Agency Theory with Parametric Uncertainty*

Konrad Grabiszewski† Bruce McWilliams‡

Abstract

Agency models assume parametric certainty, that is, that the true value of marginal productivity is known by both players. We argue that, in general, this assumption is violated. Instead, we introduce marginal productivity as a random variable in the agency model. Parametric uncertainty can explain some empirical observations that the standard model cannot. In particular, we find that, consistent with recent empirical findings, the correlation between risk and uncertainty may be positive, negative, or zero. Similarly, our model can explain the finding of an inverted-U relationship between incentives and firm performance.

Keywords: principal-agent model; moral hazard; parametric uncertainty

JEL classification: D80, D86, L22

*Support from the Asociación Mexicana de Cultura A.C. is gratefully acknowledged.

†Department of Economics, University of Miami, Coral Gables, FL 33124, USA; konrad.grabiszewski@gmail.com

‡Department of Business Administration, ITAM, Av. Camino a Sta. Teresa No.930, Magdalena Contreras, C.P. 10700 Mexico D.F., Mexico; bruce@itam.mx

1 Introduction

In agency theory, the principal designs a contract in order to motivate the agent to exert a specific level of effort. In what we will refer to as the standard model, the agent's effort a generates output B which is defined as $B = ba + \varepsilon$, where b is marginal productivity of effort and ε is a normally distributed random variable with mean zero and variance ν^2. The principal does not directly observe the agent's effort, and bases the agent's payment on the observed output B. The implication of introducing stochastic noise is that the output becomes an imprecise indicator of effort and the principal can no longer accurately deduce the agent's effort from observing the outcome.

As in the standard model, we also assume that it is impossible for the principal to deduce the agent's effort from observable output. However, in our model, this inability is due to the players' lack of precise knowledge about the *marginal productivity* b rather than because external noise confuses the signal. In the standard model, uncertainty is exogenous to the firm and there is an implicit assumption of parametric certainty; that is, the players (i.e., principal and agent) commonly know the true value of b. Here, *parametric* refers to the parameter b.

We develop a principal-agent model with parametric uncertainty in which B is defined as $B = \lambda a$, where marginal productivity λ is normally distributed with mean β and variance σ^2. This model specification can arise in the real world in at least two situations: (1) Past productivity is imperfectly measured; and (2) Productivity is uncertain in changed environments.

Past productivity is imperfectly measured refers to the fact that, even when the agent (in our case a CEO) has been employed with a firm for a long time, her productivity cannot be precisely measured since the CEO performs a variety of tasks and each of these tasks are affected by exogenous factors that are difficult to identify or may be erroneously accounted for, leading to errors in measured productivity. Furthermore, the principal's assessment of the agent's productivity is complicated by his inability to directly observe agent effort, while the agent's assessment of her own productivity are affected by personal biases (e.g., Benabou and Tirole (2002) and Compte and Postlewaite (2004)).

Productivity is uncertain in changed environments refers to the fact that even if an agent's (a CEO's) past productivity is precisely measured, changes may occur in the environment that bring into question future agent productivity. Technology changes are an important source of productivity change that make future outcomes uncertain. In a similar vein, firms may be considering changes in strategy that are expected to increase sales or reduce costs over current practices, and there is greater uncertainty of outcome the greater the changes being made. Larger changes in strategy may also require increased effort on the part of the CEO, so that effort correspondingly increases with expected output and uncertainty. Finally, this case can reflect the hiring of a new, unproven CEO, where there is great uncertainty about how she will perform.

Thus our model is more appropriate when internal shocks (uncertainties) are greater than external shocks, although one would expect both shocks to be important. Parametric uncertainty exposes the risk-averse agent to a different type of risk than does the presence of exogenous shock. In the standard model, the variance of B is the same as the variance of ε, i.e., ν^2. Consequently, the agent's effort has no impact on the level of risk. In our model the agent's effort has a more nuanced impact on output. Since the variance of B under parametric uncertainty is $\sigma^2 a^2$, this implies that the risk-averse agent must choose effort carefully as higher effort increases the level of risk. The fact that the agent's effort affects overall uncertainty is the main difference between our model and the standard model. In Section 2, we develop our model and compare it to the standard model.

Prat and Jovanovic (2014) model multiplicative error in a dynamic model where agent quality is also unknown but there is Baysian learning over time. Their focus is on the effectiveness of incentives over time. Cao and Wang (2013) also introduce a multiplicative error term, but their interest is in building a theoretical model that could be estimated empirically. Since they assume agent risk neutrality, they require multiplicative error in order to motivate non-trivial results. To our knowledge multiplicative error has not been used to address the questions we tackle in this paper.

