we take an alternate approach to understanding the effect of corporate governance on corporate cash holdings by using structural estimation. Specifically, we estimate the parameters of our model using simulated method of moments (SMM). On an intuitive level, we use observed features of managerial contracts and observed financing and investment choices to obtain estimates of the average managerial preferences for empire building and for perquisite consumption. This strategy identifies the effects of agency on governance by taking a transparent stand on the features of the managerial decision process that are important for cash holding. More generally, structural estimation recognizes that all variables in the data are endogenous and in fact exploits this endogeneity to achieve identification. It does so by imposing certain model assumptions on the data, and these modeling assumptions are usually easily understood because they are grounded in basic economics. Our estimation also allows us to isolate which types of managerial behaviors affect cash holding and identify the specific economic mechanisms that drive our results. Finally, our estimation replaces noisy measures of governance with a model of specific agency conflicts that is cast in terms of relatively easy-to-observe variables.

Topically our paper contributes to the empirical literature on corporate cash holdings by providing evidence on how agency problems affect cash policy. We thereby start to sort out the conflicting evidence, described above, that this literature has produced. Methodologically, however, our paper belongs to the empirical corporate finance literature that performs structural estimation of dynamic models. For example, Hennessy and Whited (2005, 2007) and DeAngelo, DeAngelo, and Whited (2011) also use SMM on discrete time dynamic models, but they look at different questions such as the low-leverage puzzle, market timing, the cost of external finance, and leverage rebalancing speeds. Similar to our work, Morellec, Nikolov, and Schürhoff (2009) also explore an agency issue using structural estimation, but they examine how managerial resource diversion helps us understand the low leverage puzzle, and they use simulated maximum likelihood rather than SMM. Another study of agency issues is Taylor (2010), which uses structural estimation to explore the agency issues surrounding the firing of CEOs. Finally, our paper is related to the theoretical literature on dynamic models of cash holding such as Riddick and Whited (2009), Bolton, Chen, and Wang (2011), and Anderson and Carverhill (2011).

Section 2 outlines the model, Section 3 presents several comparative statics exercises. Section 4 describes our data and presents summary statistics. Section 5 describes the estimation procedure. Section 6 contains our results and counterfactual exercises, and Section 7 concludes.

1 The Model

We present a neoclassical model of investment and cash accumulation. First we describe the real side of the firm; second, we discuss financing; and third, we specify managerial incentives. Finally, we specify the manager's maximization problem and discuss the intuition behind the solution.

We consider an infinitely lived firm in discrete time. At each time period the firm's manager chooses how much to invest in capital goods and how to finance these purchases. In a standard neoclassical model the manager acts to maximize equity value and thus acts completely in the interest of shareholders. We depart from this setting by considering a manager who faces a standard compensation package consisting of an equity stake in the firm, as well as a share of current after-tax cash flows. This latter feature of the model induces a managerial perception that capital is more productive than it actually is. The reason is that he gains more from increasing the size of the firm than do shareholders. We also assume that the manager has leeway to consume a fraction of the firm's cash stock and cash flow as a private benefit. The firm's managers select investment and liquid asset holdings to maximize their utility, which is linear in the equity stake, the bonus, and firm size. Thus, managers are risk neutral.

1.1 Production Technology

The real side of the firm is characterized by a production technology that uses only capital, k. Per period, after tax profits are given by $(1-\tau)\,\varepsilon k^{\theta}$, in which $\tau < 1$ is the corporate tax rate, $0 < \theta < 1$, and ε is a shock observed by managers each period before making any investment or cash holding decisions. The parameter θ captures a combination of market power and decreasing returns to scale. The shock ε follows an $AR(1)$ in logs,

$$\ln(\varepsilon') = \rho \ln(\varepsilon) + v', \tag{1}$$

in which a prime indicates a variable in the next period and no prime indicates a variable in the current period. The innovation v' has a truncated normal distribution with mean 0 and variance σ_v^2. Because the distribution of the shock v' has finite support, the shock ε does as well. We denote the endpoints of this support as $[\underline{\varepsilon}, \overline{\varepsilon}]$. For convenience, we define the Markov transition function associated with (1) as $q(\varepsilon', \varepsilon)$.

Investment, I, is defined as

$$I \equiv k' - (1-\delta)k, \tag{2}$$

in which δ is the capital depreciation rate, $0 < \delta < 1$. The firm purchases and sells capital at a price of 1 and incurs capital stock adjustment costs that are given by

$$\Psi(I, k) = \gamma k \mathbf{1}_{I \neq 0} + \frac{a}{2}\left(\frac{I}{k}\right)^2 k. \tag{3}$$

in which γ and a are positive constants. The functional form of (3) is from Cooper and Haltiwanger (2006), who specifically study capital stock adjustment costs. The first term captures a fixed component: it is independent of the size of investment and the indicator function implies that this cost only kicks in when investment is nonzero. The smooth component is captured by the second term. These aspects of the model are important because the smoothness or lumpiness of optimal investment policy is important for optimal cash accumulation policy.

1.2 Financing

The firm's financing environment is simple. It can fund its investment projects with current profits, cash, or external funds. Because we are interested in cash rather than in the composition of external finance, we do not distinguish between external debt and equity financing.

We have experimented with an extension of the model in which we include debt and equity separately and find that this feature has no qualitative impact on the results. We therefore opt for our simpler specification. We let c denote the stock of cash and f denote the flow of external financing, where we restrict $f \geq 0$, and $c \geq 0$. Cash earns interest at the risk-free rate, r, and interest is taxed. To make the choice set compact, we assume an arbitrary upper bound on liquid assets, $\overline{c}$. This upper bound is imposed without loss of generality because of our taxation assumptions.

External finance is costly. For every dollar of finance raised by the firm, it must pay a fee that is convex in the amount raised: $\phi(f) \equiv \phi f$. This assumption is motivated by the existence of flotation costs for equity and public debt and by origination fees for loans. As a departure from the specification in Hennessy and Whited (2007), we omit a fixed component because real world firms are likely to use lines of credit or other such low-cost sources of finance for their first dollar of outside financing.[1] Estimates of these parameters also capture possible adverse selection costs.

1.3 Compensation and Incentives

We next discuss managerial compensation and incentives. Our goal is not to derive the form of an optimal contract but to approximate contracts that we actually see in reality and that may or may not be optimal. Thus, the manager's compensation contract consists of a profit share and equity share, and we assume that the contract stays fixed over the life of the firm. Although this assumption implies that we may be missing endogenous variation in contracts, the assumption is necessary for estimation of the model because a full-blown dynamic contracting framework is too intractable to estimate directly. The

[1] It is also possible to include a quadratic cost that is motivated by the idea that the first dollar of funds raised, possibly via a revolver loan, is cheaper than subsequent dollars, which might come from bond flotations or equity offerings. Because this model feature changes our results little, we omit it for simplicity.

equity share is denoted as a fraction $\beta \in (0,1)$ of firm value. The larger β, the more managerial and shareholder incentives are aligned. The profit share is a fraction $\alpha \in (0,1)$ of current operating earnings, which are given by εk^{θ}. In this formulation we exclude cash outflows associated with either investment or cash accumulation activities; otherwise the manager would have an incentive to shut down the capital and cash accumulation programs. We also exclude income from interest on cash balances to avoid modeling a mechanical positive association between profit sharing and cash accumulation. The presence of this parameter captures the well-documented positive relation between CEO compensation and firm size. (Gabaix and Landier (2008)). This parameter also induces a managerial perception of enhanced capital productivity. From the manager's point of view, increased capital not only increases fundamental productivity, but also his profit share. Thus, although we refer to α as a profit-sharing parameter, it should be interpreted more broadly as a managerial perception of enhanced capital productivity.

The survey of CEO compensation in Murphy (1999) documents that compensation packages typically consist of a fixed wage, a profit share, straight equity, and options. Our contract deviates from this scheme in three ways. First, we omit the fixed wage component because such a lump sum would not appear in the model first-order conditions and therefore would have little effect on managerial policies. Second, the model treats the manager's profit share as linear, whereas Murphy (1999) documents that the typical annual incentive plan consists of zero payment if performance is below a specified threshold, followed by an incentive zone where the payoff is linear in performance. The profit share is then capped at some upper bound. In this model these sorts of performance thresholds essentially would serve as normalizing constants and therefore have minimal effects on decisions; so we stick with a simple proportional profit share. Third, in our model the manager's equity compensation consists entirely of common stock, whereas a typical real-world CEO typically receives both common stock and options. Therefore, in the empirical implementation of the model, each option's delta is used to compute an effective equity share. In sum, although our stylized contract departs in some respects from real-world contracts, it is nonetheless a good approximation to actual observed compensation packages. Because our aim is to estimate this model, it is important that the contract be representative of those in the data.

Finally, we model an agency problem directly related to cash holding—a managerial desire to divert a fraction, s, of sum of the firm's current cash flow and cash stock as private benefits. The interpretation of this parameter, as well as the profit sharing parameter, α, is complex. They embody a managerial perception of excess capital productivity and a proclivity toward resource diversion. However, these two parameters also embody managerial entrenchment; that is, the manager's ability to implement his decisions that are distorted by these agency issues. For example, if it is costly for shareholders to launch a control challenge against a possibly entrenched manager, then α and s will be, ceteris paribus, higher.

1.4 The Objective Function

Because part of the manager's compensation is equity, before we can specify the manager's objective, we must first write down the cash flows that go to equity holders. We use a standard accounting identity to express distributions to shareholders as

$$\begin{aligned} d\left(k, k', c, c', \varepsilon\right) \equiv & \left(1-\tau\right)\Big(1-\left(\alpha+s\right)\Big)\varepsilon k^{\theta} + \delta k \tau - I - \Psi\left(I, k\right) \\ & - c' + c\Big(\left(1-s\right) + \left(r-s\right)\left(1-\tau\right)\Big) + f\Big(1-\phi\Big). \end{aligned} \tag{4}$$

The first two terms represent after-tax operating profits, where the second term arises from the tax deductibility of depreciation. This formulation assumes that any managerial resource diversion reduces taxable income. These profits are then spent on physical and financial assets, with any deficit covered by outside financing. We restrict $d(k, k', c, c', \varepsilon) \geq 0$. Because outside financing is costly, the firm never raises outside financing to augment distributions.

Given our specification of the manager's contract and preferences, his one period utility function can be written as

$$u(k, c, \varepsilon) = (\alpha + s)\, \varepsilon k^{\theta} + sc(1 + r) + \beta d(k, k', c, c', \varepsilon). \tag{5}$$

The manager chooses (k', c') each period to maximize the present value of his future utility, discounting at the opportunity cost of funds, r. The sources and uses of fund identity (4) then implies that the manager also implicitly chooses distributions, d, and outside financing, f. The Bellman equation for the problem is given by

$$U(k, c, \varepsilon) = \max_{k',\ c'} \left\{ u(k, c, \varepsilon) + \frac{1}{1 + r} \int U(k', c', \varepsilon')\, dq(\varepsilon', \varepsilon) \right\}. \tag{6}$$

Although the model has no analytical solution, the model can satisfy the conditions for Theorem 9.6 in Stokey and Lucas (1989), which guarantee a solution for (6) in the form of a unique function $U(k, c, \varepsilon)$. Loosely, these conditions require that the parameter values ensure concavity in the term in brackets in (6), that the choice variables (k, c) lie in a compact set, and that $r > 0$ so that (6) is a contraction mapping.

The parameter estimates we obtain all imply the necessary concavity, and the upper bound $\overline{c}$ ensures compactness of the choice set for c. The choice set for k is more complex, but thinking carefully about this choice set also helps the understanding of how a managerial perception of excess capital productivity operates in the model. In a simple neoclassical model without taxation or agency conflicts, the upper bound for the capital stock is given by the condition

$$\overline{\varepsilon}\overline{k}^{\theta} - \delta\overline{k} \equiv 0.$$

As in Gomes (2001), this condition specifies the level of capital above which profits are insufficient to replace depreciated capital even in the best state of nature. Therefore, these levels of the capital stock are not economically profitable. This type of logic can be extended to our setting straightforwardly by inspecting (4) and (5) and formulating an analogous condition as:

$$(\alpha+\beta)(1-\tau)\,\overline{\varepsilon}\overline{k}^{\theta}-\delta\beta\overline{k}\equiv 0. \tag{7}$$

Therefore, k lies in the interval $[0,\overline{k}]$ implied by (7). Concavity of εk^{θ} in k and $\lim_{k\to\infty}\partial\left(\varepsilon k^{\theta}\right)/\partial k = 0$ ensure that $\overline{k}$ is well-defined. Inspection of (7) shows that a positive α broadens the interval $[0,\overline{k}]$ relative to a model without any managerial perception of excess capital productivity $(\alpha=0)$.

