Quantifying the Impact of Moral Hazard: Evidence from a Structural Estimation

Hengjie Ai, Dana Kiku, and Rui Li *

December 2, 2016

Abstract

We develop and estimate a dynamic equilibrium model with agency frictions to quantify the impact of moral hazard. Firm dynamics are driven by publicly observable productivity shocks and shocks that are private information to managers. We exploit empirical evidence on firm size, growth and managerial pay-performance sensitivity to identify the degree of moral hazard. We find that unobservable shocks account for about 10% of the total variation of firm output. Nonetheless, moral-hazard induced incentive pay is quantitatively significant and accounts for 52% of managerial compensation. Our welfare analysis suggests that eliminating moral hazard results in 1.4% increase in aggregate output.

Keywords: Structural Estimation, Dynamic Moral Hazard, CEO Compensation, General Equilibrium, Firm Dynamics

*Hengjie Ai (hengjie@umn.edu) is affiliated with the Carlson School of Management, University of Minnesota, Dana Kiku (dka@illinois.edu) is at the University of Illinois at Urbana-Champaign, and Rui Li (abclirui@gmail.com) is at the University of Massachusetts Boston. This paper supersedes a previous working paper titled "Moral Hazard, Investment, and Firm Dynamics". We would like to thank Gian Luca Clementi, Dean Corbae, Steven Durlauf, Kenichi Fukushima, Zhiguo He, Antonio Mello, Adriano Rampini, Marzena Rostek, Vish Viswanathan, Noah Williams, Jun Yang, Amir Yaron, Tao Zha, seminar participants at the Fuqua School of Business at Duke University, Federal Reserve Bank of Atlanta, University of Calgary, University of Minnesota, and University of Wisconsin-Madison, and participants of the 2011 Tel Aviv Finance Conference, the 2012 Conference on Financial Economics and Accounting, the 2012 ASU Sonoran Winter Conference, and the 2012 UBC Winter Finance Conference for useful comments on the topic. We also thank Song Ma and Hanbaek Lee for their excellent research assistance. All errors are our own.

1 Introduction

A large body of literature in finance and economics emphasizes the importance of moral hazard in determining managerial compensation, firm investment, and firm growth. The vast majority of this literature is *qualitative*. In this paper, we aim to analyze the *quantitative* impact of moral hazard. Moral hazard arises when output carries a noisy signal of managers' actions that are not directly observable by firms' owners. The degree of moral hazard, therefore, depends on the relative magnitude of shocks that are private information of managers and shocks that are publicly observable. In this paper, we develop a general equilibrium model where firm dynamics are affected by both observable and unobservable shocks and quantify the magnitude of moral hazard through a structural estimation. Our estimation suggests that about 10% of variation in firm output is driven by shocks that are privately observable by managers. We show that incentive pay accounts for roughly 52% of CEO compensation in the data and that eliminating moral hazard can result in 1.4% permanent increase in aggregate output.

To account for the key empirical features of CEO compensation, firm growth and firm size, we first develop a model of heterogenous firms operated by managers using neoclassical production technologies. As in Prescott and Visscher (1980), and Atkeson and Kehoe (2005), we model firm productivity as the accumulation of organization capital.[1] In our model, managers' decision to invest in organization capital is private information known only to themselves. Moral hazard arises because the level of organization capital depends on both managers' hidden effort and productivity shocks, which are only partially observable. The relative magnitude of observable and unobservable productivity shocks determines the degree of moral hazard.

Managers' private information is difficult to quantify because it cannot be measured directly from the data. We address this challenge and estimate the magnitude of unobservable shocks by embedding optimal contracting in a general equilibrium framework and exploiting equilibrium implications of moral hazard for observable firm characteristics. Our identification is based on two main predictions of our dynamic structural model. First, under the optimal contract, manager's continuation utility is the key state variable that determines firm investment and CEO compensation. Second, both observable and unobservable productivity shocks have the same effect on firm size but very different impact on continuation utility. Taken together, the two predictions of the model imply that the cross-sectional correlation between firm size and continuation utility and, hence, the cross-sectional

[1] Organization capital is defined as knowledge and firm-specific intangible capital that make physical capital and labor more productive within a firm. Examples of organization capital that improve firm productivity include organizational system, personnel information, team work, firm culture, etc.

relationship between firm size, investment and managerial compensation are determined by the relative magnitude of unobservable shocks. Thus, even though continuation utility is not directly observable, the magnitude of managers' private information can be identified from the joint empirical distribution of firm size, investment, and CEO compensation.

To formalize our theoretical predictions, we first show that the severity of the moral hazard problem decreases with the manager's share in the firm, which we define as the ratio of continuation utility to firm size. When a larger fraction of the firm's cash flow is promised to the manager, the incentive to shirk is lower. Consequently, pay-performance sensitivity declines and firm investment rate increases. Second, we establish that the manager's share in the firm decreases after positive observable shocks because of risk sharing but tends to increase significantly after good unobservable shocks due to incentive provision. Therefore, if the size of unobservable shocks is relatively small, firm investment is inversely related to firm size and pay-performance sensitivity is positively related to size. In contrast, a large magnitude of unobservable shocks implies a flat or even positive relationship between investment and size and a non-monotonic relationship between pay-performance sensitivity and firm size.

Building on our theoretical results, we estimate the magnitude of unobservable shocks using the Simulated Method of Moments (SMM) by exploiting moment conditions of the joint distribution of firm size, investment and managerial compensation. We find that in order to account for the strong inverse relationship between firm growth and size and the positive relationship between pay-performance sensitivity and size observed in the data, the magnitude of unobservable shocks must be relatively small compared with observable shocks. Our estimates suggest that only about 10% of the total variation in firm output is attributed to unobservable shocks. The difference in variances of observable and unobservable shocks is strongly statistically significant.

Based on our structural estimation, we evaluate the quantitative impact of moral hazard in general equilibrium. In our dynamic model, moral hazard reduces efficiency for several reasons. First, incentive provision requires managerial compensation to respond to unobservable shocks and, therefore, reduces risk sharing. Second, moral hazard leads to back-loaded compensation policies and, hence, affects the intertemporal allocation of managerial compensation. Third, moral hazard distorts firms' investment policies and lowers the steady-state capital in the economy. Our estimates suggest that incentive provision is quite costly — eliminating moral hazard would allow firms to reduce their CEO pay by 52% on average while providing the same level of utility to managers. Our decomposition analysis reveals that about two thirds of incentive pay is due to limited risk sharing and one third is the compensation for distortions in the intertemporal allocation of CEO compensation. We also evaluate the aggregate impact of investment distortions and find that eliminating moral hazard increases the total output of the economy by about 1.4%.

To better understand why the data implies a relatively limited amount of unobservable shocks, we consider a restricted version of our model specification by assigning a larger role to moral hazard. In particular, we impose the restriction that unobservable shocks contribute as much to the overall variance as observable shocks do. We find that relative to our baseline unrestricted model, the constrained specification has significant difficulties in accounting for the cross-sectional variation in growth rates and pay-performance sensitivity (PPS). Specifically, consistent with the data, our unrestricted model implies that firms in the top and bottom size quintile, on average, grow at the rate of 3.8% and 12.4%, respectively. Under the restricted specification, the cross-sectional dispersion in growth rates is reduced by almost a half, mostly due to the relatively low implied growth of small firms, of about 7.6% only. This outcome is quite intuitive — a higher degree of moral hazard diverts more capital from investment to incentive provision, and because small firms have limited resources, they are particularly affected and their growth is substantially reduced. Also, we show that in the data, pay-performance sensitivity monotonically increases from 0.22 for the small size quintile to 0.44 for the large size quintile. Our unrestricted model matches both the observed level and the cross-sectional variation in PPS, which as we show relies crucially on the presence yet a relatively small magnitude of unobservable shocks. The restricted specification that assumes a larger contribution of private information significantly overstates size elasticity of CEO compensation for small firms and, hence, understates the dispersion in PPS across size-sorted portfolios.

From a methodological point of view, the dynamic general equilibrium setup and the presence of both observable and unobservable shocks are two unique features that distinguish our paper from previous works. The dynamic general equilibrium setup allows us to exploit the implications of the optimal dynamic contract for the cross-sectional distribution of firm characteristics to identify the structural model parameters. In addition, general equilibrium is essential for understanding the welfare implications of moral hazard. A reduction in moral hazard is associated with higher levels of investment and capital accumulation. In a partial equilibrium, taking prices as given, firm profit is linear in organization capital (Hayashi (1982)). However, total output is decreasing return to scale with respect to organization capital at the aggregate level. Thus, in order to evaluate the impact of moral hazard across all firms, it is important to account for the decreasing return to scale in the aggregate production function, which requires a general equilibrium setup.

While most of the previous studies of moral hazard assume that all shocks are unobservable as a theoretical abstraction, our analysis emphasizes the importance of incorporating both public and private shocks in order to account for the cross-sectional distribution of firms' investment and compensation decisions. Our estimates, in fact, suggest that only a relatively small fraction of shocks that affect firm dynamics are unobservable, yet they have a quantitatively large impact on the level of managerial compensation and significant implications for the cross-sectional variation in firm growth and PPS.

Our model builds on the literature on optimal dynamic contracting with moral hazard, especially continuous-time models, for example, Biais, Mariotti, Plantin, and Rochet (2007), DeMarzo and Sannikov (2006), Sannikov (2008), He (2009), Biais, Mariotti, and Villeneuve (2010), DeMarzo, Fishman, He, and Wang (2012), and Zhu (2013).[2] The continuous-time methodology greatly improves numerical efficiency and makes it possible for us to estimate the model. For tractability, most papers in this literature use risk-neutral preferences. Because our focus is quantitative, we work in general equilibrium with neoclassical production technologies and constant relative risk aversion (CRRA) preferences.[3]

Our microfoundation approach based on organization capital is related to Prescott and Visscher (1980), and Becker (1993). Atkeson and Kehoe (2005) develop an equilibrium model to measure rents on organization capital. Luttmer (2011) presents a model with organization capital and firm dynamics. Eisfeldt and Papanikolaou (2013) study the implications of organization capital for equity returns. Our focus is to understand the implications of moral hazard for investment in organization capital.

Our paper is also related to the large literature that emphasizes the importance of moral hazard for understanding managerial compensation. A comprehensive survey of this literature is provided in Edmans and Gabaix (2016). While earlier work in this area is mostly static, several recent papers analyze the implications of dynamic contracting. Bond and Axelson (2015) develop a theory of overpaid jobs based on dynamic moral hazard. Edmans, Gabaix, Sadzik, and Sannikov (2012) present a dynamic moral hazard model of CEO compensation with private savings. Hoffmann and Pfeil (2010), and Li (2015) analyze the implications of observable shocks for CEO compensation in dynamic settings. Consistent with our findings, they also show that, in general, in dynamic environments, observable shocks affect continuation utility and policy decisions.

Our paper also contributes to the literature on structural estimation of dynamic models of CEO compensation and investment, for example, Taylor (2010, 2013), Nikolov and Schmid (2016), and Li, Whited, and Wu (2016). Several recent papers of Margiotta and Miller (2000), and Gayle and Miller (2009, 2015) analyze the identification and welfare implications of moral hazard in models with history-independent contracts (Fudenberg, Holmstrom, and Milgrom (1990)). Different from the above papers, the endogenous dynamics of firm size and

[2] For an excellent survey of this literature see Diais, Mariotti, Plantin, and Rochet (2004).
[3] A recent paper by Di Tella and Sannikov (2016) also relies on CRRA preferences.

promised utility are the key to our analysis. The presence of observable shocks and the history dependence of the optimal contract allow our model to match the salient features of the joint distribution of firm size, growth and managerial compensation, and provide identification for the structural parameters of our model.

The rest of the paper is organized as follows. We describe the setup of our model in Section 2. We provide the solution to the optimal contracting problem and lay out our identification strategy in Section 3. Section 4 describes our structural estimation and presents quantitative results. Section 5 concludes.

2 Setup of the Model

In this section, we set up an equilibrium model of investment and managerial compensation with moral hazard.

2.1 Preferences and Technology

Time is infinite and continuous. There is a continuum of firms indexed by j. As in Atkeson and Kehoe (2005), the output of firm j, denoted y_j, is produced from organization capital (H_j), physical capital (K_j), and labor (N_j) using a standard Cobb-Douglas production technology: $y_j = zH_j^{1-\nu}\left(K_j^{\alpha}N_j^{1-\alpha}\right)^{\nu}$, where z is the common technology parameter, and ν is the span-of-control parameter.[4] The operating profit of firm j is defined as

$$\pi\left(H_j\right) = \max\left\{zH_j^{1-\nu}\left(K_j^{\alpha}N_j^{1-\alpha}\right)^{\nu} - RK_j - WN_j\right\}, \tag{1}$$

where W is the equilibrium wage, and R is the equilibrium rental rate of physical capital. We assume that the markets for unskilled labor, N, and physical capital, K, are perfectly competitive. As a result, the constant return to scale of the production function implies that the profit function is linear in organization capital, i.e., $\pi\left(H_j\right) = AH_j$, where A is the equilibrium marginal product of organization capital. Due to the constant return to scale (CRS) technology, the optimal choice of K_j and N_j is proportional to H_j. Hence, in our model, all three are equivalent measures of firm size.