Parametric uncertainty, while more consistent with how we expect beliefs about effort productivity to be formed, also explains two important empirical phenomena, described in the literature on CEO compensation, which the standard model is unable to accommodate. First, while the standard model predicts a risk-incentive trade-off (a negative correlation between the level of risk and the level of incentives), recent empirical literature, as discussed in Prendergast (2002a), has found that this correlation may be positive, negative, or zero. In Section 3, we discuss how our model can allow for more varied correlations between risk and incentive levels. Second, one line of literature finds that there is an inverted-U relationship between incentives provided to managers and firm performance (e.g., Morck et al. (1988)). In Section 4, we analyze how such a phenomenon can occur under parametric uncertainty. To the best of our knowledge, our paper is the first to offer a principal-agent model that is able to capture both empirical phenomena at the same time.

2 Model

We consider the standard linear[1] agency problem in which the principal (hereafter, referred to as “he”) designs a contract for the agent (hereafter, referred to as “she”) consisting of two elements: the agent’s fixed wage t, and the agent’s commission s, which is the share of output B that she receives. Given the contract and her reservation value R, the agent exerts effort a. In our model, the output is defined as $B = \lambda a$ where λ is a normally distributed variable with mean β and variance σ^2.

The agent is characterized by constant absolute risk aversion and her preferences can be represented by the negative exponential utility function $u(w, a) = -e^{-\eta(w-c(a))}$, where $\eta > 0$ is the coefficient of absolute risk aversion, w is wealth expressed in monetary terms, and $c(a)$ denotes the agent’s cost function. We assume the quadratic cost function, $c(a) = \frac{1}{2}ca^2$. The agent’s profit is $t + sB - c(a)$, and since B is affected by a random element, we determine the agent’s certainty equivalent compensation which is her objective function V^A:

$$V^A = t + s\beta a - \frac{1}{2}\eta\sigma^2 s^2 a^2 - \frac{1}{2}ca^2. \tag{1}$$

Maximization of V^A with respect to a yields the agent’s best-response function a^*:

$$a^* = \frac{\beta s}{c + \eta\sigma^2 s^2}. \tag{2}$$

The principal is risk-neutral and his gain and cost are B and $t + sB$, respectively. Consequently, his net gain is $(1 - s)B - t$, and he chooses a contract (t, s) in order to maximize $V^P = (1 - s)\beta a^* - t$, subject to the constraint $V^A = R$. We determine the principal’s objective function as:

$$V^P = -R + \frac{\beta^2}{2}\frac{2s - s^2}{c + \eta\sigma^2 s^2}. \tag{3}$$

Maximizing V^P with respect to s gives the equilibrium value of the agent’s share $\hat{s}$, which

[1]We introduce parametric uncertainty into the standard linear contract model. Linear contracts have several advantages. As pointed out by various authors, linear contracts are commonly used in real life (see, for instance, Holmstrom and Milgrom (1987), Schmalensee (1989), Milgrom and Roberts (1992), and Bhattacharyya and Lafontaine (1995)). In addition, such contracts are easily tractable and allow us to determine closed-form solutions. Consequently, linear contracts are commonly used in the contract theory literature and our paper follows this approach.

allows us to obtain the equilibrium values of the agent's effort $\hat{a}$ and the principal's profit $\hat{V}^P$.

Proposition 2.1

$$\hat{s} = \frac{\sqrt{c^2 + 4c\eta\sigma^2} - c}{2\eta\sigma^2} \qquad \hat{a} = \frac{\beta}{\sqrt{c^2 + 4c\eta\sigma^2}} \qquad \hat{V}^P = -R + \frac{\beta^2}{2c}\hat{s} \tag{4}$$

The most natural way to explain the results in our model is to compare them to the results in the standard model. Let subscript S indicate the standard model. In that model, the output is defined as $B_S = ba_S + \varepsilon$, where ε is a normally distributed variable with mean zero and variance ν^2. The agent's objective function in the standard model, namely,

$$V_S^A = t + s_S b a_S - \frac{1}{2}\eta\nu^2 s_S^2 - \frac{1}{2}c a_S^2, \tag{5}$$

differs from the agent's objective function in our model, as depicted in (1). In particular, it is the third element in V^A, $-\frac{1}{2}\eta\sigma^2 s^2 a^2$, that constitutes our major departure from the standard model. In the standard model, the agent's effort negatively affects her objective function only via the costs. Consequently, in the denominator of the agent's best-response function a_S^*, the only element is the cost parameter c:

$$a_S^* = \frac{b s_S}{c}. \tag{6}$$

In our model, besides costs, there is an additional channel through which the agent's effort decreases the agent's profit, namely, parametric uncertainty. When the agent does not know the exact value of effort productivity, she considers a range of output values that her effort would yield. Greater effort implies a wider range of possible output, which increases the agent's exposure to risk. The size of that range depends on the variance of λ, σ^2, and the offered share s. The agent's dislike of that range increases with her risk aversion η. Consequently, not only c but also $\eta\sigma^2 s^2$ is in the denominator of a^* in our model.