Our model, with its estimated parameters, also satisfies the concavity conditions in Theorem 9.8 in Stokey and Lucas (1989) that ensure a unique optimal policy function, $\{k',c'\} = p(k,c,\varepsilon)$. The policy function is essentially a rule that states the best choice of k' and c' in the next period for any (k,c,ε) triple in the current period. The numerical solution for the model is described in the Appendix.

Because we are also interested in shareholder value, we need to define the value of the equity of the firm, $V(k,c,\varepsilon)$. Because it is the expected present value of cash flows given by (4), it can be expressed recursively as

$$V(k,c,\varepsilon)=\max_{k',\,c'}{}^{*}\left\{d(k,k',c,c',\varepsilon)+\frac{1}{1+r}\int V(k',c',\varepsilon')\,dq(\varepsilon,\varepsilon)\right\}. \tag{8}$$

The operator $\max^{*}$ is optimal choice of $\{k',c'\}$ given the policy rule that is the solution to the managerial utility maximization problem in (6). This choice is not, in general, the same choice of $\{k',c'\}$ that would be made if the manager were maximizing the expected present value of cash flows. Put differently, managers do not act completely in the interests of shareholders, and, therefore, for any given (k,c,ε), firm equity value is less than it would be in the absence of misaligned incentives.

1.5 Optimal Policies

Our ultimate goal is to understand whether and how agency problems affect corporate cash holdings by estimating this model directly. However, understanding the estimation results and, more importantly, identifying the model parameters require understanding the economics behind the model. To this end, we first analyze this maximization problem by examining the first-order conditions for an optimal interior financial policy. To do so, we first note that because f and $d(k, k', c, c', \varepsilon)$ cannot both be positive, we can rewrite $d(k, k', c, c', \varepsilon)$ as

$$d^*(k, k', c, c', \varepsilon) = d(k, k', c, c', \varepsilon) - f(1 - \phi).$$

Then we can rewrite $u(k, c, \varepsilon)$ as

$$u(k, c, \varepsilon) = (\alpha + s)\,\varepsilon k^{\theta} + sc(1 + r) + \beta d^*(k, k', c, c', \varepsilon)\left(1 + \phi \mathbf{1}_{d^*<0}\right).$$

Using Leibnitz' rule to differentiate (6) with respect to c' gives

$$\beta\left(1 + \phi \mathbf{1}_{d^*<0}\right) = \frac{1}{(1 + r)} \int U_c(k', c', \varepsilon')\, dq(\varepsilon', \varepsilon).$$

In words, the marginal benefit of a unit of cash today is equal to the expected discounted utility of a unit of cash tomorrow. Next we use the envelope condition to eliminate $U_c(k', c', \varepsilon')$ from the problem. Substituting in the envelope condition gives

$$\beta\left(1 + \phi \mathbf{1}_{d^*<0}\right) = \frac{1}{(1 + r)} \int s(1 + r) + \beta\left(1 - s + (r - s)(1 - \tau)\right)\left(1 + \phi \mathbf{1}_{d^{*\prime}<0}\right) dq(\varepsilon', \varepsilon) \tag{9}$$

$$1 + r = \frac{\int s(1 + r) + \beta\left(1 - s + (r - s)(1 - \tau)\right)\left(1 + \phi \mathbf{1}_{d^{*\prime}<0}\right) dq(\varepsilon', \varepsilon)}{\beta\left(1 + \phi \mathbf{1}_{d^*<0}\right)} \tag{10}$$

The right side of (10) represents the manager's marginal rate of substitution between cash tomorrow and cash today. At an optimum, this marginal rate of substitution equals the rate of return on cash. To understand the role of agency in this model, we compare (10) with the optimality condition in the absence of agency. We therefore set $\beta = 1$ and $s = 0$ in (10) to obtain

$$\frac{1+r}{1+r\,(1-\tau)} = \frac{\int \left(1+\phi \mathbf{1}_{d^{*\prime}<0}\right) dq(\varepsilon', \varepsilon)}{\left(1+\phi \mathbf{1}_{d^{*}<0}\right)}. \tag{11}$$

The fraction on the right side of (11) represents the marginal value of an extra dollar tomorrow versus an extra dollar today. Note that the marginal value of cash inside the firm can be greater than 1 because it can substitute for costly external finance.

To illustrate the implications of (11), we note that expression (4) implies that choosing to hold more cash today lowers the likelihood of needing to raise external finance tomorrow. Thus, the main financing trade-off is between the tax disadvantage of holding cash—on left side of (11)—versus the flexibility benefit of holding cash, as represented on the right side of (11). The flexibility benefit enters through the indicator function $\mathbf{1}_{d^{*}<0}$, because this indicator is more likely to be one when the firm incurs outside funding needs tied to investment opportunities.[2] Thus, this optimality condition shows that the value of cash today stems from its ability to lower the likelihood of having to pay for costly external finance in the future.

Comparing (10) to (11) shows that agency distorts this decision. First, as long as $r > s$, and $\beta << \tau$ (which is the case for all of our estimations), then the right side of (11) exceeds the right side (10), so that optimal cash holdings rise. Intuitively, if a manager hoards a unit of cash today, he increases his utility of cash within the firm tomorrow by $\beta\,(1 - s + (r - s)\,(1 - \tau))$. However, he increases the utility of cash that he can tunnel by much more, $s\,(1 + r)$, because his utility of cash within the firm is scaled down by his ownership fraction, β.

[2]Recent research (Lins, Servaes, and Tufano (2010)) has suggested that firms hoard cash primarily to guard against negative, rather than positive shocks. Our model in part incorporates this feature because the profit function already reflects the optimal choice of variable factors of production.

The intuition behind the profit sharing parameter, α, is more subtle. Recall that the manager's compensation package induces him to perceive capital as being more productive than it is. Investment policy is therefore more likely to require outside financing. The actual marginal product of capital is less likely to cover desired investment expenditures relative to a situation in which the manager chooses investment to maximize shareholder value. Thus, the firm accumulates more cash. The effect on cash is stronger than the effect on capital, however, because cash enters the manager's optimization problem linearly through (4), whereas the effect of α on the capital stock is muted because of decreasing returns.

1.6 Numerical Policy Functions

To extend the model intuition, we examine the policy function, $\{k', c'\} = p(k, c, \varepsilon)$. We plot optimal investment, cash balances, and distributions/external financing as a function of ε for various levels of cash, c, and for the steady state capital stock, k^*, at which the mean marginal product of capital equals the user cost. We also plot current cash flow, which we define precisely as $(1-\tau)\,\varepsilon k^{\theta}/k$. Similarly, investment is $(k' - (1-d)k)/k$, cash is $c'/(k+c)$, and net distributions/external financing are $d^*(k, k', c, c', \varepsilon)/(k+c)$ or $f/(k+c)$. For these exercises we parameterize the model according to the set of estimation results in Table 3 that correspond to the model in which we allow for both profit sharing and perquisite consumption. The remaining parameters, the interest rate, r, and the ownership parameter, β, are set equal to their sample averages in our data.

Figure 1 contains these plots. As seen in the first panel, for all cash levels, cash flow naturally rises with the ε shock. These cash flows are, however, distributed differently depending on the initial level of the stock of cash. From the second panel, we see first that investment responds sharply to ε. More importantly, firms with low cash stocks invest less aggressively than those with higher cash stocks, but this difference is only apparent for high productivity shocks. This differential response makes sense because those firms with higher stocks of cash need to rely less on external financing. They can therefore respond to large positive shocks without having to resort to costly external finance. The third panel depicts the policy function for the ratio of cash to assets and contains two interesting features. First, as ε rises, cash stocks fall because a higher ε implies that capital is more productive relative to cash. Second, as ε rises, firms with high initial cash stocks see a greater drop in cash than those with low initial cash stocks. This second pattern implies that firms spend their cash stocks to fund investment. However, they do not spend all of their cash stocks, except for extremely high productivity shocks. Instead, because external finance is more expensive than internal finance, firms retain some precautionary cash balances. The fourth panel depicts net distributions/external financing. For ease of reading the graph, we have indicated external financing as a negative number and distributions as a positive number. For firms with low cash balances, large productivity shocks lead them to resort to external finance to fund high optimal investment. However, for firms with large cash balances, large productivity shocks actually lead to more distributions. The reason is decreasing returns. Although investment rises with the z shock, it does not rise as fast as cash flow. Because the firm has sufficient internal liquidity to fund investment, excess funds are then distributed to shareholders.

2 Comparative Statics

We now turn to examining some of the model's comparative statics properties. Figure 2 presents the first set of these exercises. Each panel in the figure depicts the sensitivity of one or more variables to a parameter. To construct each panel, we solve and simulate the model 20 times, each time corresponding to a different value of the parameter in question. For each of these 20 simulations, we calculate the average over the 100,000 simulated time periods of the variable of interest. Finally, we plot these 20 averages against the 20 different parameter values. The model parameterization is identical to that used to construct Figure 1.

In the left column of panels, the ratio of cash to assets is on the vertical axis. In the right column of panels the vertical axis contains the fractions of investment funded by cash and by external funds. In each of the 10 panels in this figure, a parameter of interest is on the horizontal axis. We consider the five parameters that govern the basic functioning of the model: the cost of external finance, ϕ, the standard deviation of the productivity shock process, σ_v, the serial correlation of the driving process, ρ, the curvature of the profit function, θ, and the quadratic adjustment costs parameter, a.[3] We let each of these parameters take values in a range whose center is roughly the estimate from Table 3: $\phi \in [0, 0.25]$, $\sigma_v \in [0.1, 0.5]$, $\rho \in [0.5, 0.75]$, $\theta \in [0.5, 0.9]$, and $a \in [0, 0.9]$.

We first examine the cost of external finance, ϕ. These results are in the top two panels of Figure 2. The leftmost panel shows that cash increases with the cost of external finance. Intuitively, as the cost of external finance rises, internal flexibility becomes more valuable, and firms hold more cash. The right panel also reflects this trade-off. As the cost of external finance rises, the figure shows that the firm substitutes external for internal finance.

The next four rows of panels present the results from the parameters that govern technology. First, we see a positive relation between the standard deviation of the productivity shock process, σ_v, and cash holdings. As σ_v rises, the firm is more likely to see a very good realization of the productivity shock. It knows that it may have to tap external finance more often, and it holds higher cash balances to avoid costly external financing. The effect of the serial correlation of the shock process, ρ, on cash is also monotonic and positive. Because the variance of an $AR(1)$ process increases with its serial correlation, we see effects similar

[3]For brevity we omit the fixed cost of adjusting the capital stock, γ, because this parameter is never significantly different from zero in any of our estimations.

to those of increasing σ_v. In addition, as ρ increases, the frequency of large investments increases because a positive productivity shock signals not only that capital is productive today, but also that it will continue to be productive. The firm therefore wants higher cash balances to lower the probability of needing external finance when it makes these large investments. The firm also anticipates needing to fund investment in several consecutive years, which also gives rise to higher optimal cash balances. In fourth row we see that average cash rises with profit function curvature, θ. As it rises, the profit function becomes flatter, and profit shocks therefore have a larger effect on the optimal capital stock. Therefore, both the variance and average size of desired investments rises. The firm then holds more cash because large investments imply a greater likelihood of needing external finance. Finally, the fifth row shows that the convex adjustment cost, a, decreases cash holding. As a increases, the firm makes smaller, more frequent investments, which rarely require outside funds, so precautionary cash holdings fall.

For these σ_v, ρ, and θ, the fractions of investment financed with both internal cash and external finance move in the same direction as average cash. The reason is large income effects. In all three cases, higher cash balances arise because of larger optimal investments. Therefore, in response to a positive shock, the firm wants to invest more in all assets, so its need for external finance rises, and its demand for cash balances also rises. In the case of the convex adjustment cost, a, the intuition is that the fractions financed with cash and external funds fall because internal cash flow is sufficient to fund small investments.

Figure 3 examines the effects of the parameters that shape the misalignment of incentives: the fraction of the firm owned by the manager, β, the profit sharing parameter, α, and perquisite consumption parameter, s. We allow β to range from 0 to 0.1, α to range from 0 to 0.01, and s to range from 0 to 0.0002. These parameter ranges encompass the estimates that we present below.

The results on managerial ownership are in the top two panels of the figure. We see that if the manager owns a tiny fraction of the firm, his incentives to consume perquisites are high, so he holds large cash balances. This incentive effect declines monotonically as his ownership interest rises. We also see that the fractions of investment financed with both cash and external finance fall. With a small ownership interest, the manager funnels too much of the firms cash flow into liquid assets and therefore needs more external finance.