We assume that the accumulation of organization capital requires special skills and can

[4] Organization capital is firm-specific knowledge that makes physical capital and labor more productive. Examples of organization capital include corporate culture, team work, firm-specific human capital, etc. (see Prescott and Visscher (1980)).

be done only by a manager. The law of motion of organization capital is given by:

$$dH_{j,t} = H_{j,t}\left[\left(i_{j,t} - \delta\right)dt + \sigma^T dB_{j,t}\right], \tag{2}$$

where $i_{j,t} = \frac{I_{j,t}}{H_{j,t}}$ is the investment-to-organization capital ratio, δ is the depreciation rate of organization capital, $dB_{j,t}$ is a vector of firm-specific Brownian motion shocks, and σ is a vector of sensitivities of H to Brownian motion shocks. As in Atkeson and Kehoe (2005), and Luttmer (2011), firm dynamics are driven by the accumulation of organization capital. Different from the above literature, the accumulation of organization capital in our model is affected by managers' hidden actions.

We assume that investment in organization capital, $I_{j,t}$, is a decision known privately only to the manager. Firm-level productivity, $zH_j^{1-\nu}$, carries only a noisy signal about $I_{j,t}$ because the level of $H_{j,t}$ depends on both investment in organization capital and productivity shocks that are partly unobservable.[5] In particular, we assume that (suppressing the time subscript):

$$\sigma^T dB_j = \sigma_U dB_{U,j} + \sigma_O dB_{O,j}, \tag{3}$$

where the Brownian motion $B_{U,j}$ is unobservable to all but the manager who operates the firm, and the Brownian motion $B_{O,j}$ is public information. Due to the presence of unobservable productivity shocks, investment in organization capital cannot be perfectly inferred from the firm's output, and the relative magnitude of unobservable versus observable shocks, $\frac{\sigma_U^2}{\sigma_O^2}$, determines the degree of moral hazard and the is key parameter of interest in our empirical analysis.

At any time t, firm owners face the following budget constraint:

$$D_{j,t} + C_{j,t} + \phi\left(\frac{I_{j,t}}{H_{j,t}}\right)H_{j,t} = AH_{j,t}, \tag{4}$$

where $D_{j,t}$ is the amount of dividends, $C_{j,t}$ is managerial consumption, and $\phi\left(\frac{I_{j,t}}{H_{j,t}}\right)H_{j,t}$ is the total cost of investment in organization capital. We use a standard quadratic adjustment cost function, $\phi(i) = i + \phi_0 (i - i^*)^2$. We assume that consumption and organization capital investment decisions can only be observed by managers themselves. Moral hazard arises because managers can always substitute investment in organization capital for their own private consumption.

We assume that managers are risk-averse and that they are subject to health shocks that arrive at Poisson rate κ. Once hit by a health shock, a manager exits the economy and

[5] Note that firm-level productivity is observable because the stock of organization capital can be inferred from the observed output y_j, capital input K_j and labor input N_j.

organization capital of the firm evaporates. The time-t continuation utility of the manager is given by:

$$U_t = \left\{ E_t \left[\int_0^\infty e^{-(\beta+\kappa)s} (\beta+\kappa) C_{t+s}^{1-\gamma} ds \right] \right\}^{\frac{1}{1-\gamma}}, \tag{5}$$

where β is the discount rate of the manager, and γ is the coefficient of relative risk aversion. Firm owners are assumed to be well diversified and maximize the present value of dividends,

$$E_0 \left[\int_0^\infty e^{-(r+\kappa)t} D_t dt \right],$$

where r is the equilibrium interest rate.

In our empirical work, we identify the structural parameters that govern the relative magnitude of unobservable shocks from observable firms' decisions. Therefore, we set up our model to ensure that it is able to account for key features of the empirical distribution of firm investment, managerial compensation, and size. Below we briefly discuss the importance of various model assumptions that allow us to match the data.

(i) CRS technology. This feature of the model is designed to allow for long-run growth at the firm level, which is known to be important in explaining the power law and fat tail of the firm size distribution in the data (see, for example, Luttmer (2007, 2012)).

(ii) Constant relative risk aversion (CRRA) preferences. CRRA is a common assumption in quantitative work and is especially important in our setup because both firm size and CEO compensation in our model obey power law. CRRA allows for a realistic description of managers' risk aversion in the cross section and, hence, allows our model to account for the observed sensitivity of CEO pay to firm performance.[6]

(iii) Multiplicative impact of hidden actions. The manager's choice of investment rate, $i_{j,t}$, can be also interpreted as hidden effort in acquisition and accumulation of organization capital. The organization-capital accumulation in Equation (2) and the linearity of firms' profit function in Equation (1) imply that the impact of hidden actions, $i_{j,t}$, on firm size is multiplicative, which, as emphasized in Edmans, Gabaix, and Landier (2008), and Edmans and Gabaix (2016), is important for explaining the dependence of pay-performance sensitivity on size. In what follows, we will use the terminology "effort" and "investment in organization capital" interchangeably.

[6] While risk neutral or constant absolute risk aversion (CARA) preferences are theoretically convenient for their tractability, they generally lack to provide a realistic description of managers' risk aversion. For example, CARA implies that the Arrow-Pratt measure of relative risk aversion grows unboundedly as firm size and CEO pay increase. Hence, under the optimal contract with CARA preferences, the log-log measure of pay-performance sensitivity for large firms would be virtually zero, which is inconsistent with the data.

(iv) Presence of both observable and unobservable shocks. The inclusion of observable shocks is one of the key features that distinguish our paper from the existing literature. As we show below, it is important to explicitly to account for observable shocks, because in contrast to static models, in dynamic settings, they generally affect continuation utility. In fact, we show that observable shocks are critical to explain the empirical relationships between investment, pay-performance sensitivity and firm size.

(v) General equilibrium. To save space, in the main text, we focus on a single firm decision problem, and provide details of the market clearing conditions in general equilibrium in Appendix A of the paper. It is important to emphasize that the general equilibrium aspect of the model is essential in our estimation because it allows us to exploit the cross-sectional information to identify the structural parameters. It is also necessary for understanding the aggregate cost of moral hazard and its quantitative implications, which we evaluate using a counter-factual analysis in Section 4.

2.2 Profit Maximization

In our setup, high organization-capital investment accelerates firm growth and increases output. However, because it is not observable and managers have incentives to substitute investment in organization capital for consumption, shareholders' plan for organization-capital investment can be implemented only if managers find it incentive compatible to follow the plan. To induce effort and ensure that shareholders' and managers' interests are aligned, firm owners reward high output and punish low output. In a dynamic setting, these incentives are provided by conditioning future managerial compensation on past performance. Below, we formally describe the optimal contracting problem in our model.

A contract is a sequence of dividends, managerial compensation, and organization-capital investment policies, $\left(\left\{D_{j,t}\left(H_j^t, B_{O,j}^t\right)\right\}, \left\{C_{j,t}\left(H_j^t, B_{O,j}^t\right)\right\}, \left\{i_{j,t}\left(H_j^t, B_{O,j}^t\right)\right\}\right)_{t=0}^{\infty}$ that depends on the history of the realization of observables, which we denote by $H_j^t = \left\{H_{j,t}\right\}_{s=0}^{t}$ and $B_{O,j}^t = \left\{B_{O,j,s}\right\}_{s=0}^{t}$. To save notations, we write a contract as $\{D_t, C_t, i_t\}_{t=0}^{\infty}$. Given a contract $\{D_t, C_t, i_t\}_{t=0}^{\infty}$, if the manager follows the dividend payout policy, $\{D_t\}_{t=0}^{\infty}$, but chooses an alternative organization-capital investment policy, $\{\tilde{\imath}_t\}_{t=0}^{\infty}$, his continuation utility at time t can be written as:

$$U_t\left(\{\tilde{\imath}_s\}_{s=0}^{\infty}\right) = \left\{E_t\left[\int_0^{\infty} e^{-(\beta+\kappa)s}\left(\beta+\kappa\right)\left(C_{t+s} + \phi\left(i_{t+s}\right)H_{t+s} - \phi\left(\tilde{\imath}_{t+s}\right)H_{t+s}\right)^{1-\gamma} ds\right]\right\}^{\frac{1}{1-\gamma}}. \tag{6}$$

That is, at time t, if the manager chooses a lower effort $(\tilde{\imath}_t)$ than what is specified under the contract (i_t), he can privately save $\phi\left(i_{t+s}\right)H_{t+s} - \phi\left(\tilde{\imath}_{t+s}\right)H_{t+s}$ units of capital and use it for consumption without being detected by the shareholder. Therefore, a contract $\{D_t, C_t, i_t\}_{t=0}^{\infty}$ is incentive compatible if the investment policy specified by the contract is optimal from the manager's perspective, that is, if:

$$U_t\left(\{i_s\}_{s=0}^{\infty}\right) \geq U_t\left(\{\tilde{\imath}_s\}_{s=0}^{\infty}\right), \quad \text{for all } t, \tag{7}$$

for all investment policies $\{\tilde{\imath}_s\}_{s=0}^{\infty}$.

We assume double-sided limited commitment of financial contracts, as in Ai, Kiku, and Li (2015).[7] We assume that upon default, managers can take away a fraction of the firm's organization capital (for example, human capital) but are forever excluded from the credit market.[8] That is, managers are not allowed to enter into any intertemporal risk sharing contracts after default. As we show in Appendix B, the utility of the manager upon default is a linear function of organization capital: $u_{MIN}H_t$. Hence, limited commitment on the manager side requires

$$U_t\left(\{i_s\}_{s=0}^{\infty}\right) \geq u_{MIN}H_t, \quad \text{for all } t \geq 0. \tag{8}$$

The expression for u_{MIN} is provided in the appendix. Any compensation plan that violates condition (8) may lead to the manager defaulting on the contract.

We also assume that shareholders cannot commit to negative net present value (NPV) projects. This constraint requires that the net present value of firms' cash flow stays positive at all times:

$$E_t\left[\int_t^{\infty} e^{-(\mathbf{r}+\kappa)(s-t)} D_s ds\right] \geq 0, \quad \text{for all } t \geq 0. \tag{9}$$

Shareholders choose a contract, $\{D_t, C_t, i_t\}_{t=0}^{\infty}$ that maximizes the present value of firms' cash flow subject to the budget constraint in Equation (4), the incentive compatibility constraint, (7), and the limited commitment constraints in Equations (8) and (9).

2.3 Recursive Formulation

Following the standard approach in the dynamic contracting literature, we construct the solution to the optimal contracting problem recursively by using promised utility as a state variable. In our case, policy functions depend on two state variables (H, U), where H is the

[7] Some type of limited commitment is required to make the moral hazard problem non-trivial. In the absence of frictions, the principal can typically implement the efficient allocation arbitrarily closely if she is allowed to apply extremely severe punishment on the agent (see Mirrlees (1974)). In reality, managers typically have a variety of outside options and can always choose to leave a firm. We formally model outside options by limited commitment.

[8] Similar specifications are used in Albuquerque and Hopenhayn (2004), Kiyotaki and Moore (1997), and Kehoe and Levine (1993).

stock of organization capital and U is the continuation utility promised to the manager. We can think of the state variables, (H, U), as a summary of the firm's type. As in Atkeson and Lucas (1992), we construct the equilibrium allocation recursively. First, for firms of each type (H, U), we specify the flow rate of dividend payout, managerial compensation and effort level using the policy functions $D(H, U)$, $C(H, U)$, $i(H, U)$. Next, we specify the law of motion of the state variables using:

$$\frac{dH}{H} = [i(H, U) - \delta] dt + \sigma_O dB_O + \sigma_U dB_U, \tag{10}$$

and

$$\begin{aligned} \frac{dU}{U} &= \left[-\frac{\beta + \kappa}{1 - \gamma} \left(\left(\frac{C(H, U)}{U} \right)^{1-\gamma} - 1 \right) + \frac{1}{2} \gamma \left(g_O^2 \sigma_O^2 + g_U^2 \sigma_U^2 \right) \right] dt \\ &+ g_O \sigma_O dB_O + g_U \sigma_U dB_U. \end{aligned} \tag{11}$$

Equation (11) follows the formulation in Sannikov (2008) except that we use a stochastic differential utility representation of the preference so that utility is measured in consumption units (see Equation (6)). Here, g_O is the elasticity of continuation utility with respect to observable shocks, and g_U is the elasticity of U with respect to unobservable shocks. Intuitively, the policy functions g_U and g_O describe the rules of assigning continuation utilities based on the realizations of the Brownian shocks. At time t, for a given level of promised utility U_t, the principal allocates the manager's continuation utility over time and states by choosing an instantaneous consumption flow, $C(H_t, U_t)$, an elasticity of continuation utility with respect to unobservable shocks, $g_{U,t}$, and an elasticity with respect to observable shocks, $g_{O,t}$.