The analysis of equilibrium values, which for the standard model are depicted below, identifies more differences between the standard model and our model:

$$\hat{s}_S = \frac{b^2}{b^2 + c\eta\nu^2} \qquad \hat{a}_S = \frac{b^3}{cb^2 + c^2\eta\nu^2} \qquad \hat{V}_S^P = -R + \frac{b^2}{2c}\hat{s}_S \tag{7}$$

Since we want to express these differences in terms of the two variances σ^2 and ν^2, and the agent's risk aversion η, we simplify our analysis by assuming that $\beta = b = 1$.

We start with the agent's share. In our model, maximization of the principal's objective function with respect to s yields the following first-order condition: $\hat{s}^2 + \frac{c}{\eta\sigma^2}\hat{s} - \frac{c}{\eta\sigma^2} = 0$. Define $g(s) = \hat{s}^2 + \frac{c}{\eta\sigma^2}s - \frac{c}{\eta\sigma^2}$ and note that $g(\hat{s}) = 0$. Note also that g is a strictly increasing function on the interval $[0, 1]$. Given this, it is enough to determine the sign of $g(\hat{s}_S)$ in order to discover whether or not $\hat{s}_S > \hat{s}$.

[Figure 1 about here.]

Simple algebraic manipulations indicate that the sign of $g(\hat{s}_S)$ depends on the sign of $\sigma^2 - c^2\nu^2 - c^3\eta\nu^4$, which is strictly bigger than zero if and only if $\eta < \bar{\eta}$, where $\bar{\eta} = \frac{\sigma^2 - c^2\nu^2}{c^3\nu^4}$. That is, when the agent's risk aversion is not too large, then $\hat{s}_S > \hat{s}$.

When $\beta = b = 1$, then $\hat{V}_S^P > \hat{V}^P$ if and only if $\hat{s}_S > \hat{s}$. We conclude that $\hat{V}_S^P > \hat{V}^P$ whenever $\eta < \bar{\eta}$.

In order to compare the equilibrium values of the agent's effort, we do not need to rely on the first-order conditions. Analyzing the difference, $\hat{a}_S - \hat{a}$, indicates that $\hat{a}_S > \hat{a}$ if and only if $\eta < \bar{\bar{\eta}}$, where $\bar{\bar{\eta}} = \frac{4\sigma^2 - 2c^2\nu^2}{c^3\nu^4}$.

Proposition 2.2

Assume that $\beta = b = 1$. Define $\bar{\eta} = \frac{\sigma^2 - c^2\nu^2}{c^3\nu^4}$ and $\bar{\bar{\eta}} = \frac{4\sigma^2 - 2c^2\nu^2}{c^3\nu^4}$. Then $\hat{s}_S \geq \hat{s}$ if and only if $\eta \leq \bar{\eta}$, and $\hat{a}_S \geq \hat{a}$ if and only if $\eta \leq \bar{\bar{\eta}}$.

We note that $\bar{\eta} < \bar{\bar{\eta}}$ if and only if $c^2 < 3\frac{\sigma^2}{\nu^2}$; that is, when the ratio of variances is large enough or, equivalently, when c is small enough. This allows us to graphically depict Proposition 2.2 in Figures 2 and 3.

[Figure 2 about here.]

[Figure 3 about here.]

Suppose that Scenario 1 is true, that is, $c^2 < 3\frac{\sigma^2}{\nu^2}$. We consider three cases, as displayed in Figure 2. In case 1, η is smaller than $\bar{\eta}$, which implies that if our model reflects reality, then the standard model overestimates both the agent's share and her effort. In case 2, the standard model overestimates her effort and underestimates her share. Finally, in case 3, η is larger than $\bar{\bar{\eta}}$ and the standard model underestimates both variables. The analysis of Scenario 2 is done in the same fashion (see Figure 3). Note that there is only one case in which both the standard model and our model give the same equilibrium values of the agent's share and effort. This happens when $\bar{\eta} = \bar{\bar{\eta}}$, that is, when $c^2 = 3\frac{\sigma^2}{\nu^2}$. However, in general, we may expect that $\bar{\eta} \neq \bar{\bar{\eta}}$, which implies that, in comparison to our model, the standard model will either overestimate or underestimate the equilibrium values.

3 Risk-Incentives Trade-off

The standard model predicts that the correlation between the risk and incentives is negative because $\frac{d\hat{s}_s}{d\nu^2} < 0$ (see (7)). However, the empirical literature suggests that this prediction is not always correct. Prendergast (2002a), a seminal paper in this area, offers a meta-analysis of the empirical studies focusing on the risk-incentive trade-off in the contracts of CEOs, sharecropping, and franchising relations. The studies examined in his analysis reveal a variety of empirical findings, including positive, negative, and zero correlation between risk and incentives. Rather than repeating his review of those empirical studies here, we add recent empirical papers below that have come out since Prendergast (2002a) and are consistent with each of these correlations. We demonstrate how our model is able to explain a positive, negative, or zero correlation between risk and incentives.