The second row of panels shows that cash holding rises with the manager's profit sharing parameter, α. In this case, the manager perceives capital to be more productive than it actually is, so the manager invests optimally in larger amounts, which entails more outside financing. Therefore, higher precautionary cash balances are warranted. This intuition is also evident in the right panel, as the fractions of investment financed with both cash and external finance increase. If the manager is investing more than internal cash flows warrant, sources of finances other than these cash flows end up being more important.

The third panel shows a positive monotonic effect of resource diversion on average cash balances. This monotonicity is not obvious. If the manager has leeway to divert both cash flows and stocks, the manager ought to have a greater incentive to accumulate both capital and cash. However, the manager cannot "steal" capital itself, only the flow of profits generated by the capital, so diverting cash has a higher marginal benefit to the manager. Not surprisingly, the fraction of investment financed with cash rises with s, and with a large amount of cash on hand, the fraction of investment financed externally falls.

3 Data

The estimation of the model requires merging data from various sources. We collect financial statements from Compustat and managerial compensation data from ExecuComp. For some of our split sample estimations we also use governance data from IRRC (governance and blockholders), and institutional ownership data from Thomson Financial. Following the literature, we remove all regulated (SIC 4900-4999) and financial (SIC 6000-6999) firms. Observations with missing values for the SIC code, total assets, the gross capital stock, market value, and cash are also excluded from the final sample. We also require that a firm have at least two years of data because we need to lag some of the variables. As a result of these selection criteria, we obtain a panel data set with 9,274 observations for 1,438 firms, for the time period between 1992 and 2008 at an annual frequency. Table 1 contains specific definitions of the variables we use.

Table 2 provides descriptive statistics for our sample. The first panel contains firm-level accounting variables. Most of these figures are representative of generic samples from Compustat, except that our sample contains large firms. Median firm assets is 1.3 billion, a number approximately 4.5 times larger than median of all firms in Compustat over the same sample period. The reason for the bias toward large firms is the availability of compensation and governance data. The next panel contains our two compensation variables: managerial bonuses and managerial ownership. Our ownership variables show that managers on average own small fractions of the firm. This small fraction leaves a great deal of room for a manager's preference for perquisite consumption to affect his decisions. The next two panels contain variables related to corporate governance. We use these variables to determine whether groups of firms sorted on these variables have different estimates of empire building and consequently different levels of cash holdings.

4 Estimation and Identification

In this section, we explain how we estimate the model derived in Section 2. Intuitively, we use observed features of managerial contracts and observed financing and investment choices to infer estimates of the average managerial perquisite consumption and distorted perception of capital productivity. This section discusses in detail the estimation technique and the identification strategy.

4.1 Estimation

We estimate most of the model parameters using simulated method of moments. However, we estimate some of the model parameters separately. For example, we estimate the risk-free interest rate, r, to equal 0.0112, which is the average over our sample period of the three-month t-bill rate minus the rate of growth of the consumer price index. Similarly, managerial ownership, β, is the sample average of a direct component, managerial share ownership, and an indirect component, ownership due to options awards, and is equal to 0.0512. Finally, we set the corporate tax rate equal to 20%, which is an approximation of the corporate tax rate relative to personal tax rates.[4]

We then estimate the following 8 parameters using simulated method of moments: the external financing cost parameter, ϕ; the standard deviation and autocorrelation of the shock process, σ_ε and ρ; the curvature of the profit function, θ; the quadratic adjustment cost parameter, a; the depreciation rate, δ; the perquisite consumption parameter, s; and the profit sharing parameter, α. Recall that this parameter captures not only profit sharing, but also managerial perception of capital productivity. We omit the fixed adjustment cost parameter, γ, from our estimations because including it usually results in a tiny and imprecise estimate. This result makes sense given the large size of the firms in our sample, most of which have smooth investment.

Simulated method of moments, although computationally cumbersome, is conceptually simple. First, we generate a panel of simulated data using the numerical solution to the

[4]This parameter affects the firm's average cash holding, because interest taxation is the main cost of holding cash. One can add a parameter to the model that captures the discrepancy between this statutory rate and the average effective rate in the sample. However, as in DeAngelo, DeAngelo, and Whited (2011), this parameter is difficult to estimate and affects the results little.

model. Specifically, we take a random draw from the distribution of ε' (conditional on ε), and then compute $U(k,c,\varepsilon)$, $p(k,c,\varepsilon)$, $V(k,c,\varepsilon)$, and various functions of $U(\cdot)$, $V(\cdot)$, k, and c. We continue drawing values of ε and use these computations to generate an artificial panel of firms. Next, we calculate interesting moments using both these simulated data and actual data. The objective of SMM is then to pick the model parameters that make the actual and simulated moments as close to each other as possible. Details regarding the estimation are in the appendix.

4.2 Identification

The success of SMM relies on model identification. Global identification of a simulated moments estimator obtains when the moment restrictions equal zero if and only if the structural parameters equal their true values. A sufficient condition for identification is a one-to-one mapping between the structural parameters and a subset of the moment restrictions of the same dimension. Because our model does not yield such a closed-form mapping, to help ensure an identified model, we take care to choose moments that are sensitive to variations in the structural parameters such as the curvature of the profit function, θ, or the manager's preference for perquisite consumption, s.

We now describe and rationalize the 17 moments that we match. Because the firm's real and financial decisions are intertwined, all of the model parameters affect all of these moments in some way. We can, nonetheless, categorize the moments roughly as representing the real or financial side of the firm's decision-making problem.

The first two non-financial or "real" moments are the first and second central moments of the rate of investment, defined in the simulation as I/k. The first moment helps identify the capital depreciation rate. The second moment helps identify both the curvature of the profit function, θ, and the adjustment cost parameter. Higher a produces less volatile investment, and lower θ produces less volatile investment because the frictionless optimal capital stock varies less with the shock, ε. However, this one moment cannot identify two parameters at once, so we also include average operating income, which is not affected by the adjustment cost parameters and which is primarily affected by the curvature of the profit function. This relation can be seen by the definition of simulated operating income as $\varepsilon k^{\theta}/k$: the higher θ, the higher average operating income.

Our next two moments capture the important features of the driving process for ε. Here, we estimate a first-order panel autoregression of operating income on lagged operating income. The two moments that we match from this exercise are the autoregressive coefficient and the shock variance. We also match the serial correlation of investment. This moment is primarily affected by the smooth adjustment cost parameter but also by the serial correlation of the driving process, ρ. Our next moments are the mean and variance of Tobin's q. Simulated Tobin's q is constructed as $(V) / (k + c)$, and the mean and variance of q respond sharply to all parameters.

The next set of moments pertains to the firm's financing decisions. We include the mean, variance and serial correlation of the ratio of cash to capital c/k. We also include the covariance between investment and cash, the mean and variance of the ratios of distributions and security issuance to capital, d/k and f/k. Because our model does not include a debt/equity decision, we cannot attempt to match moments pertaining to the composition of external finance. These moments are useful for identifying the cost of external finance, ϕ.

We now discuss the identification of the managerial resource diversion parameter, s. It is important for our estimation that the manager can divert both profits and the stock of cash. If he could only divert profits, then given our data, it would be impossible to distinguish resource diversion from low profitability. If he could only divert the stock of cash, it would be impossible to distinguish resource diversion from a divergence between actual interest earned on cash, and the estimate given by our estimate of the real interest rate and our specification of the corporate tax rate. However, because he can divert both profits and cash and because these two variables are negatively correlated in the model (and in the actual data), identification is possible. Not surprisingly, the moment that is most important for identifying resource diversion is Tobin's q: the more resource diversion, the lower q.

Finally, we discuss the identification of the profit-sharing parameter, α. First, without our data on ownership and compensation, we would have to infer the value of this parameter solely from firm decisions. In this case, a high value of α implies low average profitability because the manager views the firm as being more profitable than it actually is and makes distorted investment decisions. However, many other parameters affect average profitability, so this moment alone cannot help identify α. Fortunately, this parameter does correspond directly to one moment from our compensation data: the average bonus.

The next issue in SMM is whether to match moments using an identity matrix or using a weighting scheme. Using an identity matrix implicitly puts the most weight on the moment that is the largest in absolute value. Because this implication rarely corresponds to a relevant economic or statistical objective, we match moments using the optimal weight matrix, which is the inverse of the covariance matrix of the moments. Roughly speaking, this scheme puts the most weight on the most precisely estimated moments, which is a sensible statistical objective. Because our moment vector contains separately estimated first and second moments, as well as regression coefficients, we use the influence-function approach in Erickson and Whited (2000) to calculate the covariance matrix of the moments. Specifically, we stack the influence functions for each of our moments and then form the covariance matrix by taking the inner product of this stack.

One final issue is unobserved heterogeneity in our data from Compustat. Recall that our simulations produce *i.i.d.* firms. Therefore, in order to render our simulated data comparable to our actual data we can either add heterogeneity to the simulations, or remove the heterogeneity from the actual data. We opt for the latter approach, using fixed firm and year effects in the estimation of our regression-based data moments and the estimation of variance and skewness.

5 Results

We first present the results of estimating several variants of our baseline model on our full sample. We then estimate our richest model on several subsamples, and we examine the recent secular upward trend in corporate cash. We then present robustness checks and counterfactual experiments.

5.1 Full Sample Results

Tables 3 and 4 present the estimation results, with Table 3 reporting moments calculated from our data, simulated moments, and t-statistics for the difference between the two. We report estimates from four models. The first contains no agency problems, that is, we set $\alpha = 0$, $\beta = 1$, and $s = 0$. In the second we only constrain the profit sharing parameter to be zero, in the third we constrain the resource diversion parameter to be zero, and in the fourth we estimate all eight parameters.

The most important result, in the second column, is that the no-agency model does a poor job of matching average cash balances, with the simulated moment less than half the actual moment. As seen in the third and fourth columns, adding either of the two agency problems improves the model fit along this dimension, but the difference between actual and simulated average cash is still significantly different from zero. The last column shows that it is only when we add both agency problems to the model that we are able to match average cash. This result alone points to the importance of agency in corporate decisions to hold liquid assets.

The model with both agency problems does a good job of matching most of the rest of the moments; large t-statistics accompany only five out of the 17 moments. Two of these are the serial correlation of investment and profits. This result is likely due to our model solution implicitly being under the risk neutral measure. However, we are calculating physical serial correlations. If risk premia are time varying, then a standard result from fixed-income asset pricing is that the risk-neutral and physical autocorrelations differ. The next two badly matched moments are the variances of distributions and external finance, which are lower in the model than in the data. This result happens because our model only has one source of uncertainty, whereas the data are driven by many. Finally, average actual distributions are significantly different from average simulated distributions, but they are not economically different. One moment that we do match statistically but not economically is the variance of Tobin's q, which is much larger in the data than in the model. This result is characteristic of many production based asset pricing models, which can match first moments of returns, but not second moment.

Table 4 contains the parameter estimates that correspond to the moments reported in Table 3. All of the parameter estimates are significantly different from zero, and the basic message of the two tables is the same. Adding agency issues to the model helps the model fit the data. For example, in the model with no agency, the estimate of the cost of external finance is too high. At a level of 9.2 cents per dollar of finance raised, this cost is much higher than SEO or loan origination fees. The reason for the failure of the vanilla model is that it only contains precautionary motives for holding cash, which are clearly insufficient for the model to fit the data. Adding either agency problem to the model by itself helps somewhat along this line. The profit sharing and perquisite consumption parameters are both significantly different from zero. (We postpone a discussion of the economic significance until later.) The estimate of the cost of external finance also falls, although the estimate of approximately 5.2% is still higher than transactions costs estimates. Finally, as in the case of Table 3, adding both agency problems to the model results in a better model fit. In this case the cost of external finance falls to a plausible 3.7%.

5.2 Governance Splits

Of course, SMM estimates the parameters of an "average firm", so it is useful to examine sample heterogeneity to see whether estimated representative firms differ across the sample. Along this line, we first examine the role of agency in corporate cash holdings by estimating our complete agency model on subsamples that have been sorted on several measures of corporate governance. We first discuss and motivate the different governance measures, and then we present the results from our split-sample estimations. We create each subsample by splitting the entire sample into thirds based on the variable of interest, and then discarding the middle third. This sorting scheme mitigates the possibility that better governed firms end up in the group of worse governed firms. This possibility is likely inasmuch as all of our governance measures are at best rough proxies for the inherently nebulous and difficult-to-measure concept of good governance.

Two of our governance measures are based on ownership. The first is the fraction of stock owned by institutional investors, with a higher value indicating better governance, presumably because institutions are more likely to be activist shareholders. For example, Hartzell and Starks (2003) find that high institutional ownership is negatively related to the level of executive compensation, and positively related to pay-for-performance sensitivity. The second measure is the fraction of stock owned by outside blockholders, with a higher value once again indicating better governance. As argued by Shleifer and Vishny (1986), the existence of large independent shareholders makes a takeover or a proxy contest easier.