Because the production technology is constant return to scale, and utility functions are homogeneous, the optimal contracting problem is homogeneous in the state variable H. As a result, the value function, denoted by $V(H, U)$, satisfies

$$V(H, U) = v\left(\frac{U}{H}\right) H, \tag{12}$$

for some function v. We define $u = \frac{U}{H}$ as the normalized continuation utility and use homogeneity to write normalized consumption, investment, and dividend as functions of u:

$$c(u) = \frac{C(H, U)}{H}; \; i(u) = \frac{I(H, U)}{H}; \; d(u) = \frac{D(H, U)}{H}. \tag{13}$$

As shown in the appendix, the elasticities g_O and g_U are also functions of u.

Note that because H is firm size and U is the value of the manager's future compensation package, we can intuitively interpret u as the manager's share in the firm. In what follows, we use the terminology of the normalized utility and the manager's share in the firm interchangeably.

2.4 Entry, Exit, and Aggregation

A unit measure of managers arrives in the economy per unit of time with an outside option that provides them with a life-time utility U_0. Upon entry, a manager enters into a contract with shareholders and starts operating a firm of the initial size of H_0, which we normalize to one. Profit maximization implies that the initial normalized utility of the manager is $u_0 = \frac{U_0}{H_0}$. New managers arrive continuously and new firms are created upon their arrival and, at the same time, existing firms continuously exit the economy. We focus on the stationary equilibrium where entry equals exit, and the total measure of firms in the economy is constant.

3 Optimal Contracting

3.1 Characterization of the Optimal Contract

As we show in Lemma 1 in Appendix B, an investment policy is incentive compatible if and only if the normalized policy function g_U satisfies

$$g_U(u) = (\beta + \kappa) c(u)^{-\gamma} u^{\gamma - 1} \cdot \phi'(i(u)) \tag{14}$$

for all u. The above condition is intuitive. The term $(\beta + \kappa) c(u)^{-\gamma} u^{\gamma-1}$ is the marginal utility of the manager, and $\phi'(i(u))$ is the marginal cost of effort measured in units of consumption goods. Thus, the right-hand side of Equation (14) is the marginal cost of effort measured in manager's utility units. Because shareholders do not observe managers' consumption and effort decisions, they assign continuation utilities according to the realized H. Therefore, from the manager's perspective, $g_U(u)$ is the increase in continuation utility for an additional unit of capital. Incentive compatibility requires that the investment policy specified by the contract is optimal from the manager's perspective. In the context of our model, it requires the marginal benefit of an additional unit of effort, $g_U(u)$, be equal to the marginal cost of effort for the manager, $(\beta + \kappa) c(u)^{-\gamma} u^{\gamma-1} \cdot \phi'(i(u))$.

Condition (14) reduces the requirement of incentive compatibility to restrictions on the policy functions for consumption and effort, and the sensitivity of the continuation utility to unobservable shocks. This allows us to characterize the value function as the solution to a HJB equation, which we describe in the following proposition.

Proposition 1. *The normalized value function, $v(u)$, satisfies the following HJB differential equation*

$$0 = \max_{\substack{c,i,g_O,\\ g_U=(\beta+\gamma)c^{-\gamma}u^{\gamma-1}\phi'(i)}} \left\{ \begin{array}{c} A-c-\phi(i)+v(u)\left(i-r-\kappa-\delta\right) \\ +uv'(u)\left[\frac{\beta+\kappa}{1-\gamma}\left(1-\left(\frac{c}{u}\right)^{1-\gamma}\right)-(i-\delta)+\frac{1}{2}\gamma\left(g_U^2\sigma_U^2+g_O^2\sigma_O^2\right)\right] \\ +\frac{1}{2}u^2v''(u)\left[\left(g_U-1\right)^2\sigma_U^2+\left(g_O-1\right)^2\sigma_O^2\right] \end{array} \right\} \tag{15}$$

on the domain $[u_{MIN}, u_{MAX}]$, with the following boundary conditions:

$$\lim_{u\to u_{MIN}} v''(u) = \lim_{u\to u_{MAX}} v''(u) = \infty, \text{ and } v\left(u_{MAX}\right)=0.$$

Proof. See Appendix B.2. □

In Figure 1, we plot the normalized value function, $v(u)$, constructed using the estimates of the model parameters that we present in Section 4 below. Under the optimal contract, the normalized continuation utility of the agent, $u = \frac{U}{H}$, stays in the bounded interval, $[u_{MIN}, u_{MAX}]$. The limited commitment constraint in Equation (8) requires $u_t \geq u_{MIN}$ because any feasible contract must provide the manager a continuation utility at least as high as his outside option, $u_{MIN}H$. As u increases, the value function, $v(u)$, monotonically declines because a higher fraction of future cash flows is promised to the manager. Limited commitment on the shareholder side that requires the NPV of the firm to be non-negative at all times imposes an upper bound on u: u_{MAX} such that $v(u_{MAX}) = 0$ and $u_t \leq u_{MAX}$ for all t. As Figure 1 shows, the value function is concave on its domain.

The key tradeoff in moral-hazard models between risk sharing and incentive provision depends on the relative importance of unobservable versus observable shocks, $\frac{\sigma_U^2}{\sigma_O^2}$. Below, we show how the dynamics of continuation utility and managers' effort choices depend on $\frac{\sigma_U^2}{\sigma_O^2}$. These theoretical results allow us to identify the relative magnitude of observable and unobservable shocks by exploiting the cross-sectional distribution of CEO compensation and firm growth that we carry out in Section 4.

3.2 Dynamics of Continuation Utility

To understand the dynamics of continuation utility under the optimal contract, consider the elasticity of promised utility U with respect to observable and unobservable shocks, g_O and g_U, respectively. Figure 2 plots g_O (solid line) and g_U (dashed line) as functions of the key state variable, the normalized utility u. Note that the elasticity with respect to observable shocks is close to zero in most of its domain and converges to one at the boundaries. Due to optimal risk sharing, managers' continuation utility responds less to observable shocks than firms' output does, hence, $g_O(u) \leq 1$ for all u. As the normalized utility approaches its left boundary: $u \to u_{MIN}$, managers' continuation utility approaches its outside option: $U \to u_{MIN}H$. Because the elasticity of managers' outside option with respect to shocks in capital is one, the limited commitment constraint on the manager side, $U \geq u_{MIN}H$, requires that $g_O(u) = 1$ as $u \to u_{MIN}$. Similarly, $g_O(u)$ approaches one as the limited commitment constraint on the shareholder side starts to bind, because this constraint is equivalent to $U_t \leq u_{MAX}H_t$ and the elasticity of $u_{MAX}H_t$ with respect to shocks is one.

As Figure 2 further shows, the sensitivity of continuation utility with respect to unobservable shocks, g_U (dashed line), is positive and is substantially larger than g_O. The fact that $g_U > 0$ is the requirement of incentive provision. The optimal contract rewards high output and punishes low output. Because high output is more likely under high investment, this condition is necessary to deter managers from shirking. Similar to observable shocks, $g_U(u)$ must converge to one at the boundaries where limited commitment constraints bind.

Several important observations follow. First, under the optimal contract, continuation utility responds to observable shocks. Hence, in order to understand firms' policy decisions, it is important to explicitly incorporate observable shocks in a structural model. The fact that $g_O(u) > 0$ for all u is in sharp contrast with the sufficient statistics result in static models where reward does not respond to observable shocks (Hölmstrom (1979)). To understand the intuition notice that in our dynamic setting, utility should be allocated to equalize the marginal cost of utility provision across observable states, where the incentive compatibility constraint does not bind. Because the presence of agency frictions implies imperfect risk sharing, the cost of utility provision is lower in states with higher cash flow (the value function is strictly concave). It is, therefore, optimal to make continuation utility an increasing function of observable productivity shocks that have persistent effect on firms' future cash flow.

Second, $g_O(u) < 1$ in the interior of its domain and g_U is much larger than g_O (in fact, $g_U(u) \geq 1$ in most of its domain). Note that under perfect risk sharing, the elasticity of continuation utility with respect to shocks is zero. In case of observable shocks, limited commitment is the only friction that prevents perfect risk sharing. Because the elasticity of outside options with respect to shocks is one, risk sharing requires that $g_O(u) < 1$ unless the constraints are binding. Further, $g_U(u)$ is significantly larger than $g_O(u)$ and is typically larger than one because incentive provision makes the continuation utility highly sensitive to innovations in output.

Third, $g_U(u)$ is a decreasing function over most of its domain. Recall that the normalized utility, u, proxies for the manager's share in the firm. When u is low, the manager is promised only a small fraction of the firm's output and, therefore, has a strong incentive to shirk. To prevent shirking, the optimal contract is designed to provide a sufficiently high reward for good performance and, therefore, features high sensitivity to output shocks. As the manager's share in the firm increases, he owns a larger fraction of the firm's cash flows, which makes the incentive problem less severe and continuation utility less sensitive. The declining pattern in $g_U(u)$ illustrated in Figure 2 suggests that pay-performance sensitivity is decreasing in the normalized utility.

3.3 Investment in Organization Capital

As shown in Figure 3, our model implies that firm investment in organization capital, $i(u_t) = \frac{I_t}{H_t}$, is an increasing function of the manager's normalized utility. Limited commitment and moral hazard both contribute to the robust positive relationship between effort and the manager's share in the firm. When u approaches u_{MAX}, the limited commitment constraint on the shareholder side is likely to bind. A binding limited commitment constraint is associated with poor risk sharing. To grow out of the constraint and improve risk sharing, firms optimally increase their investment in organization capital as u increases. Further, as the normalized utility rises, more of the firm's cash flow is promised to the manager. Hence, managers' incentives become more aligned with shareholders' interests, which motivates them to exert more effort. In contrast, when the manager's share in the firm declines, it is harder to incentivize the manager to invest in organization capital. As a result, the cost of incentive provision increases and the optimal level of effort decreases. In fact, for u close to its lower boundary, investment in organization capital becomes negative.

It is well known that empirically, small firms invest at a higher rate and grow faster than large firms. In the next section, we show how the empirical relationship between growth rate and size, and the dependence of $i(u)$ on u implied by the model can be used to identify the relative magnitude of observable and unobservable shocks.

3.4 Implications for Identification

As discussed above, under the optimal contract, effort level and therefore firm growth rates increase with the normalized utility u, while pay-performance sensitivity decreases with u. These implications could potentially help identify the structural model parameters. However, the continuation utility is not directly observable. In this section, we show that we can identify the relative amount of private information and estimate the model using the readily available data on firm size by exploiting the relationship between firm size and continuation utility, which is endogenously determined by the optimal contract. In particular, we show that moral hazard determines the equilibrium correlation between firm size and the unobservable continuation utility. If the magnitude of unobservable shocks is relatively small, firm size and continuation utility are negatively correlated, whereas large values of $\frac{\sigma_U^2}{\sigma_O^2}$ imply a zero or even positive correlation between firm size and u. Because our model imposes monotone relationships between investment rate and u, and PPS and u, we can exploit the joint empirical distribution of growth rates, PPS and firm size to identify the model parameters.

Using Equations (10) and (11), we can write the law of motion of u as

$$\frac{du}{u} = \mu(u)dt + [g_O(u) - 1]\,\sigma_O dB_O + [g_U(u) - 1]\,\sigma_U dB_U, \tag{16}$$

where the function $\mu(u)$ is given in Equation (22) in the appendix. Note that $g_O - 1$ is the elasticity of u with respect to observable shocks, and $g_U - 1$ is the elasticity of u with respect unobservable shocks. As shown in Section 3.2, $g_O \leq 1$, g_U is significantly higher than g_O, and typically $g_U(u) \geq 1$. Consider first the case in which most of the shocks are observable, i.e., $\frac{\sigma_U^2}{\sigma_O^2}$ is close to 0. In this scenario, the relationship between firm size and the normalized continuation utility is negative because it is mostly driven by observable shocks and $g_O \leq 1$. Intuitively, while a positive shock increases both firm size H and continuation utility U, risk sharing requires U to increase at a lower rate than H. Hence, the manager's share u falls, and H and u are negatively correlated.