3.1 Positive Correlation Between Risk and Incentives

Recent papers that demonstrate a positive relation between risk and incentives include Prendergast (2002b), Adams (2003) (Hypothesis 1), Raith (2003), Wright (2004), Foss and Laursen (2005) (Hypothesis 1), Oyer and Schaefer (2005), Serfes (2005) (Proposition 2.a), Guo and Ou-Yang (2006), Raith (2008), Gibbs et al. (2009), Budde and Kräkel (2011), Shi (2011), and Chen (2012). Relying on our model with parametric uncertainty, we offer two explanations of this phenomenon.

Our first explanation is based on equilibrium analysis and suggests that the aforementioned positive relationship may be driven by the CEO's risk aversion η. First, note that, as in the standard model, $\frac{\partial \hat{s}}{\partial \sigma^2} < 0$. However, the appropriate variance to analyze is not σ^2 but the variance of B, namely, $\mu = \sigma^2 a^2$. We replace a by its equilibrium value $\hat{a}$ to obtain the equilibrium value of μ as follows:

$$\hat{\mu} = \frac{\beta^2 \sigma^2}{c^2 + 4c\eta\sigma^2}. \tag{8}$$

Clearly, $\frac{\partial \hat{\mu}}{\partial \eta} < 0$ and it is easy to deduce that $\frac{\partial \hat{s}}{\partial \eta} < 0$. Consequently, if the CEO's risk aversion increases, then both risk and the CEO's share decrease. Alternatively, if η decreases, then both variables increase. That is, $\hat{\mu}$ and $\hat{s}$ are positively correlated but the correlation is due to the fact that they are both affected in the same way by η.

Thus, like Raith (2003), we argue that there is a third variable, η, which affects both risk and incentives. As is well known, the agent's risk aversion is among several variables that are used to design a contract in the principal-agent framework, but unfortunately are not observable. This causes a significant problem from the empirical perspective. As Oyer and Schaefer (2011) write, "Theory suggests a long list of unobservables that should matter for CEO pay arrangements. It is not clear how empirical researchers can control for all of these factors well enough to draw firm conclusions about the degree to which CEO pay arrangements are or are not in line with theory." Consequently, we hypothesize that the positive correlation between risk and incentives is driven by two causal relationships: risk aversion affecting the variance of B, and risk aversion affecting incentives. Our hypothesis would especially make sense in cross-sectional analyses in which the data set consists of many firms which employ distinct CEOs characterized by different levels of risk aversion. We could then expect that risk aversion has a negative impact both on the variance of output and on incentives.

Our second explanation depends on the assumption that observable contracts are not equilibrium contracts. It is possible that the positive relationship between risk and incentives is an out-of-equilibrium phenomenon. If this is the case, then note that in our model the agent's best-response function is defined as $a^* = \frac{\beta s}{c + \eta \sigma^2 s^2}$ (see (2)). While it is possible that a^* is decreasing in s (and we spend more time analyzing this possibility in Section 4) let us assume that we face the scenario in which a^* increases in s. Then, increasing the CEO's share will result in a higher CEO effort. However, higher effort means that the volatility of output, $\mu = \sigma^2 a^2$, increases as well. This suggests the following causal relationships:

$$s \nearrow \quad \Rightarrow \quad a \nearrow \quad \Rightarrow \quad \mu \nearrow$$

Such an explanation is not possible in the standard model, in which the variance of B is not affected by the agent's effort since $\mu_s = \nu^2$. This causal relationship of incentives affecting risk is also suggested in Lafontaine and Slade (2000).

3.2 Negative Correlation Between Risk and Incentives

A negative relation between risk and incentives is one of the key predictions of the standard model. A number of papers published after Prendergast (2002a) confirm that prediction, such as Adams (2003) (Hypothesis 3), Aggarwal and Samwick (2003), Garvey and Milbourn (2003), Mengistae and Xu (2004), Serfes (2005) (Proposition 2.b), Wulf (2007), Devaro and Kurtulus (2010), and Bandiera et al. (2012).

In subsection 3.1, we demonstrated that the equilibrium values of the agent's share $\hat{s}$ and the variance of B, $\hat{\mu}$, are positively correlated due to the impact of the third variable, the agent's risk aversion η. Following the same strategy as in subsection 3.1, we need to identify which parameters have opposite impacts on $\hat{s}$ and $\hat{\mu}$. Note the following.

$$\frac{\partial \hat{\mu}}{\partial \sigma^2} > 0 \quad \text{and} \quad \frac{\partial \hat{s}}{\partial \sigma^2} < 0$$

$$\frac{\partial \hat{\mu}}{\partial c} < 0 \quad \text{and} \quad \frac{\partial \hat{s}}{\partial c} > 0$$

Hence, the impact of both σ^2 (idiosyncratic risk) and c on $\hat{\mu}$ and $\hat{s}$ explains why the empirically observed correlation between risk and incentives is negative. While the impact of idiosyncratic risk is the standard explanation for this negative correlation, the impact of the cost of effort introduces yet another unobserved variable that can affect this correlation.