In this case, the cost of a control challenge is smaller, and the market for corporate control puts more discipline on the manager.

Our next two measures are the commonly used governance indexes from Bebchuk, Cohen, and Ferrell (2009) and Gompers, Ishii, and Metrick (2003)—the E-index and G-index, respectively. Both indices count provisions recorded by the Investor Responsibility Research Center (IRRC) that describe shareholder rights. The E-index includes only those provisions argued by Bebchuk, Cohen, and Ferrell (2009) to be the most important for entrenching managers: staggered boards, limits to shareholder bylaw amendments, supermajority requirements for mergers, supermajority requirements for charter amendments, poison pills, and golden parachutes. The G-index counts all of the provisions documented by the IRRC. As argued in Coates (2000), all firms' possession of a latent poison pill can seriously undermine the informativeness of these indices. Nonetheless, we examine them because of their widespread use in the rest of the governance literature.

The results from our split sample estimations are in Figure 4 and Table 5. The former plots the actual ratio of cash to assets versus the model-implied ratio of cash to assets for each of our 8 samples, and the latter reports our parameter estimates. The striking result from this figure is the good job our fairly parsimonious model does of matching the average ratio of cash to assets in samples of firms with widely different levels of average cash. At least as interesting are the parameter estimates that generate these good fits in diverse subsamples.

The first two panels of Table 5 present the results from dividing the sample by institutional ownership. The estimates of the profit sharing and resource diversion parameters conform to intuition. We find a lower profit sharing parameter (a less distorted perception of capital productivity) in the sample with high institutional ownership and slightly (and statistically insignificantly) less managerial resource diversion, presumably because of a greater tendency for such firms to experience shareholder activism. Interestingly, as shown in Figure 4, this difference in the estimated agency parameters is accompanied by almost identical ratios of cash to assets for the two groups of firms. This result is puzzling because our model predicts that the firms with worse agency problems have higher cash balances. However, this apparent puzzle can be understood by looking further at the parameter estimates in Table 5. The main dimension along which these two groups differ is the capital depreciation rate. Thus, although the larger agency parameters for the low-ownership group would indicate higher cash, lower capital depreciation implies lower optimal cash for two reasons. The first is simply that slowly depreciation capital is more valuable, relative to cash, than rapidly depreciating capital; so the firm holds less cash. Second, slowly depreciating capital needs to be replaced less often, so the firm needs fewer precautionary cash balances to insure against having to obtain external finance. Although it is difficult to gauge the relative strength of these parameter differences, they obviously combine to produce almost identical cash holdings.

The next two panels contain the results for the high and low blockholder groups. Here, although the profit sharing parameter is almost identical for the two groups, it is statistically significant only for the low-blockholder group. The perquisite consumption parameter is also larger for this group, but the estimates of this parameter are not statistically significant for either group, almost certainly because of the small sample size of 1,863. Nonetheless, these results are intuitive inasmuch as firms with independent blockholders are likely to be better governed. In contrast to the case of the previous sample split, these two groups have similar parameters, so that their nearly equal average cash balances are not at all puzzling.

The next four panels contain the results for the samples split by the two governance indices. These results contradict the predictions of our model. First, all four estimates of the profit sharing parameter are significantly different from zero, with the high G-index firms having a slightly lower estimate than the other four groups. However, the estimates of the resource diversion parameter are only significant for the samples split by the E-index, with the low-index firms having a much higher estimate. Although counterintuitive, this result arises from other characteristics that differ across the groups, such as size and growth opportunities. For example, the low-index firms tend to be much younger and smaller than the high-index firms, and both of these characteristics are associated, by themselves, with higher cash holding. Finally, given the evidence in Bhagat, Bolton, and Romano (2008) that governance indices are uninformative about corporate performance, we view this result at least as much as casting doubt on the indices rather than on our model.

In sum, one of the themes that runs throughout this table is that governance in the form of either perquisite consumption or a distorted perception of capital productivity is not the only determinant of corporate cash holdings. Because accountants do not record fundamental technological characteristics such as returns to scale, adjustment costs, or serial correlation of an unobservable driving process, it is almost impossible to include all of the appropriate controls when trying to understand the relation between governance and cash via linear regressions. This table, therefore, provides some insight into the difficulty that the literature has had with finding conclusive evidence concerning the relation between governance and cash holdings.

5.3 Secular Increase in Cash

Next, we try to understand the increase in corporate cash over the last two decades documented by Bates, Kahle, and Stulz (2009). Table 6 reports the moment-matching results from two estimations of the full model that includes both resource diversion and empire building. In the first the sample runs from 1992 to 1999, and in the second from 2000 to 2008. In our sample this increase is evident in the difference between the average cash balances in the early and late parts of our sample period. In the 1990s average cash is approximately 10% of assets, and in the 2000s average cash is just over 15% of assets. Three other moments change noticeably: average investment falls, average operating profits fall, and the use of external finance decreases sharply. The model fits the data slightly worse for the second sub-period, with six moments poorly matched instead of five.

To understand these shifts, we examine our parameter estimates in Table 7. The estimates from the two sub-periods are similar. Indeed, the estimates of the two governance parameters are nearly identical, and the higher estimate of the smooth adjustment cost parameter in the second half of the sample indicates that cash ought to, if anything, be lower in the second half of the sample. However, the ownership fraction (estimated outside of SMM) falls from 0.054 to 0.048 in the second half of the sample, mostly because many executive stock options fell far out of the money during the steep stock market decline of the early 2000s. To understand whether ownership is behind the secular increase in cash, we parameterize the model according to the late sample estimates, except that we hold ownership where it was in the early sample. We find an average ratio of cash to assets of approximately 0.06, and we can match the high cash levels observed in the late sample only by decreasing ownership. Thus, we conclude that although the basic agency issues within firms remained constant, the misalignment of incentives from the lower ownership shares was an important force behind the cash buildup in the 2000s. This result is not in accord with those in Bates, Kahle, and Stulz (2009), who find increase in uncertainty has accompanied the increase in cash over the last several years. The difference in results is likely due to different sample periods because we find in Table 6 that profit variances actually decrease in the second half of the sample.

5.4 Small and Large Firms

Although we see a difference between cash holdings in the early and late periods, a larger gap exists between large and small firms. To understand this gap, we estimate our model on two subsamples, where the large firms are the upper tercile in terms of book assets and the small firms are the lower tercile. Table 8 contains the actual and simulated moments. The first striking result is that small firms hold twice as much cash as large firms. These larger cash balances are accompanied by higher variances of cash, investment, profits, Tobin's q, and external finance. The evidence thus points to substantial uncertainty surrounding the operations of small firms. Interestingly, for both groups, the model fits the data better than it does for the full sample—a result that makes sense, given that SMM estimates the parameters of an average firm, and given that the small and large groups are substantially less heterogeneous than the full sample.

Table 9 contains the corresponding parameter estimates. Three results stand out. First, small firms not surprisingly experience a higher cost of external finance. Second, the variance of the profit shock is much higher for the small than for the large firms. These two results are intuitive, both point to a higher precautionary motive in small firms. Third, both agency parameters are much higher for the small than the large firms. This result at first appears puzzling in that the managers of small firms hold a much greater fraction of the firm's equity (7.8%) than do the managers of large firms (2.5%). However, it is precisely this difference in ownership that can help explain the results. In our model agency is a combination of modeled incentives to misbehave, combined with limited ownership. The higher is ownership, the lower is the agency effect on cash. In fact, when we set the agency parameters to neutral values ($\beta = 1$, $\alpha = s = 0$), we find similar decreases in cash for both groups. We therefore conclude that uncertainty and the cost of external funds, and not agency that helps explain the difference between cash holdings of large and small firms.

5.5 Robustness Checks

We have made several simplifying assumptions to construct our model. Two of the most important are an infinitely lived manager and profits that are never negative. An infinite horizon is potentially troublesome because it might influence the manager's decision to consume perquisites. The second feature is potentially a problem given the survey evidence in Lins, Servaes, and Tufano (2010) that one of managers' primary self-reported reasons for holding cash is to avoid cash flow shortfalls. To deal with these concerns, we reestimate three versions of our model with features that capture these concerns. First, we introduce a probability that the manager is fired (η), which naturally makes the manger discount the future at a higher rate. Second, we model a fixed cost of production (κ_c), and third, we introduce a "flow" fixed cost (κ_p) that is proportional to the capital stock.[5] Table 10 shows that all four models do approximately as well at matching our set of moments. Table 11 shows that these three new parameters are all significantly different from zero. However, the estimates of the two agency parameters are largely unchanged. Interestingly, the firing probability is small at 0.0007. Because we infer this probability from managerial decisions, we interpret this result as indicating that managers perceive the probability of being fired as near zero. In sum, we conclude that although these new model parameters are statistically significant, modifying the model along these lines does little to change our basic conclusions about agency and cash.

5.6 Counterfactuals

In this section we quantify the extent to which a misalignment of incentives destroys shareholder value and changes average cash holdings. This exercise is useful because it assists in the interpretation of the economic magnitude of the profit sharing and the resource diversion parameters. It also helps answer the question of just how much governance affects corporate

[5] A further concern is risk aversion. When we allow the manager to have constant relative risk aversion utility, we cannot identify the risk aversion parameter with our current set of moments. High risk aversion and a flat production function produce results similar to those from low risk aversion and a very concave production function. Intuitively, a managerial preference to smooth cash flows can be accomplished with concave technology. We conclude that leaving risk aversion out of the model affects our estimates of production function curvature, which is not our main focus.

cash holding. We consider several scenarios: 5% increases and decreases in α and s individually, and 5% increases and decreases in these two parameters jointly. The results are in Table 12. Not surprisingly, if we change the managerial perquisite consumption parameter, we see large changes in corporate cash holding. If shareholders tolerate more wastefulness, and if managers can divert funds from cash stocks, then in a dynamic model, average cash holdings rise. Interestingly, this large change in cash is accompanied by a relatively small value loss. The reason is that the firm value function is relatively flat in the cash dimension. This result echoes the finding in Korteweg (2010) that firm value is largely flat with respect to leverage. In contrast, a 5% change in the profit sharing parameter is associated with a much smaller change in cash, but a much larger change in firm value. The reason is that this agency problem operates through distortion of real investment policy, and firm value decreases sharply when investment policy is suboptimal. Finally, when we change both parameters at once, we see effects that are marginally more than the sum of the individual effects alone, that is, the two agency problems complement each other in the firm's value function.

One obvious concern with all of these exercises is the possibility that governance and the cost of external finance are intertwined. In particular, common intuition suggests that investors would demand a higher return on invested capital if they suspected that management would not use these funds optimally. The intuition that changing governance should also change the cost of external finance cannot be captured in our model because the cost of external finance is a parameter that does not depend on governance. Therefore, our model might understate the deleterious effects of governance on shareholder value. So it is important to interpret these counterfactual exercises as lower bounds on the effects of governance.

6 Summary and Conclusions

We use structural estimation of a dynamic model of firm investment and cash accumulation to ascertain whether agency problems affect corporate cash holding decisions. We model three specific mechanisms that misalign managerial and shareholder incentives: managerial compensation based on current profits, managerial private benefits from diverting liquid resources, and limited managerial ownership of the firm. We take the model to the data and find that all three agency problems are statistically and economically important for understanding corporate cash holding. In particular, models that do not contain these features struggle to match simple data moments, whereas models that do contain these features match these moments much better. The first of the three agency problems leads managers to think that capital is more productive than it is because they get direct compensation from having a larger firm. We find a low elasticity of firm value with respect to resource diversion but a much higher elasticity with respect to the dependence of compensation on firm size.

We also use our model as a laboratory to examine whether groups of firms characterized by different measures of corporate governance produce different estimates of agency problems. Intuitively, we find that managers of firms with low institutional ownership on average divert more liquid resources than do managers of firms with high institutional ownership. However, these two groups of firms hold almost identical cash balances. The positive effect of agency on cash holding is offset by other firm characteristics that lead to lower cash balances. On the whole, we thus conclude that agency issues are but one of many important determinants of cash holding.

Because we try to quantify empirically the effects of agency on cash, we model contracts that are actually used, and we use compensation data to estimate directly parameters that describe these contracts. A separate but also interesting question is whether these contracts are optimal. Thus one extension of our approach would be to estimate a model that derives contracts as the result of a dynamic principal-agent problem. Unfortunately, these models are often couched in terms of unobservables, such as the manager's continuation utility, and they are therefore extremely difficult to estimate. Thus, one interesting avenue for future research would be to adapt dynamic principal-agent models so that they can be taken directly to the data.