If the contribution of unobservable shocks increases, that is, if $\frac{\sigma_U^2}{\sigma_O^2}$ becomes large, the negative correlation between firm size and manager's share weakens because $g_U > g_O$. In fact, because $g_U \geq 1$ in most of its domain, the correlation between H and u might be even positive. As moral hazard becomes severe, the cost of incentive provision increases; therefore, as firms grow and become larger, they have to provide a higher fraction of their value to managers. We illustrate the relationship between firm size and the normalized utility in Figure 4. We consider several cases for the break-down between observable and unobservable shocks: $\frac{\sigma_U^2}{\sigma_O^2} = 0$ (dotted line), $\frac{\sigma_U^2}{\sigma_O^2} = 0.1$ (dashed line), $\frac{\sigma_U^2}{\sigma_O^2} = 1$ (dash-dotted line), and $\frac{\sigma_U^2}{\sigma_O^2} = 10$ (solid line). As the figure shows, when most shocks are observable, u is monotonically decreasing in firm size. As the relative magnitude of unobservable shocks increases, the negative relationship between u and H becomes considerably weaker and eventually changes its sign.

Figure 5 shows how the equilibrium relationship between firm size and the normalized utility translates into the relationship between size and investment in organization capital. For each of the four cases, we plot the effort level $i(u)$ as a function of firm size. For low levels of $\frac{\sigma_U^2}{\sigma_O^2}$ (0 and 0.1), our model features a strong inverse relationship between effort level and firm size, which is due to a strong negative correlation between the normalized utility and size. As the magnitude of unobservable shocks increases, the negative relationship between effort level and size disappears and ultimately reverses to positive. Although effort is not observable, firm growth rates are. The above discussion implies that the inverse relationship between firm growth and firm size in the data can be used to identify the magnitude of $\frac{\sigma_U^2}{\sigma_O^2}$.

Figure 6 shows the cross-sectional distribution of the average elasticity of continuation utility to productivity shocks. We compute average elasticity, ξ, as a weighted average of elasticities with respect to observable and unobservable shocks, $\xi = \sqrt{\frac{\sigma_U^2}{\sigma^2} g_U^2 + \frac{\sigma_O^2}{\sigma^2} g_O^2}$, where $\sigma^2 = \sigma_U^2 + \sigma_O^2$ is the total variance. Although the average elasticity of continuation utility with respect to shocks, ξ, is not observable, it determines pay-performance sensitivity (i.e., the elasticity of CEO pay to firm performance), which we can measure using the available data.[9] Figure 6 illustrates two important implications that help identify the relative magnitude of observable and unobservable shocks. First, as the amount of unobservable shocks increases, the overall elasticity of managerial pay to firm performance rises due to the higher cost of incentive provision. Second, without moral hazard, that is, when $\frac{\sigma_U^2}{\sigma_O^2} = 0$, ξ is a U-shaped function of size. A relatively modest amount of unobservable shocks (eg., $\frac{\sigma_U^2}{\sigma_O^2} = 0.1$) implies an increasing (almost monotone) relationship between ξ and firm size. As the contribution of unobservable shocks gets larger, the relationship between the elasticity of managerial pay and firm size virtually disappears. To summarize, Figure 6 shows that both the overall level and the cross-sectional variation of PPS provide important information about the relative magnitude of observable versus unobservable shocks.

Building on these theoretical insights, in the next section, we estimate the model by exploiting moment conditions of the joint distribution of firm size, firm growth rate and managerial compensation.

[9]The mapping between the elasticity of utility with respect to shocks, ξ, and the elasticity of managerial compensation with respect to firm performance (that is, the log-log based measure of PPS) is monotone, but nonlinear. For example, for $\frac{\sigma_U^2}{\sigma_O^2} = 0$, the log-log PPS is zero in the interior of (u_{MIN}, u_{MAX}) due to risk sharing. However, $\xi > 0$ because continuation utility accounts for a possibility of a binding constraint in the future, which is associated with a positive response of managerial pay with respect to shocks. Estimating the magnitude of $\frac{\sigma_U^2}{\sigma_O^2}$ from the level of PPS is thus a quantitative issue, which we address formally in the next section.

4 Quantitative Evidence

4.1 Data

We use the standard panel of data that consist of US non-financial firms and come from the Center for Research in Securities Prices (CRSP) and Compustat. For each firm in our sample, we collect its size, investment in physical capital, and managerial compensation. We measure firm size by the gross value of property, plant and equipment, and firm investment in physical capital by capital expenditure.[10] Executive compensation comprises salary, bonuses, the value of restricted stock and options granted, and long-term incentive payouts. All nominal quantities are converted to real using the consumer price index provided by the Bureau of Labor Statistics. The data are sampled on the annual frequency and cover the period from 1993 till 2013.

4.2 Structural Estimation

We estimate the model parameters using the simulated method of moments (McFadden (1989), and Pakes and Pollard (1989)). Our primary focus is on volatility parameters that govern the magnitude of observable and unobservable shocks, σ_O and σ_U, respectively. Following the discussion in Section 3, our identification strategy is to exploit the dynamics of firm growth and executive pay-performance sensitivity.

In addition to volatility parameters, we also estimate other structural parameters that are specific to our model. They include the equilibrium marginal product of capital (A), the organization-capital depreciation rate (δ), the adjustment cost parameters (ϕ_0 and i^*), and parameters that determine the outside option of managers and the initial normalized utility (u_{MIN} and u_0, respectively).[11] Several parameters that are common in standard macro models are fixed consistently with the literature. In particular, following Kydland and Prescott (1982), Rouwenhorst (1995), King and Rebelo (1999), and Atkeson and Kehoe (2005), we set the degree of risk aversion at 2, the time-discount rate (hence, the annual interest rate) at 0.04 per year, the span-of-control parameter to 0.85, and we calibrate the death rate to be 0.05 to match the average firms' exit rate in the data.

Let $\Theta = \{\sigma_O, \sigma_U, A, \delta, \phi_0, i^*, \bar{u}_{MIN}, \bar{u}_0\}$ denote the vector of parameters to estimate, and let $\mathcal{M}_D$ and $\mathcal{M}_M(\Theta)$ denote the vectors of data-based and model-implied moments,

[10] Recall that in our model, physical capital, organization capital and the number of employees are equivalent measures of firm size.

[11] It is more convenient to estimate the equilibrium marginal product of capital directly. We can always use the equilibrium relationship, $A = (1-\nu)\, z\mathbf{H}^{-\nu}$, where $\mathbf{H}$ denotes the aggregate stock of organization capital, to back out the primitive productivity parameter, z (see Appendix A).

respectively.[12] We estimate the model parameters by minimizing the following objective function:

$$\widehat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \big[\mathcal{M}_D - \mathcal{M}_M(\Theta)\big]' W \big[\mathcal{M}_D - \mathcal{M}_M(\Theta)\big], \tag{17}$$

where W is the weighting matrix. The model-based moments are computed via simulations. Specifically, we draw a panel of shocks, and for a given parameter configuration we solve the model numerically, discretize it and simulate a cross-section of firms. The simulated panel consists of 20,000 firms per year and the length of time series is set at 150 years. We discard the first 100 years of data and use the remaining 50 years of simulated data to calculate the vector of moments $\mathcal{M}_M(\Theta)$. Because our panel is fairly large, the simulated moments represent the population moments sufficiently well. We confirm that increasing either the length of the simulated sample or the size of the cross-section has virtually no effect on the model-implied moments and the parameter estimates. We estimate the model parameters using the optimal weight matrix. Hence, the estimation is carried out in two stages: in the first stage, we obtain the initial estimates by weighting each moment condition by the inverse of the variance of the sampling distribution of the corresponding statistic, and in the second stage we use the inverse of the variance-covariance matrix of the moment conditions evaluated at the first-stage estimates. We use the Newey and West (1987) estimator of the spectral density matrix at frequency zero with a truncation lag of two.

Guided by the model's implications discussed in Section 3, we select moments that are informative about the degree of moral hazard and help identify the structural parameters that we seek to estimate. The first set of moments characterizes aggregate and the cross-sectional dynamics of firms' growth. It consists of the average aggregate growth rate, the time-series average of the cross-sectional standard deviation of firms' growth rates, the time-series mean of median Tobin's Q, and average growth and physical capital investment rates of size-sorted portfolios. The second set of moments comprises CEO pay-performance sensitivity of firms with different size characteristics. The last moment is the power law in firm size. To calculate moments of the joint distribution of firm size, growth and managerial compensation, we construct five size-sorted portfolios using breakpoints that are equally-spaced in log size. Portfolios are re-balanced at the annual frequency, and for consistency, in the data and in the model, firm size is measured by the stock of organization capital.[13] The full list of moments that we exploit in estimation is presented in Table 2. In all, we use 15 moments to estimate

[12] For computational convenience, in estimation we use the following parameterization: $\bar{u}_{MIN} \equiv u_{MIN}/u_{MAX}$, and $\bar{u}_0 \equiv u_0/u_{MAX}$.

[13] Because in estimation we exploit the cross-sectional differences in the dynamics of CEO compensation, in constructing portfolios, we only use firms with available executive compensation data. Also, to ensure that the variance-covariance matrix of the moment conditions is well behaved, in estimation, we exploit growth and investment rate moments for only three out of five size-sorted portfolios; yet we present the model implications for the entire cross section.

eight model parameters.

It is important to note that we do not exploit investment in organization capital in estimation because it is not observable. Instead, we use the data on investment in physical capital that are readily available. To match the moments in the model and the data, we measure model-implied investment in physical capital as a depreciation-adjusted change in physical capital. While at a firm level, change in physical capital carries only a noisy signal of managerial effort (i.e., investment in organization capital), at a portfolio level, firm-specific shocks cancel out and the average investment rate of physical capital corresponds to the average investment rate of organization capital. Hence, average investment rates in physical capital and their dispersion across size-sorted portfolios contain information about structural model parameters that we exploit in estimation.

4.3 Parameter Estimates and Implications

Panel A of Table 1 presents the SMM estimates of the model parameters. First, notice that the set of moment conditions that we exploit in estimation allows us to identify the structural parameters sufficiently well; with only one exception, the estimates have relatively small standard errors. Second, our estimates reveal a significant difference in the magnitude of observable and unobservable shocks. The estimates of volatility of observable and unobservable shocks are 0.28 (SE=0.014) and 0.09 (SE=0.014), respectively. That is, observable shocks account for about 90% of the overall variation in productivity while the unobservable shocks contribute a much modest 10%. The difference in volatilities is strongly statistically significant with a robust t-statistic of 7.6.

To evaluate the fit of the model, in Table 2 we report sample moments alongside moments implied by the model estimates. For completeness, in addition to the moments exploited in estimation we report four additional moments that characterize growth and investment rates of firms in the second and fourth size-sorted quintiles (those are marked with asterisks). In the last column, we present t-statistics for the difference between the data and the model-implied moments. Note first that our estimated model matches well the dynamics at the aggregate level. The model implies an average growth rate of firm size of about 4.7% and a median Tobin's Q of about 1.4. Both of them are statistically similar to the corresponding sample statistics. Below we focus on the cross-sectional moments that allow us to identify the relative magnitude of observable and unobservable shocks.

Physical Investment, Growth Rates, and Firm Size As Table 2 shows, our model is able to account for the cross-sectional variation in firms' physical investment and growth observed in the data. As is well known, small firms, on average, invest more and grow at a much higher rate relative to large firms. Consistent with the data, our model generates a significant amount of dispersion in growth rates and investment rates across firms. In the model, the average investment rate declines monotonically from 19.1% for the bottom size-quintile to about 9.3% for the top size-sorted quintile. Similarly, the average growth rates vary from about 12.4% for the small-size cohort to 3.8% for large firms. The cross-sectional standard deviation of growth rates in the model is about 26.7%, which is statistically similar to 23% in the data. Although the model does not fully account for the very large physical investment and growth rates of the smallest size-quintile observed in the data, the difference between the model and the data is not statistically significant.

Limited commitment allows our model to account for the robustly negative relationship between firm size and firm growth observed in the data. Equally important is the finding that the amount of unobservable shocks is modest compared to the size of observable shocks. As we explained in Section 3.4, as long as the amount of unobservable shocks is relatively small, risk sharing dominates incentive provision, which makes managerial compensation less sensitive to shocks than productivity. As a result, the manager's share in the firm, u, is negatively correlated with firm size. Because managers in small firms have a claim to a larger share of firm's cash flow, they have stronger incentives to invest in organization capital than those in large firms. Hence, the negative relationship between size and u translates into a negative relationship between size and investment.