3.3 Zero Correlation Between Risk and Incentives

Foss and Laursen (2005) (Hypothesis 2b) and Edmans et al. (2009) suggest that there is no relationship between risk and incentives. Therefore, we want to determine if there is a parameter in our model such that there is no relationship between $\hat{s}$ and $\hat{\mu}$. Note that while $\frac{\partial \hat{\mu}}{\partial \beta} > 0$, it is the case that $\frac{\partial \hat{s}}{\partial \beta} = 0$. Hence, if the data set consists of agents with distinct expected effort productivity parameters, then the variability in β may explain why there is no correlation between risk and incentives. Obviously, an alternative explanation is that the positive and negative impacts mentioned in the previous subsections may counteract each other and generate a net zero correlation between risk and incentives.

In summary of subsections 3.1-3.3, because the variance under parametric uncertainty is affected by effort, the correlation between risk and incentives may be positive, negative, or zero. The correlation is positive if risk aversion drives the correlation, since it has the same qualitative impact on both incentives and variance. The correlation is negative if either idiosyncratic risk or the cost of effort drives the correlation, since they both have the opposite qualitative impact on incentives and variance. The correlation can be zero if (i) expected effort productivity drives the correlation, since it has no impact on incentives; or (ii) the positive and negative impacts of the other variables offset each other.

4 Increase in CEO's Share Can Reduce Effort

Equation (6) indicates that the standard model posits the convergence-of-interest hypothesis according to which there is a positive relationship between CEO incentives and firm performance. At the core of this hypothesis is a belief that increasing the CEO's share aligns the interests of the manager and the firm, and therefore it is beneficial for the firm. However, the empirical literature suggests that this hypothesis works well only at low ownership levels and that firm performance may decline for high incentives. Consequently, the relationship between incentives and performance may exhibit an inverted-U shape.

Morck et al. (1988) find that firm performance increases as management ownership increases up to a given level (in their study, 5 percent), but that performance decreases as management ownership increases beyond that level. Since then, other researchers have confirmed this inverted-U relationship between the share of ownership by corporate insiders (usually defined as managers and members of the board) and corporate performance (see McConnell and Servaes (1990), Hermalin and Weisbach (1991), Hubbard and Palia (1995), Holderness et al. (1999), Fogelberg and Griffith (2000), Claessens et al. (2002), Anderson and Reeb (2003), Tian (2004), Davies et al. (2005), Adams and Santos (2006), Pukthuanthong et al. (2007), McConnell et al. (2008), Benson and Davidson (2009), Fahlenbrach and Stulz (2009), Kim and Lu (2011), and Zhang (2013)).

Why would firm performance decrease when the CEO's share is high? The entrenchment hypothesis (e.g., Morck et al. (1988)) suggests that when managers accumulate sufficient ownership of a firm, they have the power to guarantee their continued employment at the firm while indulging their preference for non-value-maximizing behavior (e.g., promoting projects that serve personal interests). This entrenchment effect overcomes the convergence-of-interest effect when managers own a large percent of the company they run, leading to the inverted-U relationship. For instance, Stulz (1998) develops a model in which management teams vie for control of a firm through ownership. He concludes that there is an optimal level of management ownership, above and below which the firm would be managed inefficiently, leading to the inverted-U.

An alternative explanation for declining performance at high incentive levels is given by the risk-reduction hypothesis (e.g., Benson and Davidson (2009)), which exploits the fundamental assumption in agency theory that incentives expose the risk-averse agent to risk which she dislikes. Managers' risk-reducing behavior will hurt firm performance when the incentives are too high. Thus, the risk-reduction effect can overcome the convergence-of-interest effect at high ownership levels, yielding the inverted-U relationship. For instance, Lambert (1986) develops a model in which the manager exerts effort to learn the expected value and riskiness of different projects and then selects the project to be undertaken by the firm. Because the manager is risk-averse, she may select a project that is less risky than another with a higher expected outcome. High incentives can impose excessive risk on the manager, inducing the executive to exert less effort on identifying the appropriate project and to select projects that are suboptimal from the principal's perspective.

Our model is consistent with the risk-reduction hypothesis. In our model, incentives and effort interact to generate greater risk for the risk-averse agent. At excessive incentive levels, the agent may choose to reduce effort in order to moderate the increase in risk caused by the high incentives. Thus, firm performance can suffer at high incentive levels.

Since expected output is increasing in effort in both the standard model and our model, we focus our attention on the agent's effort. We consider the implications of increasing CEO incentives in both models. In the standard model, expressing CEO effort in terms of the incentives she receives yields the best-response function a_S^* given in (6). In this model, the agent's effort increases linearly with the share offered by the firm. However, when we consider parametric uncertainty, increasing incentives may not always lead to greater effort. Expressing effort as a function of the share that the agent receives yields a^*, given in (2). Taking the derivative of a^* with respect to s yields:

$$\frac{\partial a^*}{\partial s} = \frac{\beta(c - \eta\sigma^2 s^2)}{(c + \eta\sigma^2 s^2)^2}. \tag{9}$$

Hence, $\frac{\partial a^*}{\partial s} > 0$ if and only if:

$$s < \sqrt{\frac{c}{\eta\sigma^2}}. \tag{10}$$

This implies that the larger the CEO share, the less likely the CEO is going to increase her effort if incentives increase.