Appendix

This appendix describes the numerical solution to the model and the details of our estimation procedure.

Model solution

To find a numerical solution, we need to specify a finite state space for the three state variables. We let the capital stock lie on a grid of 25 points centered at the steady state capital stock from a model with no financing or agency friction. Denoted k^*, this is the point at which the marginal product of capital for a neutral ε shock equals $r + \delta$. We let the productivity shock ε have 11 points of support, transforming (1) into a discrete-state Markov chain on the interval $[-4\sigma_v, 4\sigma_v]$ using the method in Tauchen (1986). We let c have 10 equally spaced points in the interval $[0, k^*]$. The optimal choice of c never hits the upper endpoint, although it is occasionally optimal for the firm to hold no cash. We allow the firm to choose policies in between these grid points, with 10 equally spaced choices between each grid point.

We solve the model via value function iteration on the Bellman equation (6), which produces the value function $U(k, c, \varepsilon)$ and the policy function $\{k', c'\} = p(k, c, \varepsilon)$. We solve for equity value by value function iteration on (8), using the policy function corresponding to

(6). In the subsequent model simulation, the space for ε is expanded to include 120 points, with interpolation used to find corresponding values of U, k, and c.

Estimation

We now give a brief outline of the estimation procedure, which closely follows Ingram and Lee (1991). Let x_i be an *i.i.d.* data vector, $i = 1, \ldots, n$, and let $y_{ik}(b)$ be an *i.i.d.* simulated vector from simulation k, $i = 1, \ldots, n$, and $k = 1, \ldots, K$. Here, n is the length of the simulated sample, and K is the number of times the model is simulated. We pick $n = 53,677$ and $K = 10$, following Michaelides and Ng (2000), who find that good finite-sample performance of a simulation estimator requires a simulated sample that is approximately ten times as large as the actual data sample.

The simulated data vector, $y_{ik}(b)$, depends on a vector of structural parameters, b. In our application $b \equiv (\phi, \sigma_v, \rho, \theta, a, \delta, s, \alpha)$. The goal is to estimate b by matching a set of *simulated moments*, denoted as $h(y_{ik}(b))$, with the corresponding set of actual *data moments*, denoted as $h(x_i)$. The candidates for the moments to be matched include simple summary statistics, OLS regression coefficients, and coefficient estimates from non-linear reduced-form models. Define

$$g_n(b) = n^{-1} \sum_{i=1}^{n} \left[h(x_i) - K^{-1} \sum_{k=1}^{K} h(y_{ik}(b)) \right].$$

The simulated moments estimator of b is then defined as the solution to the minimization of

$$\hat{b} = \arg\min_{b} g_n(b)' \hat{W}_n g_n(b),$$

in which $\hat{W}_n$ is a positive definite matrix that converges in probability to a deterministic positive definite matrix W. In our application, we use the inverse of the sample covariance matrix of the moments, which we calculate using the influence function approach in Erickson and Whited (2000).

The simulated moments estimator is asymptotically normal for fixed K. The asymptotic distribution of b is given by

$$\sqrt{n}\left(\hat{b}-b\right)\overset{d}{\longrightarrow}\mathcal{N}\left(0,\text{avar}(\hat{b})\right)$$

in which

$$\text{avar}(\hat{b})\equiv\left(1+\frac{1}{K}\right)\left[\frac{\partial g_n\left(b\right)}{\partial b}W\frac{\partial g_n\left(b\right)}{\partial b'}\right]^{-1}\left[\frac{\partial g_n\left(b\right)}{\partial b}W\Omega W\frac{\partial g_n\left(b\right)}{\partial b'}\right]\left[\frac{\partial g_n\left(b\right)}{\partial b}W\frac{\partial g_n\left(b\right)}{\partial b'}\right]^{-1} \tag{12}$$

in which W is the probability limit of $\hat{W}_n$ as $n\rightarrow\infty$, and in which Ω is the probability limit of a consistent estimator of the covariance matrix of $h\left(x_i\right)$.

References

Anderson, Ronald W., and Andrew Carverhill, 2011, Liquidity and capital structure, *Review of Financial Studies*, forthcoming.

Bates, Thomase W., Kathleen M. Kahle, and Rene M. Stulz, 2009, Why do U.S. firms hold so much more cash than they used to?, *Journal of Finance* 64, 1985–2021.

Bebchuk, Lucian, Alma Cohen, and Allen Ferrell, 2009, What matters in corporate governance?, *Review of Financial Studies* 22, 783–827.

Bhagat, Sanjai, Brian J. Bolton, and Roberta Romano, 2008, The promise and peril of corporate governance indices, *Columbia Law Review* 108, 1803–1881.

Bolton, Patrick, Hui Chen, and Neng Wang, 2011, A unified theory of tobin's q, corporate investment, financing, and risk management, *Journal of Finance* 66, 1545–1578.

Coates, John C., 2000, Takeover defenses in the shadow of the pill: A critique of the scientific evidenc, *Texas Law Review* 79, 271–382.

Cooper, Russell, and John C. Haltiwanger, 2006, On the nature of capital adjustment costs, *Review of Economic Studies* 73, 611–634.

DeAngelo, Harry, Linda DeAngelo, and Toni M. Whited, 2011, Capital structure dynamics and transitory debt, *Journal of Financial Economics* 99, 235–261.

Dittmar, Amy, and Jan Mahrt-Smith, 2007, Corporate governance and the value of cash holdings, *Journal of Financial Economics* 83, 599–634.

Erickson, Timothy, and Toni M. Whited, 2000, Measurement error and the relationship between investment and q, *Journal of Political Economy* 108, 1027–1057.

Gabaix, Xavier, and Augustin Landier, 2008, Why has ceo pay increased so much?, *Quarterly Journal of Economics* 123, 49–100.

Gomes, Joao F., 2001, Financing investment, *American Economic Review* 91, 1263–1285.

Gompers, Paul A., Joy L. Ishii, and Andrew Metrick, 2003, Corporate governance and equity prices, *Quarterly Journal of Economics* 118, 107–155.

Harford, Jarrad, 1999, Corporate cash reserves and acquisitions, *Journal of Finance* 54, 1969–1697.

——— , Sattar Mansi, and William Maxwell, 2008, Corporate governance and firm cash holdings, *Journal of Financial Economics* 87, 535–55.

Hartzell, Jay C., and Laura T. Starks, 2003, Institutional investors and executive compensation, *Journal of Finance* 58, 2351–2374.

Hennessy, Christopher A., and Toni M. Whited, 2005, Debt dynamics, *Journal of Finance* 60, 1129–1165.

———, 2007, How costly is external financing? Evidence from a structural estimation, *Journal of Finance* 62, 1705–1745.

Ingram, Beth F., and Bong-Soo Lee, 1991, Simulation and estimation of time series models, *Journal of Econometrics* 47, 197–205.

Jensen, Michael C., 1986, The agency costs of free cash flow: Corporate finance and takeovers, *American Economic Review* 76, 323–329.

Korteweg, Arthur G., 2010, The net benefits to leverage, *Journal of Finance* 65, 2137–2170.

Lins, Karl, Henri Servaes, and Peter Tufano, 2010, What drives corporate liquidity? An international survey of cash holdings and lines of credit, *Journal of Financial Economics* 98, 160–176.

Michaelides, Alexander, and Serena Ng, 2000, Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators, *Journal of Econometrics* 96, 231–266.

Mikkelson, Wayne, and Megan Partch, 2003, Do persistent large cash reserves hinder performance?, *Journal of Financial and Quantitative Analysis* 38, 275–294.

Morellec, Erwan, Boris Nikolov, and Norman Schürhoff, 2009, Corporate governance and capital structure dynamics, *Joural of Finance*, forthcoming.

Murphy, Kevin, 1999, Executive compensation, in Orley Ashenfelter, and David Card, ed.: *Handbook of Labor Economics* (Elsevier: Amsterdam).

Opler, Tim, Lee Pinkowitz, Rene M. Stulz, and Rohan Williamson, 1999, The determinants and implications of corporate cash holdings, *Journal of Financial Economics* 52, 3–43.

Riddick, Leigh A., and Toni M. Whited, 2009, The corporate propensity to save, *Journal of Finance* 64, 1729–1766.

Shleifer, Andrei, and Robert W. Vishny, 1986, Greenmail, white knights, and shareholders' interest, *RAND Journal of Economics* 17, 293–309.

Stokey, Nancy L., and Robert E. Lucas, 1989, *Recursive Methods in Economic Dynamics* (Harvard University Press: Cambridge, MA).

Tauchen, George, 1986, Finite state markov-chain approximations to univariate and vector autoregressions, *Economics Letters* 20, 177–181.

Taylor, Lucian, 2010, Why are CEOs rarely fired? Evidence from structural estimation, *Journal of Finance* 65, 2051–2087.

Table 1: Variable Definitions.

Table 1 presents variable definitions and data sources.

Variable (Data Source)	Variable Definition
Investment and Financial Characteristics (Compustat)	
Cash	Cash and Short-Term Investments (CHE) / Assets - Total (AT)
Investment	Capital Expenditures (CAPX) - Sale of Property (SPPE) / Property Plant and Equipment - Total (Gross) (PPEGT)
Cash Flow	Earnings Before Interest (EBITDA) / Assets - Total (AT)
Book Equity	Stockholders Equity - Total (SEQ) + Deferred Taxes and Investment Tax Credit (TXDITC) - Preferred/Preference Stock (Capital) - Total (PSTK) if (PSTK) missing then Preferred Stock Redemption Value (PSTKRV) if (PSTKRV) missing then Preferred Stock Liquidating Value (PSTKL)
Book Debt	Assets - Total (AT) - Book Equity
Market-to-Book	(Common Shares Outstanding (CSHO) * Price Close - Annual Fiscal Year (PRCC_F) + Book Debt (BD)) / Assets - Total (AT)
External Financing	(Long-Term Debt Issuance (DLTIS) - Long-Term Debt Reduction (DLTR) + Sale of Common and Preferred Stock (SSTK)) / Assets - Total (AT)
Distributions	(Dividends Common/Ordinary (DVC) + Dividends - Preferred/Preference (DVP) Purchase of Common and Preferred Stock (PRSTKC)) / Assets - Total (AT)
Executive Compensation (ExecuComp)	
Managerial compensation variables are computed for the 5 highest paid executives of the firm.	
Managerial Bonus	Bonus (BONUS) / Assets - Total (AT)
Managerial Ownership	Shares Owned - Options Excluded (SHROWN_EXCL_OPTS) / Common Shares Outstanding (CSHO)
Managerial Own. & Options	(Shares Owned - Options Excluded (SHROWN_EXCL_OPTS) + Unexercised Exercisable Options (OPT_UNEX_EXER_NUM)) / Common Shares Outstanding (CSHO)
Managerial Own. & Options II	(Shares Owned - Options Excluded (SHROWN_EXCL_OPTS) + Unexercised Exercisable Options (OPT_UNEX_EXER_NUM) + Unexercised Unexercisable Options (OPT_UNEX_UNEXER_NUM)) / Common Shares Outstanding (CSHO)
Institutional Ownership (Thompson Financial)	
Institutional ownership	Fraction of stock owned by institutional investors
Blockholders (IRRC blockholders)	
Blockholder ownership	Fraction of stock owned by outside blockholders
Anti-Takeover Provisions (IRRC governance)	
G-index	24 anti-takeover provisions index by Gompers, Ishii, and Metrick (2003)
E-index	6 anti-takeover provisions index by Bebchuk, Cohen, and Farell (2004)
Risk-free rate (FED)	
Risk-free rate	Average T-bill rate

Table 2: Descriptive Statistics.

Table 2 presents descriptive statistics for the main variables used in the estimation. The sample is based on Compustat Annual Industrial Files, ExecuComp, IRRC (governance and blockholders), and Thompson Financial. The sample covers the period from 1992 to 2008 at the annual frequency. Table 1 provides a detailed definition of the variables.

	Mean	S.D.	25%	50%	75%	Obs
Investment and Financial Characteristics						
Cash	0.134	0.163	0.019	0.064	0.189	9,274
Investment	0.124	0.104	0.061	0.097	0.154	9,274
Cash Flow	0.158	0.111	0.099	0.154	0.218	9,274
Market-to-Book	2.012	1.304	1.222	1.591	2.305	9,274
External financing	0.037	0.110	-0.007	0.011	0.050	9,274
Distributions	0.045	0.064	0.001	0.020	0.059	9,274
Depreciation	0.107	0.071	0.065	0.087	0.125	9,274
Book Assets (in billions)	4.903	10.179	0.522	1.305	3.987	9,274
Managerial Characteristics						
Bonus (in bps)	0.114	0.163	0.014	0.057	0.144	9,274
Ownership	0.036	0.075	0.003	0.008	0.028	9,274
Ownership + Options	0.051	0.074	0.011	0.025	0.055	9,274
Ownership + Options II	0.062	0.081	0.017	0.036	0.072	9,274
Ownership Structure						
Institutional Ownership	0.665	0.186	0.558	0.690	0.803	4,708
Blockholder Ownership	0.169	0.142	0.060	0.147	0.256	1,863
Anti-Takeover Provisions						
G-index	7.341	2.597	6.000	7.000	9.000	5,370
E-index	2.391	1.272	1.000	2.000	3.000	5,370

Table 3: Simulated Moments Estimation.