Pay-Performance Sensitivity The level and the cross-sectional variation of pay-performance sensitivity also play an important role in identifying the amount of private information. In the data and in the model, we measure PPS by regressing changes in log compensation on log returns controlling for firm fixed effects. As Table 2 shows, in the data, pay-performance sensitivity varies substantially with size — small firms feature relatively low PPS while large firms are characterized by relatively high size sensitivity of managerial compensation. The empirical estimate of PPS almost doubles from about 0.22 (SE=0.01) for the bottom quintile to 0.44 (SE=0.05) for the top size-sorted portfolio. The difference in PPS between the largest and smallest size cohorts of about 0.22 is statistically significant with a t-statistics of 2.24. Our model implies similar magnitudes of PPS and a similar monotonically increasing pattern in pay-performance sensitivity across size-sorted portfolios. The model-implied spread between the top and the bottom quintile is 0.22, which replicates the observed dispersion.

The ability of the model to account for both the level and the cross-sectional variation in PPS relies crucially on the presence yet a relatively small magnitude of unobservable shocks. At the point estimates, unobservable shocks account for about 10% of the total firm-level volatility. As explained in Section 3.4, such a relatively low magnitude of unobservable shocks implies a negative relationship between firm size and manager's normalized utility, u. Recall that the sensitivity of manager's utility to productivity shocks is decreasing in u. Taken together, the two implications lead to a positive relationship between pay-performance sensitivity and firm size and allow the model to simultaneously match the level and the cross-sectional pattern in PPS observed in the data.

As Table 2 shows, overall, the model accounts quite well for the joint distribution of firm size, growth and managerial compensation. None of 15 moment conditions exploited in estimation are statistically significant at the conventional five-percent level, and the model is not rejected by the over-identifying conditions. The p-value of the chi-square test of over-identifying restrictions is 0.41.

4.4 The Role of Moral Hazard

In order to better understand why the data implies a relatively small magnitude of unobservable shocks, we consider a restricted version of the model specification that assigns a larger role to moral hazard. In particular, we impose the constraint that the magnitudes of observable and unobservable shocks are equal, i.e., $\sigma_U^2 = \sigma_O^2 = 0.5\sigma^2$. Instead of evaluating the restriction directly (without re-estimating other model parameters), we give the constrained specification a fair chance to match the data and estimate the rest of its parameters by exploiting the same set of moment conditions. Table 3 presents the estimates of the restricted specification, and Table 4 reports its implications.

First notice that the constrained specification has significant difficulties in accounting for aggregate growth. As shown in Table 4, it implies a 3.1% aggregate growth, on average, which is much lower than in the data. But a more pronounced deterioration in the fit is in the cross-sectional dimension. Imposing the constraint significantly limits the model's ability to generate a sizable cross-sectional variation in growth rates and pay-performance sensitivity.

Our benchmark unrestricted model generates a spread of 8.9% in average growth rates of firms in the bottom and the top size quintiles, and a 9.8% dispersion in their physical investment rates. Under the constraint, the difference in average growth rates between the two cohorts of firms shrinks to 4.5%, and the spread in physical investments rates declines to 5.3%. These implications are quite intuitive. Keeping everything else constant, a larger magnitude of unobservable shocks and a higher degree of moral hazard reduce the overall investment and growth in the economy because a significant share of capital is allocated to incentive provision. Because small firms have limited resources, this happens to hurt small firms more relative to large firms. Therefore, investment and growth rates of small firms are substantially reduced and so is the variation in average growth rates across size-sorted portfolios. The constrained specification is able to generate only a 7.6% growth rate and a 15.3% physical investment rate of small firms, which are significantly lower than the corresponding statistics in the data. While other model parameters (such as productivity and the parameters of the adjustment cost function) try to adjust and compensate the negative impact of moral hazard on growth, their adjustment is bound by the discipline imposed by other moment conditions.

Further notice the impact of moral hazard on pay-performance sensitivity. Consistent with the data, in the unrestricted specification, PPS almost doubles from 0.20 for the small quintile to 0.42 for the large size quintile. Under the constraint, the model-implied dispersion shrinks significantly. Once again, deviations in PPS implied by the restricted model specification and the data are more pronounced in the left tail of the size distribution — the restricted specification significantly overstates size elasticity of CEO compensation for small firms. In particular, it implies a 0.28 PPS of the smallest portfolios, which is significantly larger the corresponding empirical estimate. Recall that in order to stay away from the shareholder-side limited commitment constraint and to improve risk sharing, small firm try to allocate as much resources as possible to invest and grow. Thus, in the absence of moral hazard, CEO compensation in small firms is sensitive to negative shocks but is inelastic to positive productivity shocks. Introducing moral hazard makes CEO compensation also respond to positive innovations and, therefore, magnifies pay-performance sensitivity of small firms. This is why the restricted specification fails to match a relatively low PPS of small firms observed in the data.

4.5 Counter-factual Exercises

In this section, we use our estimated model to quantify the impact of moral hazard on CEO compensation and aggregate output. We consider a counter-factual exercise of eliminating all moral hazard in the economy (for example, by implementing more transparent accounting rules and a more stringent legal system). To carry out our analysis, we keep all preference and technology parameters, including the total volatility of shocks, fixed at their estimated values but assume that all shocks are observable, that is, we set $\sigma_U = 0$. We compare the steady-state of the economy without moral hazard ($\sigma_U = 0$) with our benchmark economy with moral hazard.

Note that because we estimate the equilibrium marginal product of organization capital A directly, the span of control parameter ν does not affect our estimation. However, ν has to be factored in our analysis of the general equilibrium effect of moral hazard. Our counterfactual exercise varies the information parameter σ_U while keeping the fundamental technology parameter z constant. Eliminating moral hazard makes the accumulation of organization capital more efficient at the firm level and increases the total stock of organization capital at the aggregate level. This implies that the equilibrium marginal product of organization capital must decline due to decreasing return to scale at the aggregate level: $A = (1-\nu)\, z\mathbf{H}^{-\nu}$. Following Atkeson and Kehoe (2005), we set $\nu = 0.85$ in the counterfactual analysis below.

Impact of moral hazard on CEO compensation Moral hazard is often considered a key determinant of the level and the dynamics of CEO compensation. Our counter-factual exercise allows us to quantify the total fraction of managerial compensation that is attributed to incentive pay. To emphasize the dependence of policy functions on parameters, we denote the CEO compensation policy by $C\,(H, U|\,\Theta)$ and the steady-state distribution of the state variables by $\Phi\,(H, U|\,\Theta)$, where Θ represents the vector of parameter values of the model. We define λ_{CEOPAY} as the faction of incentive pay in total CEO compensation. Formally, we calculate the total amount of CEO compensation in the economy without moral hazard as a fraction of the total CEO pay in the economy with moral hazard and compute λ_{CEOPAY} as:

$$1 - \lambda_{CEOPAY} = \frac{\int C\,(H, U|\,\Theta_0)\,\Phi\,(dH, dU|\,\Theta_0)}{\int C\left(H, U|\,\hat{\Theta}\right)\Phi\left(dH, dU|\,\hat{\Theta}\right)}, \tag{18}$$

where $\hat{\Theta} = \{\hat{\sigma}_O,\, \hat{\sigma}_U,\, \hat{z},\, \hat{\delta},\, \hat{\phi},\, \hat{\imath}^*,\, \hat{\bar{u}}_{MIN},\, \hat{\bar{u}}_0\}$ is the estimated parameter vector, and Θ_0 is obtained from $\hat{\Theta}$ by setting $\sigma_U = 0$ and keeping all other parameters, including the total volatility of shocks, $\sqrt{\hat{\sigma}_U^2 + \hat{\sigma}_O^2}$, the same.

Our estimates imply that $\lambda_{CEOPAY} = 52.4\%$. That is, moral hazard accounts for about half of the overall CEO compensation. In other words, eliminating all moral hazard allows firms to save about 52% of CEO compensation while keeping managers' utility unchanged. In our model, eliminating moral hazard makes managerial compensation contract more efficient for two reasons. First, in the presence of moral hazard, incentive provision requires CEO compensation to respond to unobservable idiosyncratic shocks. This arrangement reduces welfare because managers are risk averse. In our model with moral hazard, shocks to firm output induce variation in CEO compensation of about 20% per year, whereas perfect risk sharing implies that CEO pay is constant over time. Note that the fact that moral hazard distorts the allocation of consumption across states of the world and limits risk sharing is true in both static and dynamic models.

Second, unique to our dynamic model, moral hazard also distorts the intertemporal allocation of managerial compensation. Because managers and shareholders have the same discount rate, efficiency requires CEO compensation to be constant over time. However, consistent with the previous literature (eg., Sannikov (2008), DeMarzo and Sannikov (2006)), the presence of moral hazard typically implies a back-loaded CEO compensation package. Intuitively, delayed compensation allows shareholders to condition future compensation on realized output to provide proper incentives to invest. In our model with moral hazard, the average growth rate of CEO compensation relative to the aggregate economy is about 1.8% per year, whereas it is zero in the model without moral hazard.

To decompose the overall impact of moral hazard into its effect on risk sharing and its impact on the intertemporal allocation of consumption, let $\bar{C}(H,U)$ denote CEO compensation policy implied by the optimal contract in the moral hazard model that has no idiosyncratic risk (that is, by setting the diffusion coefficient of $C\left(H,U|\,\hat{\Theta}\right)$ to zero but keeping the drift). Also, let $\bar{\Phi}(H,U)$ be the stationary distribution of CEO compensation under the policy $\bar{C}(H,U)$. Then, the total impact of moral hazard can be decomposed as follows:

$$[1-\lambda_{CEOPAY}] = \frac{\int \bar{C}(H,U)\,\bar{\Phi}(dH,dU)}{\int C\left(H,U|\,\hat{\Theta}\right)\Phi\left(dH,dU|\,\hat{\Theta}\right)} \times \frac{\int C(H,U|\,\Theta_0)\,\Phi(dH,dU|\,\Theta_0)}{\int \bar{C}(H,U)\,\bar{\Phi}(dH,dU)}. \quad (19)$$

The first term of the product on the right-hand side measures the efficiency loss due to limiting risk sharing, and the second component measures the efficiency loss due to distortions in the intertemporal allocation of CEO compensation. We find that $1-\frac{\int \bar{C}(H,U)\bar{\Phi}(dH,dU)}{\int C\left(H,U|\hat{\Theta}\right)\Phi\left(dH,dU|\hat{\Theta}\right)} = 39\%$ and $1-\frac{\int C(H,U|\Theta_0)\Phi(dH,dU|\Theta_0)}{\int \bar{C}(H,U)\bar{\Phi}(dH,dU)} = 22\%$. Thus, our estimates imply that 52.4% of CEO compensation is the result of incentive provision. About two thirds of the incentive pay is compensation for the additional risk managers have to bear and one third of it is compensation for distortions in the intertemporal allocation of their compensation packages. Our decomposition highlights the importance of a dynamic setting in estimating efficiency losses — static models would miss distortions in the intertemporal allocation, which we find to be a qualitatively significant part of incentive compensation.

Impact of moral hazard on aggregate output Our general equilibrium framework also allows us to quantify the impact of moral hazard on aggregate output. We compute it as a percentage increase in output that can be achieved by eliminating all moral hazard in the economy:

$$\lambda_{OUTPUT} = \frac{\int zH^{1-\nu}\left[K(H,U|\,\Theta_0)^{\alpha}\,N(H,U|\,\Theta_0)^{1-\alpha}\right]^{\nu} d\Phi(dH,dU|\,\Theta_0)}{\int zH^{1-\nu}\left[K\left(H,U|\,\hat{\Theta}\right)^{\alpha} N\left(H,U|\,\hat{\Theta}\right)^{1-\alpha}\right]^{\nu}\Phi\left(dH,dU|\,\hat{\Theta}\right)} - 1.$$

Under the estimated parameter values of the model, we find that $\lambda_{OUTPUT} = 1.44\%$. That is, eliminating all moral hazard results in a permanent increase in aggregate output of the economy by about 1.4%.

Note the importance of general equilibrium for our welfare analysis. Taking as given the equilibrium level of marginal product of organization capital, A, a reduction in moral hazard improves the efficiency of production at a firm level. This could potentially lead to an infinitely large increase in firm value because revenue is linear in A. However, a lower degree of moral hazard implies that firms invest more and organization capital accumulates faster. Hence, the equilibrium level of marginal product must fall as the economy accumulates more organization capital. Thus, the welfare impact of moral hazard is determined by the return to scale of the technology with respect to organization capital at the aggregate level that our general equilibrium model properly takes into account.

5 Conclusion

We quantify the impact of moral hazard using a structural estimation of a dynamic general equilibrium model with agency frictions. The degree of moral hazard is defined by the relative magnitude of unobservable versus observable productivity shocks. We show that moral hazard has important implications for the cross-sectional relationships between firm size and firm growth rate, and firm size and pay-performance sensitivity. We exploit the predictions of our model and identify the amount of observable and unobservable shocks by exploiting moment conditions of the joint empirical distribution of firm size, growth and PPS. We find that the magnitude of unobservable shocks is relatively small and accounts for about 10% of the total variation. Our estimates imply that moral-hazard induced incentives are quantitatively significant and explain 52% of managerial compensation. Our welfare analysis suggests that eliminating moral hazard results in about 1.4% increase in aggregate output.