Figure 4 depicts the case in which increasing share (incentives) always has a positive impact on effort. Figure 5 depicts the case in which increasing share increases effort up to some maximal point, reached where $\bar{s} = \sqrt{\frac{c}{\eta\sigma^2}}$, beyond which effort decreases with additional increases in incentives.

[Figure 4 about here.]

[Figure 5 about here.]

The inverted-U relationship depicted in Figure 5 occurs when $\bar{s} < 1$; that is, when $c < \eta\sigma^2$. Consequently, Figure 5 describes the true relationship between incentives and effort when the agent's risk aversion is large, the idiosyncratic risk is large, or the cost of effort is low.

Recall that the variance of output under parametric uncertainty is $\sigma^2 a^2$, and the negative utility from variance is given by $\frac{1}{2}\eta\sigma^2 s^2 a^2$. Therefore, at high incentive levels, the agent's effort and incentives interact to generate excessive negative utility from risk. The risk-averse agent responds by reducing effort, which in turn reduces her risk. As long as the contract design is based on parametric uncertainty, excessive incentives are not an issue since they cannot arise in the equilibrium. However, if other factors enter into the selection of the manager's ownership of the firm, incentives may be either too low or too high, and we would observe an inverted-U in performance similar to that shown in Figure 5. For example, if the principal designs the contract based on the standard model while the agent bases her response function on parametric uncertainty, equation (10) may be violated and firm performance will decline. In this example, if firms are heterogeneous in the parameters given in (10), some firms may offer excessively low incentives while others may offer excessively high incentives, leading to the inverted-U relationship.

5 Conclusions

We argue that the players cannot, in general, know the true value of the effort productivity parameter. Therefore, we introduce effort productivity as a random, rather than fixed, parameter in the principal-agent model. This parametric uncertainty provides an explanation for some empirical observations that the standard model cannot explain. While the standard model predicts only a negative correlation between risk and incentives, our model reveals that a positive, or even zero, correlation between risk and incentives can also arise. This variety of theoretical outcomes is consistent with the variety of empirical findings for this relationship that we have cited. Furthermore, the standard model predicts a consistently positive relationship between the level of incentives and effort. Under parametric uncertainty however, the agent can actually reduce her effort if incentives are excessive. Since performance is positively related to effort, this implies that firm performance will initially be increasing as CEO incentives increase, but at some point additional incentives may lead to declining performance.

A benefit of modeling parametric uncertainty is that it is more consistent with how we believe players form their beliefs about effort productivity. That is, neither side can know the true value of b. However, introducing a variance for b should not be seen as a substitute for the exogenous error typically modeled. The two types of errors arise from different sources and have different implications.

References

Adams, Christopher P., "Agent Discretion, Adverse Selection and the Risk-Incentive Trade-Off," *FTC Working Paper No. 255*, 2003.

Adams, Renée B. and João A. C. Santos, "Identifying the Effect of Managerial Control on Firm," *Journal of Accounting and Economics*, 2006, *41*, 55–85.

Aggarwal, Rajesh K. and Andrew A. Samwick, "Endogenous Matching and the Empirical Determinants of Contract Form," *Journal of Finance*, 2003, *58*, 1613–1650.

Anderson, Ronald C. and David M. Reeb, "Founding-Family Ownership and Firm Performance: Evidence from the S&P 500," *Journal of Finance*, 2003, *58*, 1301–1328.

Bandiera, Oriana, Luigi Guiso, Andrea Prat, and Raffaella Sadun, "Matching Firms, Managers, and Incentives," *CEP Discussion Paper No. 1144*, 2012.

Benabou, Roland and Jean Tirole, "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 2002, *117*, 871–915.

Benson, Bradley W. and Wallace N. Davidson, "Reexamining the managerial ownership effect on firm value," *Journal of Corporate Finance*, 2009, *15*, 573–586.

Bhattacharyya, Sugato and Francine Lafontaine, "Double-Sided Moral Hazard and the Nature of Share Contracts," *Rand Journal of Economics*, 1995, *26*, 761–781.

Budde, Jörg and Matthias Kräkel, "Limited Liability and the Risk-Incentive Relationship," *Journal of Economics*, 2011, *102*, 97–110.

Cao, Melanie and Rong Wang, "Optimal CEO Compensation with Search: Theory and Empirical Evidence," *Journal of Finance*, 2013, *68*, 2001–2058.

Chen, Ying-Ju, "Risk-Incentives Trade-off and Outside Options," *OR Spectrum*, 2012, pp. 1–20.

Claessens, Stijn, Simeon Djankov, Joseph P. H. Fan, and Larry H. P. Lang, "Disentangling the Incentive and Entrenchment Effects of Large Shareholdings," *Journal of Finance*, 2002, *57*, 2741–2771.