Calculations are based on a sample of nonfinancial, unregulated firms from the annual 2009 COMPUSTAT industrial files. The sample period is from 1992 to 2008. The estimation is done with SMM, which chooses structural model parameters by matching the moments from a simulated panel of firms to the corresponding moments from the data. The table reports the simulated and actual moments and the t-statistics for the differences between the corresponding moments. All moments are self-explanatory, except the serial correlation and innovation to income. These moments are the slope coefficient and error variance from a first order autoregression of the ratio of income to assets. T-statistics for the equality of the simulated moments with the data moments are under the simulated moments in parentheses.

		Simulated moments			
	Actual moments	No Agency	Tunneling	Profit Sharing	Both
Average cash (c/k)	0.1336	0.0624	0.0714	0.1041	0.1304
		(6.4178)	(6.8423)	(2.7302)	(0.4450)
Variance of cash (c/k)	0.0042	0.0039	0.0042	0.0037	0.0047
		(0.5580)	(-0.2405)	(1.7237)	(-1.4944)
Serial correlation of cash (c/k)	0.9297	0.9307	0.8966	0.9230	0.9087
		(-0.0203)	(0.7156)	(0.1729)	(0.5768)
Average investment (I/k)	0.1243	0.1379	0.1328	0.1287	0.1322
		(-1.4031)	(-0.8565)	(-0.4589)	(-0.8221)
Variance of investment (I/k)	0.0058	0.0041	0.0044	0.0025	0.0036
		(0.8600)	(0.8170)	(2.1362)	(1.2209)
Serial correlation of investment (I/k)	0.5701	0.5166	0.4371	0.4144	0.3791
		(0.5462)	(1.3917)	(1.6148)	(2.0673)
Average profits (εk^{θ})	0.1585	0.1603	0.1500	0.1443	0.1446
		(-0.1504)	(0.6887)	(1.1683)	(1.1738)
Error variance of profits (εk^{θ})	0.0042	0.0022	0.0019	0.0017	0.0022
		(0.3054)	(0.3575)	(0.3767)	(0.3278)
Serial correlation of profits (εk^{θ})	0.8037	0.6379	0.6180	0.5682	0.5764
		(2.3439)	(3.3577)	(3.9566)	(3.2237)
Average Tobin's q $(V/(k+c))$	2.0120	1.9533	1.7432	1.8112	1.8803
		(0.4801)	(2.2375)	(1.4660)	(1.0730)
Variance of Tobin's q $(V/(k+c))$	0.5153	0.0244	0.0083	0.0042	0.0053
		(1.7989)	(1.8507)	(1.8595)	(1.8591)
Average external finance (f/k)	0.0369	0.0447	0.0305	0.0216	0.0240
		(-1.0545)	(0.7455)	(1.7045)	(1.6093)
Variance of external finance (f/k)	0.0092	0.0023	0.0009	0.0005	0.0006
		(2.0506)	(2.5020)	(2.5778)	(2.5984)
Average of distributions (d)	0.0447	0.0376	0.0382	0.0321	0.0358
		(1.8235)	(1.0656)	(1.5902)	(2.4438)
Variance of distributions (d)	0.0021	0.0009	0.0011	0.0006	0.0009
		(1.8551)	(1.6562)	(2.1352)	(2.1799)
Covariance of cash and investment $cov(c/k, I/k)$	0.0002	0.0000	0.0002	0.0002	0.0004
		(0.4296)	(0.0513)	(0.0853)	(-0.2615)
Average bonus	0.0011	0.0002	0.0010	0.0002	0.0011
		(4.0734)	(1.6381)	(4.5141)	(1.3230)

Table 4: Structural Parameter Estimates.

Table 4 reports the estimated structural parameters, with standard errors in parentheses. ϕ is the linear costs of external finance. σ_v is the standard deviation of the innovation to $\ln(\varepsilon)$, in which ε is the shock to the revenue function. ρ is the serial correlation of $\ln(\varepsilon)$. θ is the curvature of the revenue function, εk^{θ}. a is the convex adjustment cost parameter. δ is the capital depreciation rate. s is the tunneling parameter. α is the profit sharing parameter. Standard errors are in parentheses under the parameter estimates.

	ϕ	σ_v	ρ	θ	a	δ	$s*10$	α
No Agency								
	0.0917	0.2728	0.6314	0.9000	0.7999	0.1359		
	(0.0443)	(0.0392)	(0.0308)	(0.0122)	(0.1371)	(0.0079)		
Tunneling								
	0.0518	0.2619	0.5304	0.8966	0.6315	0.1275	0.0011	
	(0.0259)	(0.0466)	(0.0456)	(0.0135)	(0.3444)	(0.0080)	(0.0003)	
Profit sharing								
	0.0802	0.2722	0.6214	0.8041	0.5586	0.1306		0.0069
	(0.0452)	(0.0495)	(0.0477)	(0.0271)	(0.1868)	(0.0076)		(0.0015)
Both								
	0.0367	0.2943	0.5565	0.7758	0.5853	0.1304	0.0011	0.0073
	(0.0172)	(0.0655)	(0.0387)	(0.0319)	(0.2983)	(0.0080)	(0.0001)	(0.0016)

Table 5: STRUCTURAL PARAMETER ESTIMATES: SAMPLE SPLITS.

Table 5 reports the estimated structural parameters, with standard errors in parentheses. ϕ is the linear costs of external finance. σ_v is the standard deviation of the innovation to $\ln(\varepsilon)$, in which ε is the shock to the revenue function. ρ is the serial correlation of $\ln(\varepsilon)$. θ is the curvature of the revenue function, εk^{θ}. a is the convex adjustment cost parameter. δ is the capital depreciation rate. s is the tunneling parameter. α is the profit sharing parameter. Standard errors are in parentheses under the parameter estimates. Panel A to H present results from split-sample estimations. The sample is split with respect to high and low i) institutional ownership, ii) blockholder ownership, iii) the G-index, and iv) the E-index. High and low refer to the top and bottom 35% of the distribution, respectively.

ϕ	σ_v	ρ	θ	a	δ	$s*10$	α
PANEL A : Institutional Ownership - High							
0.0445	0.3310	0.5646	0.7552	0.5854	0.1533	0.0006	0.0063
(0.0451)	(0.0982)	(0.0857)	(0.0588)	(0.2516)	(0.0129)	(0.0004)	(0.0024)
PANEL B : Institutional Ownership - Low							
0.0388	0.2900	0.6069	0.8221	0.6372	0.1189	0.0014	0.0078
(0.0442)	(0.1100)	(0.0682)	(0.0435)	(0.2199)	(0.0134)	(0.0011)	(0.0031)
PANEL C : Blockholder Ownership - High							
0.0347	0.2518	0.6110	0.7846	0.6765	0.1332	0.0008	0.0065
(0.3730)	(0.1893)	(0.1536)	(0.0909)	(0.5230)	(0.0213)	(0.0028)	(0.0042)
PANEL D : Blockholder Ownership - Low							
0.0339	0.2655	0.6310	0.7502	0.6990	0.1339	0.0009	0.0073
(2.0329)	(0.1347)	(0.0854)	(0.0797)	(0.6283)	(0.0214)	(0.0003)	(0.0045)
PANEL E : G-index - High							
0.0286	0.2989	0.5912	0.7720	0.5537	0.1315	0.0007	0.0062
(0.0468)	(0.1215)	(0.1055)	(0.0532)	(0.4290)	(0.0138)	(0.0009)	(0.0024)
PANEL F : G-index - Low							
0.0513	0.2765	0.6483	0.8133	0.7919	0.1407	0.0012	0.0073
(0.0998)	(0.1461)	(0.1128)	(0.0576)	(0.6103)	(0.0192)	(0.0008)	(0.0034)
PANEL G : E-index - High							
0.0427	0.2472	0.5930	0.7437	0.4462	0.1316	0.0007	0.0075
(0.0270)	(0.1218)	(0.0733)	(0.0478)	(0.2180)	(0.0130)	(0.0003)	(0.0023)
PANEL H : E-index - Low							
0.0355	0.2973	0.5938	0.7869	0.7260	0.1481	0.0015	0.0074
(0.0510)	(0.2969)	(0.1247)	(0.0751)	(1.2525)	(0.0206)	(0.0008)	(0.0042)

Table 6: Simulated Moments Estimation: Early versus Late Sample Period.

Calculations are based on a sample of nonfinancial, unregulated firms from the annual 2009 COMPUSTAT industrial files. The sample period is from 1992 to 2008. The estimation is done with SMM, which chooses structural model parameters by matching the moments from a simulated panel of firms to the corresponding moments from the data. The first panel reports the simulated and actual moments and the t-statistics for the differences between the corresponding moments. All moments are self-explanatory, except the serial correlation and innovation to income. These moments are the slope coefficient and error variance from a first order autoregression of the ratio of income to assets. T-statistics for the equality of the simulated moments with the data moments are under the simulated moments in parentheses.

	1992–1999		2000–2008	
	Actual Moments	Estimated Moments	Actual Moments	Estimated Moments
Average cash (c/k)	0.0996	0.0925 (0.8350)	0.1540	0.1464 (1.3446)
Variance of cash (c/k)	0.0019	0.0019 (0.1582)	0.0036	0.0038 (-0.3516)
Serial correlation of cash (c/k)	0.9194	0.8838 (0.8159)	0.9271	0.9183 (0.2298)
Average investment (I/k)	0.1414	0.1472 (-0.6125)	0.1120	0.1248 (-1.6184)
Variance of investment (I/k)	0.0050	0.0026 (1.9179)	0.0037	0.0023 (1.4356)
Serial correlation of investment (I/k)	0.5499	0.3547 (2.1947)	0.5622	0.3981 (1.8300)
Average profits (εk^{θ})	0.1738	0.1643 (1.0794)	0.1490	0.1365 (1.1225)
Error variance of profits (εk^{θ})	0.0045	0.0022 (0.4934)	0.0039	0.0020 (0.3154)
Serial correlation of profits (εk^{θ})	0.8097	0.5416 (3.7626)	0.7928	0.5763 (3.7356)
Average Tobin's q ($V/(k+c)$)	2.0138	1.8631 (1.4487)	1.9961	1.8570 (1.4068)
Variance of Tobin's q ($V/(k+c)$)	0.3264	0.0071 (2.1812)	0.4079	0.0051 (2.0236)
Average external finance (f/k)	0.0471	0.0249 (2.5164)	0.0292	0.0194 (1.2695)
Variance of external finance (f/k)	0.0092	0.0007 (2.8023)	0.0072	0.0003 (2.5812)
Average of distributions (d)	0.0419	0.0362 (1.4286)	0.0469	0.0343 (3.4101)
Variance of distributions (d)	0.0011	0.0009 (0.8978)	0.0021	0.0008 (2.9578)
Covariance of cash and investment $cov(c/k, I/k)$	0.0006	0.0001 (1.2332)	0.0004	0.0003 (0.2264)
Average bonus	0.0012	0.0011 (0.2790)	0.0011	0.0010 (2.5409)

Table 7: Structural Parameter Estimates: Early versus Late Sample Period.

Table 7 reports the estimated structural parameters, with standard errors in parentheses. ϕ is the linear costs of external finance. σ_v is the standard deviation of the innovation to $\ln(\varepsilon)$, in which ε is the shock to the revenue function. ρ is the serial correlation of $\ln(\varepsilon)$. θ is the curvature of the revenue function, εk^{θ}. a is the convex adjustment cost parameter. δ is the capital depreciation rate. s is the tunneling parameter. α is the profit sharing parameter. Standard errors are in parentheses under the parameter estimates.

	ϕ	σ_v	ρ	θ	a	δ	$s*10$	α
1992-1999								
	0.0369	0.2671	0.5365	0.7989	0.5430	0.1459	0.0011	0.0070
	(0.0444)	(0.1085)	(0.0359)	(0.0218)	(0.2348)	(0.0082)	(0.0002)	(0.0012)
2000-2008								
	0.0374	0.2988	0.5533	0.7709	0.6924	0.1237	0.0011	0.0070
	(0.0336)	(0.0568)	(0.0421)	(0.0292)	(0.2416)	(0.0073)	(0.0000)	(0.0014)

Table 8: Simulated Moments Estimation: Small versus Large Firms.