Appendix (For Online Publication)

A Details of General Equilibrium

Here we provide the details of aggregation and general equilibrium in our model.

A.1 The product market

The optimality condition for the choice of labor for (1) implies

$$\nu\alpha z\left(\frac{H_j}{K_j^{\alpha}N_j^{1-\alpha}}\right)^{1-\nu}\left(\frac{K_j}{N_j}\right)^{\alpha-1} = R; \tag{20}$$

$$\nu(1-\alpha)z\left(\frac{H_j}{K_j^{\alpha}N_j^{1-\alpha}}\right)^{1-\nu}\left(\frac{K_j}{N_j}\right)^{\alpha} = W. \tag{21}$$

By combining (20) and (21), we have

$$\alpha^{\alpha}\left(1-\alpha\right)^{1-\alpha}\nu z H^{1-\nu} = R^{\alpha}W^{1-\alpha}$$

with $\mathbf{H} \equiv \frac{H_j}{K_j^{\alpha}N_j^{1-\alpha}}$ which is constant across all firms. Furthermore, $\mathbf{K} \equiv \frac{K_j}{N_j}$ is constant across all firms as well according to (20). Therefore, $R = \nu\alpha z\mathbf{H}^{1-\nu}\mathbf{K}^{\alpha-1}$ and $W = \nu\left(1-\alpha\right)z\mathbf{H}^{1-\nu}\mathbf{K}^{\alpha}$. We can therefore compute the profit function in (1) as

$$\pi\left(H_j\right) = z\left(1-\nu\right)\mathbf{H}^{-\nu}H_j$$

In our notion, this means $A = z\left(1-\nu\right)\mathbf{H}^{-\nu}$ is the equilibrium marginal product of capital.

A.2 Aggregation

The value and policy functions of the optimal contracting problem satisfy the homogeneity property (12) and (13). This implies that aggregate quantities can be computed by using the following "summary measure":

$$m\left(u\right) = \int H\Phi\left(H,u\right)dH,$$

where $\Phi(H,u)$ is the joint distribution of normalized utility and firm size. For example total managerial compensation can be calculated as

$$\int C(H,uH)\Phi(H,u)\,dH = \int Hc(u)\Phi(u,H)\,dH = \int c(u)\phi(u)\,du.$$

Also, because output at the firm level satisfies $y_j = \frac{1}{\alpha}AH_j$ by the result from the last section, total output of the economy can be calculated as

$$\int \frac{1}{1-\nu}AH_j\Phi(H,u)\,dH = \frac{1}{1-\nu}A\int \phi(u)\,du.$$

Notice that the law of motion of the normalized continuation utility satisfies

$$du = u\mu(u)dt + u\left[g_O(u)-1\right]\sigma_O dB_O + u\left[g_U(u)-1\right]\sigma_U dB_U$$

with

$$\mu(u) = \begin{bmatrix} -\frac{\beta+\kappa}{1-\gamma}\left(\left(\frac{c(u)}{u}\right)-1\right)+\frac{1}{2}\gamma\left(g_O^2(u)\sigma_O^2+g_U^2(u)\sigma_U^2\right) \\ +\left(g_O(u)-1\right)\sigma_O^2+\left(g_U(u)-1\right)\sigma_U^2 \end{bmatrix}. \tag{22}$$

Therefore, numerically, we can compute the summary measure $m(u)$ by solving the following forward equation:

$$\begin{aligned} 0 \;=\;& \bar{m}(u)-(\kappa-(i(u)-\delta))\,m(u)-\frac{d}{du}\left[m(u)\mu_u(u)+u\left(\left(g_O(u)-1\right)\sigma_O^2+\left(g_U(u)-1\right)\sigma_U^2\right)\right] \\ &+\frac{1}{2}\frac{d^2}{du^2}\left(m(u)u^2\left(\left(g_O(u)-1\right)^2\sigma_O^2+\left(g_U(u)-1\right)^2\sigma_U^2\right)\right). \end{aligned}$$

Here $\bar{m}(u) = \int H\bar{\Phi}(H,u)\,dH$ with $\bar{\Phi}(H,u)$ being the entrance measure of new firms.

B Optimal Contracting

B.1 Incentive Compatibility

Lemma 1. [Incentive compatibility]*A contract constructed form the allocation rule, $C(H,U)$, $I(H,U)$, $D(H,U)$, $N(H,U)$, $g_U(H,U)$, $g_O(H,U)$ satisfies the obedience constraint* (7) *if and only if for all $t\in[0,\infty)$*

$$\frac{g_U(H_t,U_t)}{U_t^{\gamma-1}H_t} = (\beta+\kappa)\,C(H_t,U_t)^{-\gamma}H_I\left(I(H_t,U_t),H_t\right) \tag{23}$$

or in normalized terms

$$g_U(u_t) = (\beta+\kappa)\, c_t^{-\gamma} u_t^{\gamma-1} \phi'(i_t)\,. \tag{24}$$

Lemma 2. *We show that the obedience constraint* (7) *is satisfied if and only if for all* $t \in [0,\infty)$

$$(C(H_t,U_t), I(H_t,U_t)) \in \arg \max_{\substack{C,I \ s.t \\ C+\phi(I,H_t)=AH_t-D(H_t,U_t)}} \frac{\beta+\kappa}{1-\gamma} C^{1-\gamma} + \frac{1}{U_t^{\gamma-1} H_t} g_U(H_t,U_t)\, I \tag{25}$$

which, along with concavity of $\phi(I,H)$*, implies* (23) *and* (24)*.*

To simply notation, we focus on the case $\sigma_O = 0$, that is, all shocks are unobservable. We also omit the arguments H and U in the policy functions in the statement of the lemma and define

$$W_t = \frac{1}{1-\gamma} U_t^{1-\gamma} = E_t\left[\int_t^\tau e^{-\beta(s-t)} (\beta+\kappa)\, C_s^{1-\gamma} ds\right], \tag{26}$$

which is the expected-utility representation of the agent's continuation utility. The Martingale representation theorem implies

$$dW_t = (\beta+\kappa)\left[W_t - \frac{C_t^{1-\gamma}}{1-\gamma}\right] dt + G_U(H_t,U_t)\, \sigma_U dB_{U,t} \tag{27}$$

with

$$G_U(H_t,U_t) = \frac{g_U(H_t,U_t)}{U_t^{\gamma-1}} \tag{28}$$

and $g_U(H_t,U_t)$ being the sensitivity term in Equation (11). Suppose that $\left\{\tilde{C}_t, \tilde{I}_t\right\}_{t=0}^{\infty}$ is an alternative consumption and investment plan other than $\{C_t, I_t\}_{t=0}^{\infty}$, such that

$$\tilde{C}_t + \phi\left(\tilde{I}_t, H_t\right) = C_t + \phi(I_t, H_t) = AH_t - D_t$$

We define

$$\mathcal{G}_t^{\tilde{C},\tilde{I}} = \int_0^t e^{-(\beta+\kappa)s} (\beta+\kappa) \frac{1}{1-\gamma} \tilde{C}_s^{1-\gamma} ds + e^{-(\beta+\kappa)t} W_t.$$

So $\mathcal{G}_t^{\tilde{C},\tilde{I}}$ is the time-t conditional expected utility of the agent's life-time utility if he follows plan $\left\{\tilde{C}_t, \tilde{I}_t\right\}_{t=0}^{\infty}$ over $[0,t]$ and switches to $\{C_t, I_t\}_{t=0}^{\infty}$ at t. Obviously, $\mathcal{G}_0^{\tilde{C},\tilde{I}} = W_0$ and

$$e^{(\beta+\kappa)t} d\mathcal{G}_t^{\tilde{C},\tilde{I}} = (\beta+\kappa) \frac{1}{1-\gamma} \tilde{C}_t^{1-\gamma} dt - (\beta+\kappa) W_t dt + dW_t. \tag{29}$$

Let $\left\{B_{U,t}^{C,I}\right\}_{t=0}^{\infty}$ and $\left\{B_{U,t}^{\tilde{C},\tilde{I}}\right\}_{t=0}^{\infty}$ be the Itô's processes which are standard Brownian motion under the probability measures induced by $\{C_t, I_t\}_{t=0}^{\infty}$ and $\left\{\tilde{C}_t, \tilde{I}_t\right\}_{t=0}^{\infty}$ respectively according to Girsanov theorem. Then (10) implies

$$\sigma_U dB_{U,t}^{C,I} = \sigma_U dB_{U,t}^{\tilde{C},\tilde{I}} + \frac{1}{H_t}\left(\tilde{I}_t - I_t\right) dt. \tag{30}$$

Then (29) and (27) implies

$$\begin{aligned} e^{(\beta+\kappa)t} d\mathcal{G}_t^{\tilde{C},\tilde{I}} &= \frac{\beta+\kappa}{1-\gamma}\left(\tilde{C}_t^{1-\gamma} - C_t^{1-\gamma}\right) dt + G_U\left(H_t, U_t\right)\sigma_U dB_{U,t}^{C,I} \\ &= \left[\frac{\beta+\kappa}{1-\gamma}\left(\tilde{C}_t^{1-\gamma} - C_t^{1-\gamma}\right) + G_U\left(H_t, U_t\right)\frac{1}{H_t}\left(\tilde{I}_t - I_t\right)\right] dt \\ &\quad + G_U\left(H_t, U_t\right) dB_{U,t}^{\tilde{C},\tilde{I}}. \end{aligned} \tag{31}$$

The last equality above is due to (30). Suppose that (25) is not satisfied and, according to (28),

$$\frac{\beta+\kappa}{1-\gamma}C_t^{1-\gamma} + \frac{1}{H_t}G_U\left(H_t, U_t\right) I_t < \frac{\beta+\kappa}{1-\gamma}\tilde{C}_t^{1-\gamma} + \frac{1}{H_t}G_U\left(H_t, U_t\right)\tilde{I}_t$$

over a time interval with a positive measure. Then there exists a $\bar{t}$ such that $\left\{\mathcal{G}_t^{\tilde{C},\tilde{I}}\right\}$ is a sub-martingale over $[0, \bar{t}]$ and

$$W_0 = \mathcal{G}_0^{\tilde{C},\tilde{I}} < E_0\left[\mathcal{G}_{\bar{t}}^{\tilde{C},\tilde{I}}\right]$$

so that $\{C_t, I_t\}_{t=0}^{\infty}$ is dominated by following $\{C_t, I_t\}_{t=0}^{\infty}$ from the beginning and switching to $\{C_t, I_t\}_{t=0}^{\infty}$ at $\bar{t}$. Conversely, if (25) is satisfied for all t, there is no such profitable deviation for all $t \in [0, \infty)$. Proof of (23) is straightforward.

B.2 Proof of Proposition 1

According to the definition of $V(H, U)$, the laws of motion of H and U, (10) and (11), when neither limited commitment constraint is binding, $V(H, U)$ satisfies the following HJB differential equation.

$$(\mathbf{r}+\kappa) V(H, U) = \max_{c,i,g_U,g_O}\left\{\begin{array}{c} AH - \phi(i) H - cH + V_H(H, U) H[i-\delta] \\ +\frac{1}{2}V_{HH}(H, U) H^2\left(\sigma_U^2 + \sigma_O^2\right) \\ +V_U(H, U) U\left[\frac{r+\kappa}{1-\gamma}\left(1 - \left(\frac{C}{U}\right)^{1-\gamma}\right) + \frac{1}{2}\gamma\left(g_U^2\sigma_U^2 + g_O^2\sigma_O^2\right)\right] \\ +\frac{1}{2}V_{UU}(H, U) U^2\left(g_U^2\sigma_U^2 + g_O^2\sigma_O^2\right) \\ +V_{H,U}(H, U) HU\left(g_U\sigma_U^2 + g_O\sigma_O^2\right) \end{array}\right\}. \tag{32}$$

Notice that the incentive constraint (24) implies the following restriction on the maximization problem on the right hand side of (32).

$$g_U = (\beta + \gamma)\, c^{-\gamma} u^{\gamma-1} \phi'(i). \tag{33}$$

Furthermore, according to the normalization, we have $V_H(H,U) = v(u) - uv'(u)$, $V_U(H,U) = v'(u)$, $V_{HH}(H,U) = \frac{1}{H} u^2 v''(u)$, $V_{UU}(H,U) = \frac{1}{Hv''(u)}$, and $V_{HU}(H,U) = -\frac{1}{H} uv''(u)$. Therefore we have (15).