Compte, Olivier and Andrew Postlewaite, "Confidence-Enhanced Performance," *American Economic Review*, 2004, *94*, 1536–1557.

Davies, J. R., David Hillier, and Patrick McColgan, "Ownership structure, managerial behavior and corporate value," *Journal of Corporate Finance*, 2005, *11*, 645–660.

Devaro, Jed and Fidan Ana Kurtulus, "An Empirical Analysis of Risk, Incentives and The Delegation of Worker Authority," *Econometrica*, 2010, *63*, 641–661.

Edmans, Alex, Xavier Gabaix, and Augustin Landier, "A multiplicative model of optimal CEO incentives in market equilibrium," *Review of Financial Studies*, 2009, *22*, 4881–4917.

Fahlenbrach, Rüdiger and René M. Stulz, "Managerial ownership dynamics and firm value," *Journal of Financial Economics*, 2009, *92*, 342–361.

Fogelberg, Lawrence and John M. Griffith, "Control and Bank Performance," *Journal of Financial and Strategic Decisions*, 2000, *13*, 63–69.

Foss, Nicolai J. and Keld Laursen, "Performance pay, delegation and multitasking under uncertainty and innovativeness: An empirical investigation," *Journal of Economic Behavior and Organization*, 2005, *58*, 246–276.

Garvey, Gerald and Todd Milbourn, "Incentive Compensation When Executives Can Hedge the Market: Evidence of Relative Performance Evaluation in the Cross Section," *Journal of Finance*, 2003, *58*, 1557–1582.

Gibbs, Michael, Kenneth A. Merchant, Wim A. Van der Stede, and Mark E. Vargus, "Performance Measure Properties and Incentive System Design," *Industrial Relations*, 2009, *48*, 237–264.

Guo, Ming and Hui Ou-Yang, "Incentives and Performance in the Presence of Wealth Effects and Endogenous Risk," *Journal of Economic Theory*, 2006, *129*, 150–191.

Hermalin, Benjamin E. and Michael S. Weisbach, "The Effects of Board Composition and Direct Incentives on Firm Performance," *Financial Management*, 1991, *20*, 101–112.

Holderness, Clifford G., Randall S. Kroszner, and Dennis P. Sheehana, "Were the Good Old Days That Good? Changes in Managerial Stock Ownership Since the Great Depression," *Journal of Finance*, 1999, *54*, 435–469.

Holmstrom, Bengt and Paul Milgrom, "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 1987, *55*, 303–328.

Hubbard, R. Glenn and Darius Palia, "Benefits of Control, Managerial Ownership, and the Stock Returns of Acquiring Firms," *RAND Journal of Economics*, 1995, *26*, 782–792.

Kim, E. Han and Yao Lu, "CEO ownership, external governance, and risk-taking," *Journal of Financial Economics*, 2011, *102*, 272–292.

Lafontaine, Francine and Margaret E. Slade, "Incentive Contracting and the Franchise Decision," in K. Chatterjee and W. Samuelson, eds., *Advances in business applications of game theory*, Kluwer Academic Press, 2000.

Lambert, Richard, "Executive Effort and Selection of Risky Projects," *RAND Journal of Economics*, 1986, *17*, 77–88.

McConnell, John J. and Henri Servaes, "Additional evidence on equity ownership and corporate value," *Journal of Financial Economics*, 1990, *27*, 595–618.

_ , _ , and Karl V. Lins, "Changes in insider ownership and changes in the market value of the firm," *Journal of Corporate Finance*, 2008, *14*, 92–106.

Mengistae, Taye and Lixin Colin Xu, "Agency Theory and Executive Compensation: The Case of Chinese State-Owned Enterprises," *Journal of Labor Economics*, 2004, *22*, 615–637.

Milgrom, Paul and John Roberts, *Economics, Organization and Management*, Englewood Cliffs, NJ: Prentice Hall, 1992.

Morck, Randall, Andrei Shleifer, and Robert W. Vishny, "Management Ownership and Market Valuation: An Empirical Analysis," *Journal of Financial Economics*, 1988, *20*, 293–315.

Oyer, Paul and Scott Schaefer, "Why do some firms give stock options to all employees?: An empirical examination of alternative theories," *Journal of Financial Economics*, 2005, *76*, 99–133.

_ **and** _ , "Personnel Economics: Hiring and Incentives," in David Card and Orley Ashenfelter, eds., *Handbook of Labor Economics*, Vol. 4, Part B, Elsevier, 2011.

Prat, Julien and Boyan Jovanovic, "Dynamic contracts when agent's quality is unknown," *Theoretical Economics*, 2014, *9*, 865–914.

Prendergast, Canice, "The Tenuous Trade-Off between Risk and Incentives," *Journal of Political Economy*, 2002, *110*, 1071–1102.

_ , "Uncertainty and Incentives," *Journal of Labor Economics*, 2002, *20*, S115–S137.

Pukthuanthong, Kuntara, Richard Roll, and Thomas Walker, "How employee stock options and executive equity ownership affect long-term IPO operating performance," *Journal of Corporate Finance*, 2007, *13*, 695–720.