Calculations are based on a sample of nonfinancial, unregulated firms from the annual 2009 COMPUSTAT industrial files. The sample period is from 1992 to 2008. The estimation is done with SMM, which chooses structural model parameters by matching the moments from a simulated panel of firms to the corresponding moments from the data. The first panel reports the simulated and actual moments and the t-statistics for the differences between the corresponding moments. All moments are self-explanatory, except the serial correlation and innovation to income. These moments are the slope coefficient and error variance from a first order autoregression of the ratio of income to assets. T-statistics for the equality of the simulated moments with the data moments are under the simulated moments in parentheses. The sample is split with respect to small and large firms. Small and large refer to the top and bottom 35% of the size distribution, respectively.

	Small Firms		Large Firms	
	Actual Moments	Estimated Moments	Actual Moments	Estimated Moments
Average cash (c/k)	0.2076	0.1843 (2.0286)	0.0848	0.0866 (-0.2181)
Variance of cash (c/k)	0.0062	0.0059 (0.7023)	0.0020	0.0030 (-2.8322)
Serial correlation of cash (c/k)	0.9224	0.9156 (0.1613)	0.9326	0.8966 (0.3731)
Average investment (I/k)	0.1363	0.1384 (-0.1860)	0.1106	0.1398 (-1.9038)
Variance of investment (I/k)	0.0069	0.0039 (1.7334)	0.0028	0.0037 (-0.6621)
Serial correlation of investment (I/k)	0.5128	0.3967 (1.0607)	0.6602	0.4016 (1.2869)
Average profits (εk^{θ})	0.1353	0.1447 (-0.5488)	0.1677	0.1562 (0.5904)
Error variance of profits (εk^{θ})	0.0064	0.0027 (0.4345)	0.0029	0.0017 (0.1355)
Serial correlation of profits (εk^{θ})	0.8116	0.5983 (4.0991)	0.7767	0.5852 (1.5539)
Average Tobin's q ($V/(k+c)$)	2.1078	1.9074 (1.3880)	1.9671	1.8479 (0.5267)
Variance of Tobin's q ($V/(k+c)$)	0.6106	0.0076 (1.7739)	0.4119	0.0066 (0.9002)
Average external finance (f/k)	0.0380	0.0242 (1.0097)	0.0325	0.0262 (0.5630)
Variance of external finance (f/k)	0.0090	0.0006 (1.8567)	0.0063	0.0006 (1.6722)
Average of distributions (d)	0.0385	0.0354 (0.6640)	0.0521	0.0377 (2.0911)
Variance of distributions (d)	0.0018	0.0009 (1.5462)	0.0018	0.0011 (0.8930)
Covariance of cash and investment $cov(c/k, I/k)$	0.0008	0.0004 (0.3823)	-0.0001	0.0003 (-0.4804)
Average bonus	0.0021	0.0017 (3.2121)	0.0004	0.0005 (-2.5941)

Table 9: Structural Parameter Estimates: Small versus Large Firms.

Table 9 reports the estimated structural parameters, with standard errors in parentheses. ϕ is the linear costs of external finance. σ_v is the standard deviation of the innovation to $\ln(\varepsilon)$, in which ε is the shock to the revenue function. ρ is the serial correlation of $\ln(\varepsilon)$. θ is the curvature of the revenue function, εk^{θ}. a is the convex adjustment cost parameter. δ is the capital depreciation rate. s is the tunneling parameter. α is the profit sharing parameter. Standard errors are in parentheses under the parameter estimates. The sample is split with respect to small and large firms. Small and large refer to the top and bottom 35% of the size distribution, respectively.

	ϕ	σ_v	ρ	θ	a	δ	$s * 10$	α
Small Firms								
	0.0620	0.3169	0.5713	0.7665	0.6887	0.1365	0.0017	0.0117
	(0.0446)	(0.1061)	(0.0584)	(0.0328)	(0.4491)	(0.0123)	(0.0002)	(0.0023)
Large Firms								
	0.0269	0.2451	0.5717	0.8054	0.4907	0.1380	0.0004	0.0030
	(0.0184)	(0.0685)	(0.0654)	(0.0308)	(0.1513)	(0.0078)	(0.0002)	(0.0007)

Table 10: Simulated Moments Estimation: Robustness Tests.

Calculations are based on a sample of nonfinancial, unregulated firms from the annual 2009 COMPUSTAT industrial files. The sample period is from 1992 to 2008. The estimation is done with SMM, which chooses structural model parameters by matching the moments from a simulated panel of firms to the corresponding moments from the data. The table reports the simulated and actual moments and the t-statistics for the differences between the corresponding moments. All moments are self-explanatory, except the serial correlation and innovation to income. These moments are the slope coefficient and error variance from a first order autoregression of the ratio of income to assets. T-statistics for the equality of the simulated moments with the data moments are under the simulated moments in parentheses.

		Simulated moments			
	Actual Moments	Main Specification	Probability to Fire the Manager	Fixed Cost of Production	Proportional Cost of Production
Average Cash (c/k)	0.1336	0.1304 (0.4450)	0.1201 (1.6615)	0.1329 (0.1175)	0.1231 (1.1971)
Variance of Cash (c/k)	0.0042	0.0047 (-1.4944)	0.0063 (-6.3185)	0.0054 (-5.1159)	0.0041 (0.3130)
Serial correlation of cash (c/k)	0.9297	0.9087 (0.5768)	0.8855 (1.1377)	0.9051 (0.5394)	0.9168 (0.3764)
Average Investment (I/k)	0.1243	0.1322 (-0.8221)	0.1295 (-0.5428)	0.1324 (-0.8732)	0.1335 (-1.0138)
Variance of Investment (I/k)	0.0058	0.0036 (1.2209)	0.0068 (-0.6260)	0.0050 (0.5633)	0.0050 (0.6435)
Serial correlation of investment (I/k)	0.5701	0.3791 (2.0673)	0.3699 (1.9761)	0.3814 (1.8575)	0.3977 (1.8983)
Average profits (εk^{θ})	0.1585	0.1446 (1.1738)	0.1412 (1.4167)	0.1419 (1.3491)	0.1467 (1.2131)
Error variance of profits (εk^{θ})	0.0042	0.0022 (0.3278)	0.0022 (0.3112)	0.0024 (0.2911)	0.0050 (-0.1281)
Serial correlation of profits (εk^{θ})	0.8037	0.5764 (3.2237)	0.5786 (4.7181)	0.5690 (3.6652)	0.5601 (3.7422)
Average Tobin's q ($V/(k+c)$)	2.0120	1.8803 (1.0730)	1.8144 (1.5134)	1.8282 (1.5601)	1.9506 (0.5831)
Variance of Tobin's q ($V/(k+c)$)	0.5153	0.0053 (1.8591)	0.0068 (1.8529)	0.0055 (1.8881)	0.0143 (1.9009)
Average external finance (f/k)	0.0369	0.0240 (1.6093)	0.0295 (0.8855)	0.0285 (1.0083)	0.0433 (-0.6880)
Variance of external finance (f/k)	0.0092	0.0006 (2.5984)	0.0009 (2.4771)	0.0008 (2.5482)	0.0015 (2.3160)
Average of distributions (d)	0.0447	0.0358 (2.4438)	0.0427 (0.2842)	0.0407 (0.7998)	0.0523 (-2.0564)
Variance of distributions (d)	0.0021	0.0009 (2.1799)	0.0015 (1.2760)	0.0012 (1.9143)	0.0019 (0.5914)
Covariance of cash and investment $cov(c/k, I/k)$	0.0002	0.0004 (-0.2615)	0.0007 (-0.9563)	0.0005 (-0.6854)	0.0002 (0.1830)
Average bonus	0.0011	0.0011 (1.3230)	0.0011 (1.4329)	0.0009 (3.6834)	0.0001 (4.6175)

Table 11: Structural Parameter Estimates: Robustness Tests.

Table 11 reports the estimated structural parameters, with standard errors in parentheses. ϕ is the linear costs of external finance. σ_v is the standard deviation of the innovation to $\ln(\varepsilon)$, in which ε is the shock to the revenue function. ρ is the serial correlation of $\ln(\varepsilon)$. θ is the curvature of the revenue function, εk^{θ}. a is the convex adjustment cost parameter. δ is the capital depreciation rate. s is the tunneling parameter. α is the profit sharing parameter. η is the probability that the manager is fired. κ_c is the constant production cost. κ_p is the proportional production cost. Standard errors are in parentheses under the parameter estimates.

ϕ	σ_v	ρ	θ	a	δ	$s*10$	α	η	κ_c	κ_p
Panel A : Main Specification										
0.0367	0.2943	0.5565	0.7758	0.5853	0.1304	0.0011	0.0073			
(0.0172)	(0.0655)	(0.0387)	(0.0319)	(0.2983)	(0.0080)	(0.0001)	(0.0016)			
Panel B : Probability to Fire the Manger										
0.0403	0.3019	0.5788	0.7574	0.4124	0.1261	0.0010	0.0086	0.0007		
(0.1302)	(0.0593)	(0.0611)	(0.0789)	(0.1328)	(0.0082)	(0.0010)	(0.0054)	(0.0002)		
Panel C : Fixed Cost of Production										
0.0217	0.3038	0.5489	0.8092	0.5009	0.1300	0.0010	0.0046		66.16	
(0.0173)	(0.0549)	(0.0506)	(0.0536)	(0.1522)	(0.0084)	(0.0000)	(0.0034)		(16.23)	
Panel D : Proportional Cost of Production										
0.0484	0.2863	0.5661	0.7717	0.7027	0.1310	0.0010	0.0075			0.0746
(0.0446)	(0.0512)	(0.0680)	(0.0117)	(0.3318)	(0.0088)	(0.0000)	(0.0017)			(0.0204)

Table 12: Value Loss and Changes in Cash Holding.

Table 12 reports the value loss and differences in cash balances due to the misalignment of incentives between managers and shareholders. The value loss is expressed in terms of Tobin's q. $\hat{s}$ and $\hat{\alpha}$ are respectively the estimated tunneling and profit sharing parameters as reported in Table 4.

	Value loss (in %)	Cash Difference (in %)
5% decrease in $\hat{s}$	1.65	-18.84
5% increase in $\hat{s}$	-2.17	28.01
5% decrease in $\hat{\alpha}$	4.48	-3.36
5% increase in $\hat{\alpha}$	-4.42	5.05
5% decrease in $\hat{s}$ and $\hat{\alpha}$	6.58	-26.29
5% increase in $\hat{s}$ and $\hat{\alpha}$	-6.69	35.60

Figure 1: Policy Functions.

Figure 1 depicts optimal cash flow, investment, cash, and distributions/external financing as a function of the productivity shock z, in the revenue function zk^{θ}.

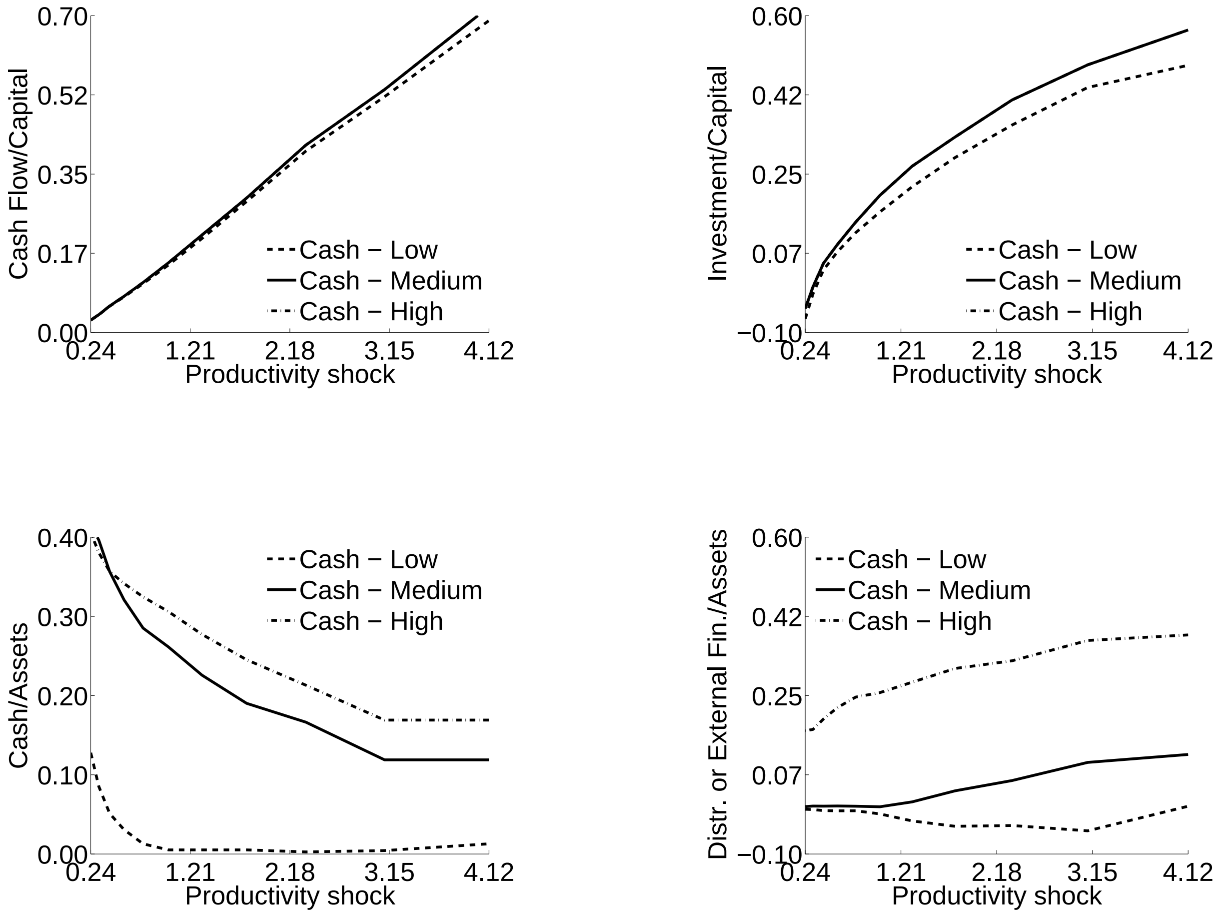

Figure 2: COMPARATIVE STATICS - TECHNOLOGY AND FINANCING.