The boundary conditions can be shown by following the argument in the proof of Lemma 1 in Ai and Li (2015).

References

Ai, Hengjie, Dana Kiku, and Rui Li, 2015, A Mechanism Design Model of Firm Dynamics: The Case of Limited Commitment, *Working paper, University of Minnesota*.

Ai, Hengjie, and Rui Li, 2015, Investment and CEO Compensation under Limited Commitment, *Journal of Financial Economics.* 116, 452–472.

Albuquerque, Rui, and Hugo Hopenhayn, 2004, Optimal Lending Contracts and Firm Dynamics, *Review of Economic Studies* 71, 285–315.

Atkeson, Andrew, and Patrick J. Kehoe, 2005, Modeling and Measuring Intangible Capital, *The Journal of Political Economy* 113.

Atkeson, Andrew, and Robert Lucas, 1992, On Efficient Distribution with Private Information, *Review of Economic Studies* 59, 427–453.

Becker, Gary S., 1993, *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. (Chicago University Press, 3rd Edition).

Biais, Bruno, Thomas Mariotti, Cuillaume Plantin, and Jean-Charles Rochet, 2007, Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications, *The Review of Economic Studies* 74, 345–390.

Biais, Bruno, Thomas Mariotti, and Jean-Charles Rochet Stéphane Villeneuve, 2010, Large Risks, Limited Liability, and Dynamic Moral Hazard, *Econometrica* 78, 73–118.

Bond, Philip, and Ulf Axelson, 2015, Wall Street Occupations, *Journal of Finance* 70(5), 1949–1996.

DeMarzo, Peter, Michael Fishman, Zhiguo He, and Neng Wang, 2012, Dynamic Agency and q Theory of Investment, *Journal of Finance, forthcoming* 67, 2295–2340.

DeMarzo, Peter, and Yuliy Sannikov, 2006, Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model, *Journal of Finance* 61, 2681–2724.

Di Tella, Sebastian, and Yuliy Sannikov, 2016, Optimal asset management contracts with hidden savings, *Working paper, Stanford University.*

Diais, Bruno, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet, 2004, Dynamic Security Design, *Working Paper, IDEI*.

Edmans, Alex, and Xavier Gabaix, 2016, Executive Compensation: A Modern Primer, *Journal of Economic Literature* forthcoming.

Edmans, Alex, Xavier Gabaix, and Augustin Landier, 2008, A Multiplicative Model of Optimal CEO Incentives in Market Equilibrium, *The Review of Financial Studies* 22.

Edmans, Alex, Xavier Gabaix, Tomasz Sadzik, and Yuliy Sannikov, 2012, Dynamic CEO Compensation, *Journal of Finance* 67.

Eisfeldt, Andrea L., and Dimitris Papanikolaou, 2013, Organization Capital and the Cross-Section of Expected Returns, *Journal of Finance* 68, 1365–1406.

Fudenberg, Drew, Bengt Holmstrom, and Paul Milgrom, 1990, Short-Term Contracts and Long-Term Agency Relationships, *Journal of Economic Theory* 51, 1–31.

Gayle, George-Levi, and Rober A. Miller, 2009, Has Moral Hazard Become a More Important Factor in Managerial Compensation?, *American Economic Review* 99.

Gayle, George-Levi, and Rober A. Miller, 2015, Identifying and Testing Models of Managerial Compensation, *Review of Economic Studies* 82, 1074–1118.

Hayashi, Fumio, 1982, Tobin's Marginal q and Average q: A Neoclassical Interpretation, *Econometrica* 50, 213–224.

He, Zhiguo, 2009, Optimal Executive Compensation when Firm Size Follows Geometric Brownian Motion, *Review of Financial Studies* 22, 859–892.

Hoffmann, Florian, and Sebastian Pfeil, 2010, Reward for Luck in a Dynamic Agency Model, *Review of Financial Studies* 23(9), 3329–3345.

Hölmstrom, B., 1979, Moral Hazard and Observability, *The Bell Journal of Economics* 10, 74–91.

Kehoe, Timothy, and David Levine, 1993, Debt-Constrained Asset Markets, *Review of Economic Studies* 60, 865–888.

King, Robert G., and Sergio T. Rebelo, 1999, Resuscitating Real Business Cycles, in *Handbook of Macroeconomics* (Elsevier,).

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit Cycles, *Journal of Political Economy* 105, 211–284.

Kydland, Finn, and Edward Prescott, 1982, Time to Build and Aggregate Fluctuations, *Econometrica* 50, 1345–1370.

Li, Rui, 2015, Dynamic Agency with Persistent Exogenous Shocks, *Working paper, University of Massachusetts Boston*.

Li, Shaojin, Toni Whited, and Yufeng Wu, 2016, Collateral, Taxes, and Leverage, *Review of Financial Studies*.

Luttmer, Erzo, 2007, Selection, Growth, and the Size Distribution of Firms, *Quarterly Journal of Economics* 122, 1103–1144.

Luttmer, Erzo, 2011, On the Mechanics of Firm Growth, *The Review of Economic Studies* 78, 1042–1068.

Luttmer, Erzo, 2012, Technology diffusion and growth, *Journal of Economic Theory* 147, 602–622.

Margiotta, M. M., and R. A. Miller, 2000, Managerial Compensation and the Cost of Moral Hazard, *International Economic Review* 41, 669–719.

McFadden, Daniel, 1989, A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration, *Econometrica* 57, 995–1026.

Mirrlees, James, 1974, *Notes on Welfare Economics, Information, and Uncertainty,* American Elsevier Pub. Co, pp. 243–258 in M. Balch, D. McFadden and S. Wu (Eds), Essays in Economic Behavior under Uncertainty edn.

Newey, Whitney, and Kenneth West, 1987, A Simple, Positive Semi-Definite, Heteroscedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703–708.

Nikolov, Boris, and Lukas Schmid, 2016, Testing Dynamic Agency Theory via Structural Estimation, *Working paper, Duke University.*

Pakes, Ariel, and David Pollard, 1989, Simulation and the Asymptotics of Optimization Estimators, *Econometrica* 57, 1027–1057.

Prescott, Edward C., and Michael Visscher, 1980, Organization Capital, *The Journal of Political Economy* 88.

Rouwenhorst, K. Geert, 1995, *Asset Pricing Implications of Equilibrium Business Cycle Models.* (Princeton University Press).

Sannikov, Yuliy, 2008, A Continuous-Time Version of the Principal-Agent Problem, *Review of Economic Studies* 75, 957–984.

Taylor, Lucian A., 2010, Why Are CEOs Rarely Fired? Evidence from Structural Estimation, *The Journal of Finance* 65, 2051–2087.

Taylor, Lucian A., 2013, CEO Wage Dynamics: Estimates from a Learning Model, *Journal of Financial Economics* 108, 79–98.

Zhu, John, 2013, Optimal Contracts with Shirking, *Review of Economic Studies* 80, 812–839.

Table 1

Model Parameters

Panel A: Estimated Parameters

Parameter	Estimate	
σ_O	0.2753	(0.0139)
σ_U	0.0909	(0.0135)
A	0.1691	(0.0021)
δ	0.0570	(0.0022)
ϕ_0	0.9684	(0.0943)
i^*	0.0019	(0.0001)
$\bar{u}_{MIN}$	0.0232	(0.0029)
$\bar{u}_0$	0.2257	(0.3225)

Panel B: Calibrated Parameters

Parameter	Value
γ	2
β	0.04
κ	0.05
ν	0.85

Table 1 presents the estimates of the model parameters and their standard errors in parentheses (Panel A) and the values of the calibrated parameters (Panel B). The estimated parameters include volatility of observable and unobservable shocks (σ_O and σ_U, respectively), the equilibrium marginal product of organization capital (A), the depreciation rate of organization capital (δ), the adjustment cost parameters (ϕ_0 and i^*), and parameters that determine the outside option of managers and the initial normalized utility ($\bar{u}_{MIN}$ and $\bar{u}_0$, respectively). The set of calibrated parameters consists of risk aversion (γ), the discount rate (β), the death rate of managers (κ), and the span of control parameter, ν.

Table 2
Sample and Model-Implied Moments

Moments		Data	Model	t-stat(Diff)
Growth:	Aggregate	0.048	0.047	0.10
	P1 (Small)	0.194	0.124	1.70
	P2*	0.108	0.082	1.26
	P3	0.071	0.059	1.30
	P4*	0.053	0.044	1.35
	P5 (Large)	0.032	0.038	−1.02
CS-Std of Growth Rates		0.231	0.267	−1.55
Tobin's Q		1.446	1.401	0.84
Inv-Rate:	P1 (Small)	0.251	0.191	1.64
	P2*	0.160	0.144	1.10
	P3	0.120	0.120	0.01
	P4*	0.101	0.102	−0.24
	P5 (Large)	0.092	0.093	−0.11
PPS:	P1 (Small)	0.215	0.202	0.58
	P2	0.234	0.193	1.48
	P3	0.240	0.279	−0.86
	P4	0.309	0.376	−1.84
	P5 (Large)	0.443	0.422	0.22
Power Law		1.180	1.198	−0.69

Table 2 shows the set of moments exploited in estimation. We report moments in the data and in the model, and t-statistics for the difference between sample- and model-implied moments. The set of moments consists of average growth rates (Growth) in the aggregate economy and across five size-sorted portfolios (P1–P5), the cross-sectional standard deviation of firms' growth rates, median Tobin's Q, average investment rates in physical capital (Inv-Rate), and executive pay-performance sensitivity (PPS) for the cross section of size-sorted portfolios, and the power law in firm size. Growth rates are measured in logs, PPS is measured in a panel regression of log change in CEO compensation on log firm return, controlling for firm fixed effects, and the power law is estimated using firms in the top size-sorted portfolio. Moments with asterisks are not exploited in estimation. The data are annual, measured in real terms and cover the period from 1993 to 2013.

Table 3

Estimates of the Restricted Model Specification

Parameter	Estimate
σ	0.3846 (0.0088)
A	0.2042 (0.0022)
δ	0.0701 (0.0015)
ϕ_0	2.5611 (0.1422)
i^*	0.0004 (0.0001)
$\bar{u}_{MIN}$	0.0176 (0.0043)
$\bar{u}_0$	0.5002 (0.1144)
σ_O^2	$0.5\sigma^2$
σ_U^2	$0.5\sigma^2$

Table 3 presents the estimates of the restricted model specification and their standard errors in parentheses. The estimated parameters include total volatility (σ), the equilibrium marginal product of organization capital (A), the depreciation rate of organization capital (δ), the adjustment cost parameters (ϕ_0 and i^*), and parameters that determine the outside option of managers and the initial normalized utility ($\bar{u}_{MIN}$ and $\bar{u}_0$, respectively). The bottom rows show the constrains imposed on volatilities of observable and unobservable shocks (σ_O and σ_U, respectively).

Table 4
Model-Implied Moments of the Restricted Specification

Moments		Data	Model	t-stat(Diff)
Growth:	Aggregate	0.048	0.031	2.26
	P1 (Small)	0.194	0.076	2.16
	P2*	0.108	0.080	1.31
	P3	0.071	0.044	2.09
	P4*	0.053	0.025	2.30
	P5 (Large)	0.032	0.032	0.14
CS-Std of Growth Rates		0.231	0.350	−2.44
Tobin's Q		1.446	1.512	−1.15
Inv-Rate:	P1 (Small)	0.251	0.153	2.09
	P2*	0.160	0.154	0.48
	P3	0.120	0.117	0.33
	P4*	0.101	0.095	1.21
	P5 (Large)	0.092	0.100	−1.73
PPS:	P1 (Small)	0.215	0.278	−2.05
	P2	0.234	0.284	−1.62
	P3	0.240	0.275	−0.80
	P4	0.309	0.309	−0.01
	P5 (Large)	0.443	0.374	0.71
Power Law		1.180	1.169	0.41

Table 4 shows the implications of the restricted model specification detailed in Table 3. We report moments in the data and in the model, and t-statistics for the difference between sample- and model-implied moments. The set of moments consists of average growth rates (Growth) in the aggregate economy and across five size-sorted portfolios (P1–P5), the cross-sectional standard deviation of firms' growth rates, median Tobin's Q, average investment rates in physical capital (Inv-Rate), and executive pay-performance sensitivity (PPS) for the cross section of size-sorted portfolios, and the power law in firm size. Growth rates are measured in logs, PPS is measured in a panel regression of log change in CEO compensation on log firm return, controlling for firm fixed effects, and the power law is estimated using firms in the top size-sorted portfolio. Moments with asterisks are not exploited in estimation. The data are annual, measured in real terms and cover the period from 1993 to 2013.