Raith, Michael, "Competition, Risk, and Managerial Incentives," *American Economic Review*, 2003, *93*, 1425–1436.

_ , "Specific Knowledge and Performance Measurement," *RAND Journal of Economics*, 2008, *39*, 1059–1079.

Schmalensee, Richard, "Good Regulatory Regimes," *Rand Journal of Economics*, 1989, *20*, 417–436.

Serfes, Konstantinos, "Risk sharing vs. incentives: Contract design under two-sided heterogeneity," *Economics Letters*, 2005, *88*, 343–349.

Shi, Lan, "Respondable risk and incentives for CEOs: The role of information-collection and decision-making," *Journal of Corporate Finance*, 2011, *17*, 189–205.

Stulz, René M., "Managerial Control of Voting Rights: Financing Policies and the Market for Corporate Control," *Journal of Financial Economics*, 1998, *20*, 25–54.

Tian, Yisong S., "Too much of a good incentive? The case of executive stock options," *Journal of Banking and Finance*, 2004, *28*, 1225–1245.

Wright, Donald J., "The Risk and Incentives Trade-off in the Presence of Heterogeneous Managers," *Journal of Economics*, 2004, *83*, 209–223.

Wulf, Julie, "Authority, Risk, and Performance Incentives: Evidence from Division Manager Positions Inside Firms," *Journal of Industrial Economics*, 2007, *55*, 169–196.

Zhang, Feng, "Shareholder Rights, Managerial Incentives, and Firm Value," *working paper*, 2013.

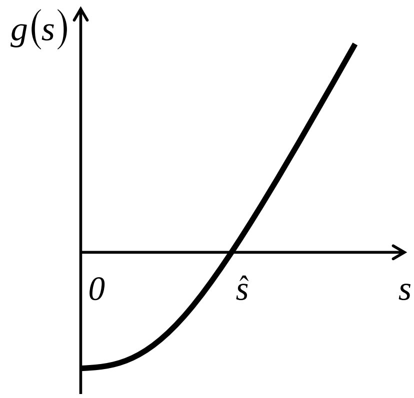

Figure 1

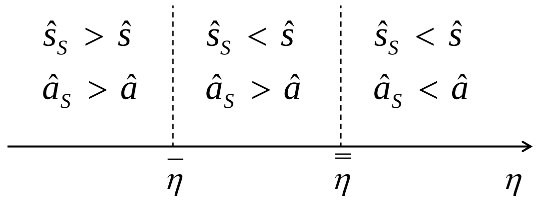

Figure 2: Scenario 1, $c^2 < 3\frac{\sigma^2}{\nu^2}$.

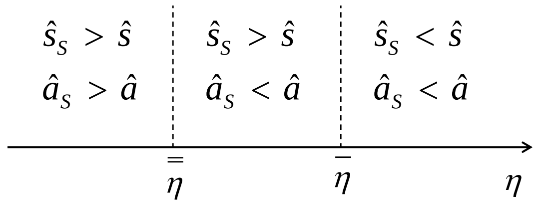

Figure 3: Scenario 2, $c^2 > 3\frac{\sigma^2}{\nu^2}$.

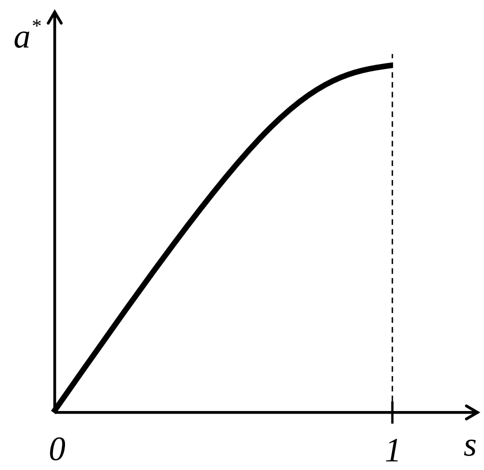

Figure 4: Increasing share always increases effort.

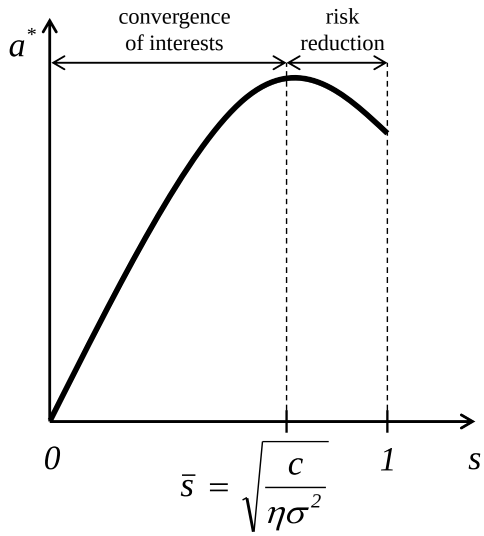

Figure 5: Increasing share can decrease effort.