Figure 2 depicts the relation between the cost of external financing, ϕ, the standard deviation of the innovation to $\ln(\varepsilon)$, σ_v, the serial correlation of $\ln(\varepsilon)$, ρ, the curvature of the revenue function, θ, and the convex adjustment cost a and i) the cash to assets ratio and ii) the fraction of investment funded by cash or external financing.

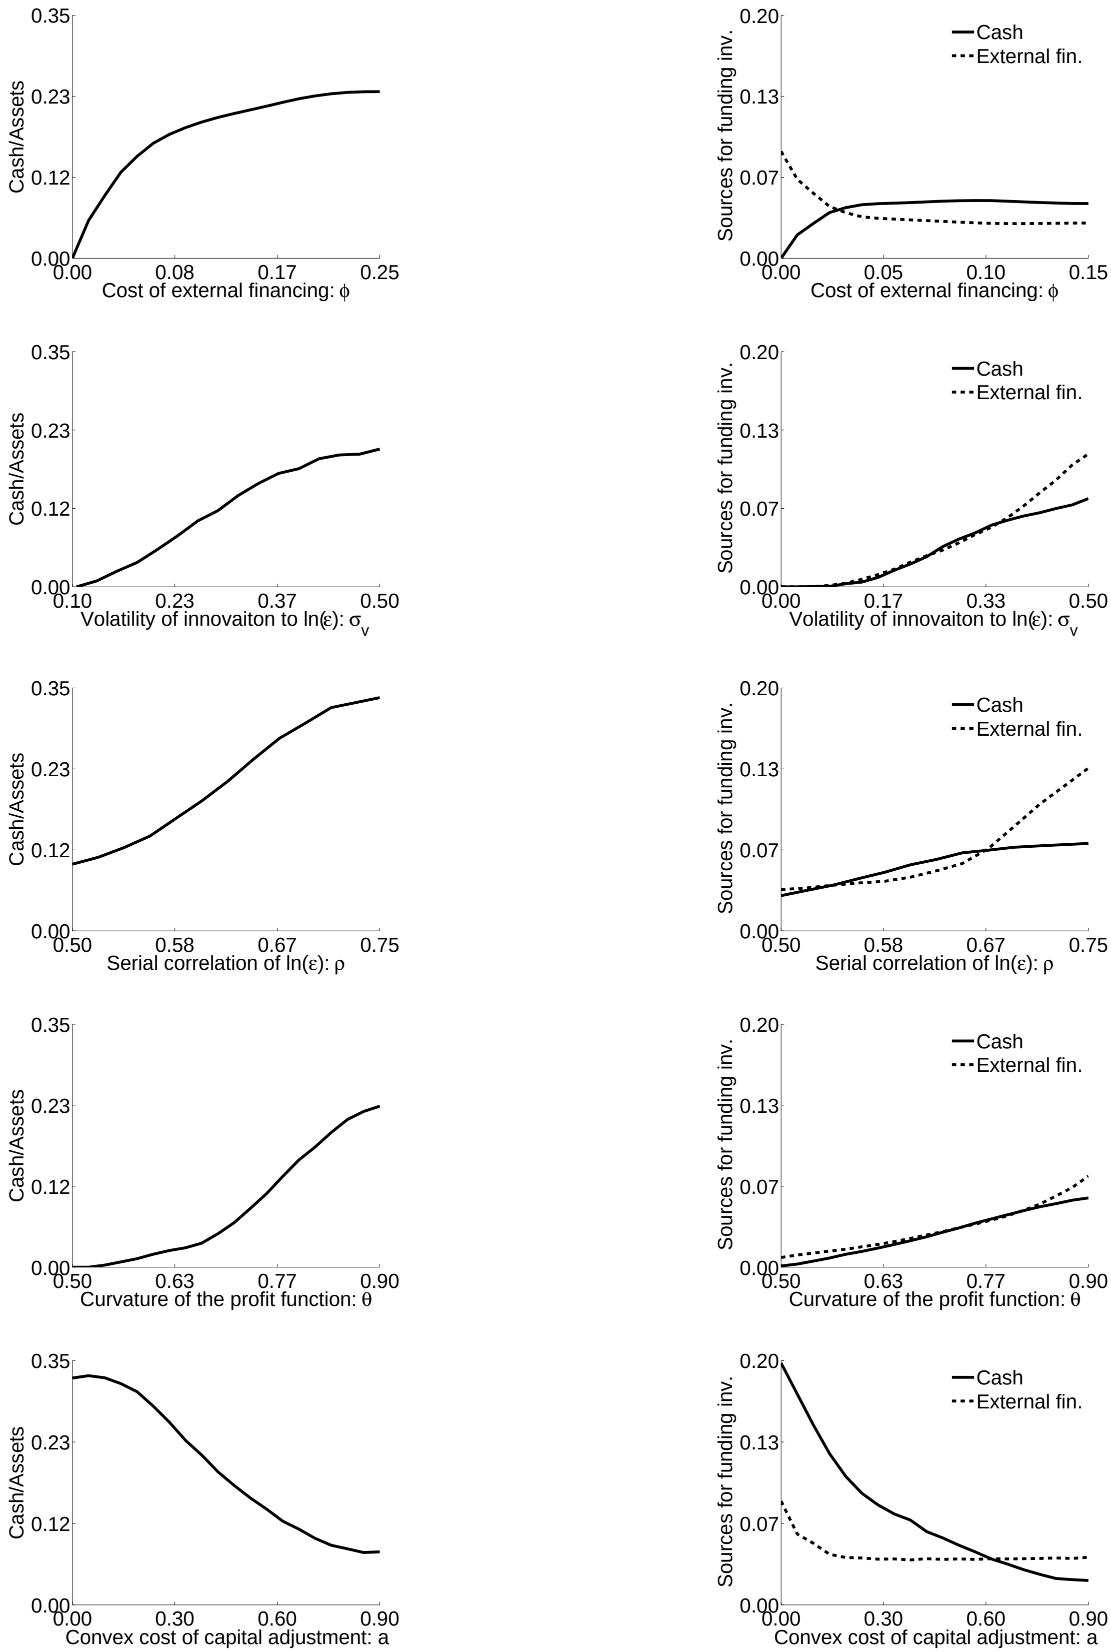

Figure 3: Comparative Statics - Agency.

Figure 3 depicts the relation between managerial ownership, β, the profit sharing parameter, α, and the tunneling parameter, s and i) the cash to assets ratio and ii) the fraction of investment funded by cash or external financing.

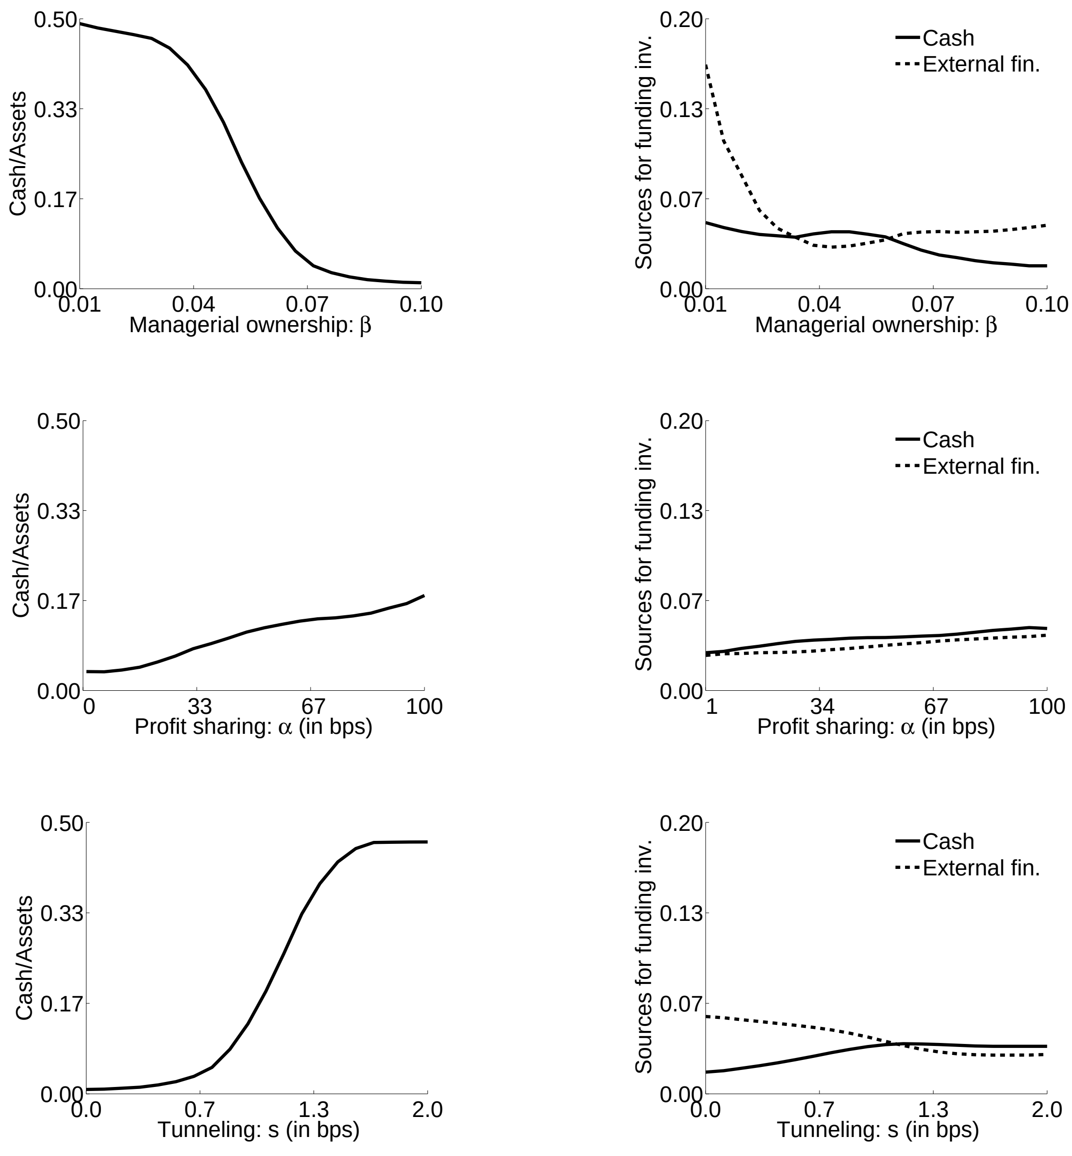

Figure 4: Estimated and Actual Moments: Governance Sample Splits.

Figure 4 depicts the relation between the estimated and the actual average cash moment for the sub-samples. The sample is split with respect to high and low i) institutional ownership, ii) blockholder ownership, iii) the G-index, iv) the E-index, and v) firm size, and early and late time periods. High and low refer to the top and bottom 35% of the distribution, respectively. Early and late time periods refer to 1992-1999 and 2000-2008, respectively.