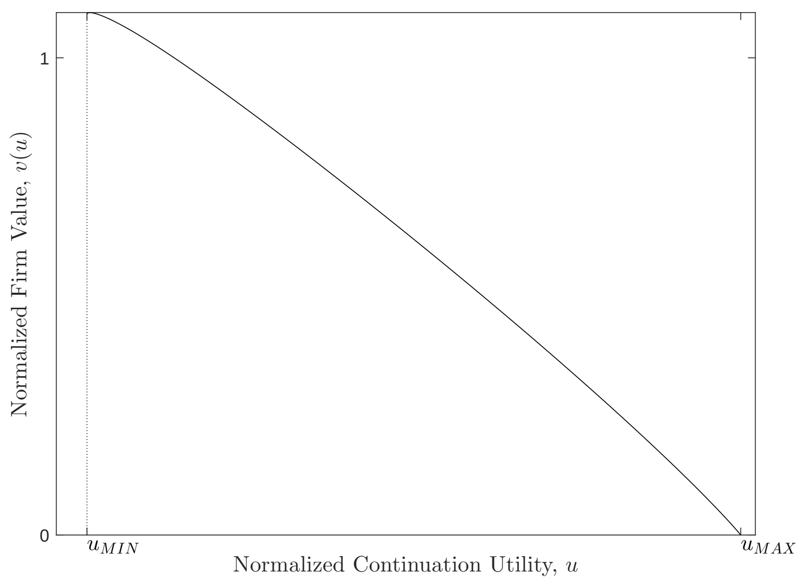

Figure 1. The Normalized Value Function

Figure 1 plots the normalized value function. The horizontal axis represents the normalized continuation utility of the manager, u, and the vertical axis represents the normalized firm value. u_{MIN} and u_{MAX} are the lower and upper bounds of u respectively.

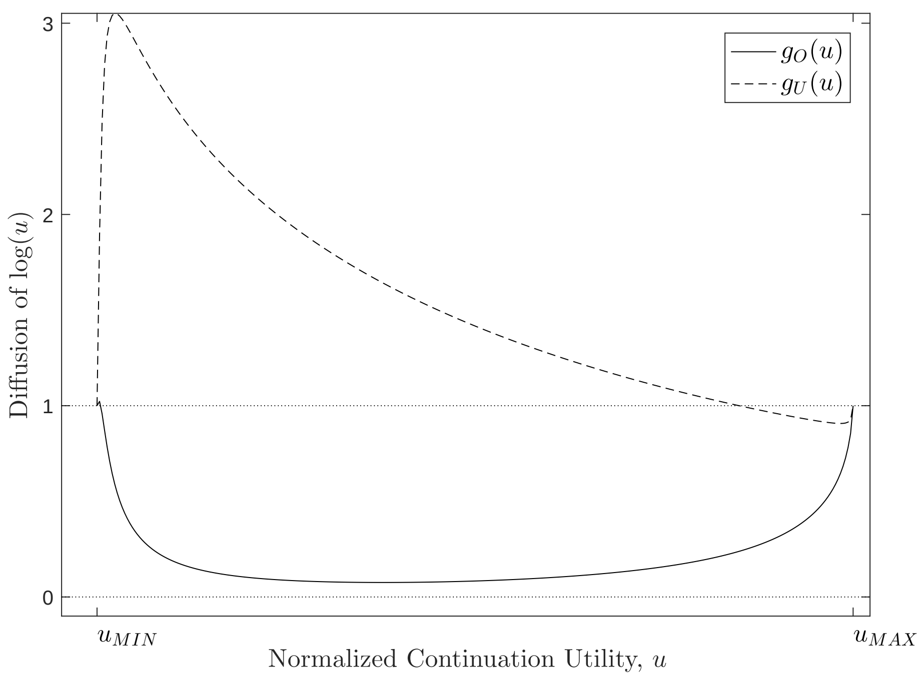

Figure 2. Elasticity of Continuation Utility

Figure 2 plots the sensitivity of continuation utility with respect to observable shocks, g_O (solid line), and with respect to unobservable shocks, g_U (dashed line), under the optimal contract.

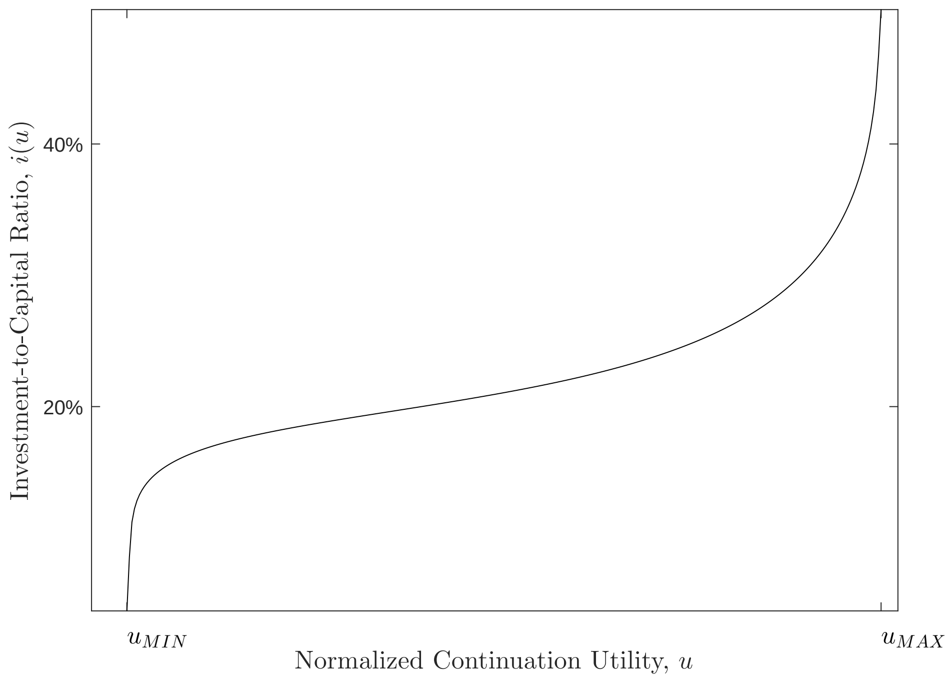

Figure 3. Organization-Capital Investment Rate

Figure 3 plots the organization-capital investment rate, $i(u)$, under the optimal contract as a function of the normalized utility.

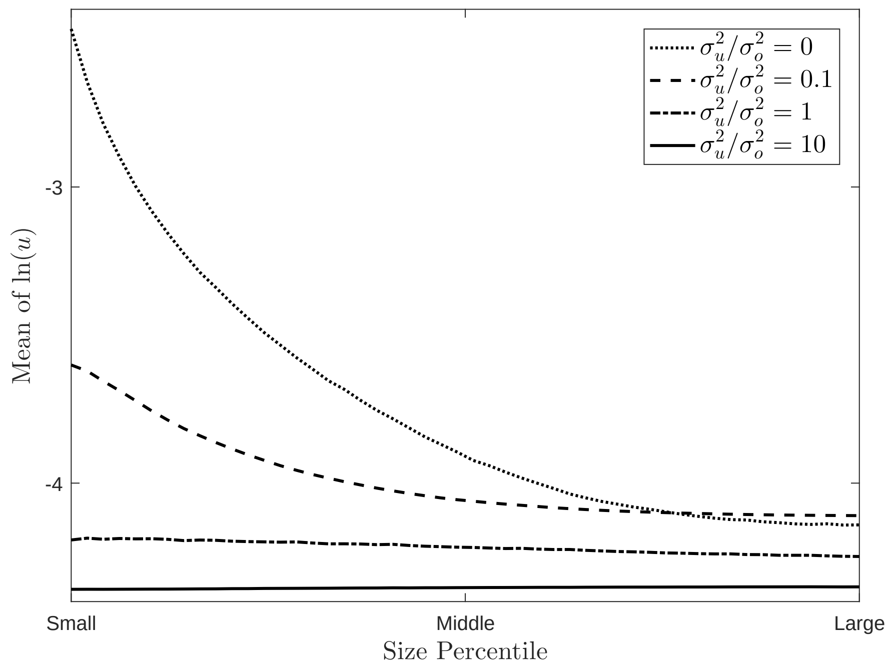

Figure 4. Equilibrium Correlation Between Manager's Share and Firm Size

Figure 4 plots the average manager's share (in logs) as a function of firm size in the equilibrium stationary distribution under different assumptions about the relative magnitude of observable and unobservable shocks.

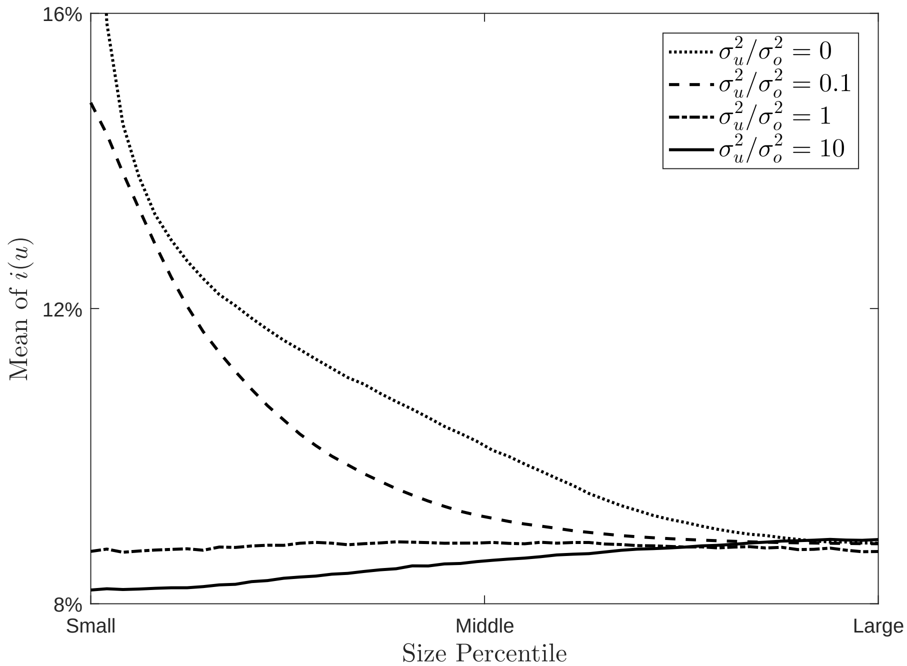

Figure 5. Organization-Capital Investment and Firm Size

Figure 5 plots the average organization-capital investment rate, $i(u)$, as a function of firm size in the equilibrium stationary distribution under different assumptions about the relative magnitude of observable and unobservable shocks.

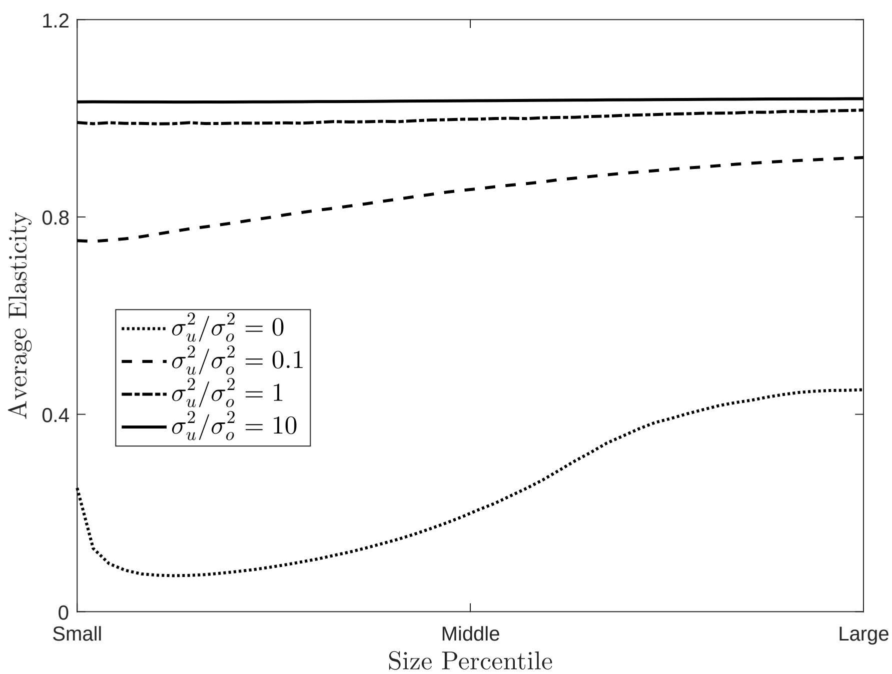

Figure 6. Pay-Performance Sensitivity and Firm Size

Figure 6 plots the average elasticity of continuation utility with respect to observable and unobservable shocks as a function of firm size in the equilibrium stationary distribution under different assumptions about the relative magnitude of observable and unobservable shocks.