

On the Role of Uncertainty in the Risk-Incentives Tradeoff

Heikki Rantakari*
Marshall School of Business
University of Southern California

This Version: December 29, 2007

Abstract

We use a simple agency model to clarify and characterize the various avenues through which changes in the level of uncertainty impact the optimal strength of linear incentives. Instead of attempting to characterize different "types" of uncertainty, which has been the approach in the literature so far, we base our characterization on the link between uncertainty and the agent's action choice. We then use this characterization to provide conditions under which the relationship between uncertainty and incentives *can* be positive and relate it back to the existing models in the literature.

*Contact: rantakar@marshall.usc.edu. I would like to thank my advisors Bengt Holmström and Robert Gibbons for their continuous encouragement and many helpful discussions, and the editor and two anonymous referees for their extremely helpful comments. All remaining errors are my own.

1 Introduction

One of the first predictions of the theoretical agency literature was the prediction of a negative trade-off between risk and incentives. An increase in the volatility of the available performance measures increases the volatility of the agent's pay for any given level of pay-for-performance. As a consequence, when the agent is risk-averse, incentives become costlier to provide, and so the equilibrium level of incentives should go down as the volatility of the performance measures goes up, other things constant. This negative relationship, however, has not found strong empirical support, and in his review of the empirical literature, Prendergast (2002) concludes that the aggregate evidence for a negative trade-off between risk and incentives is weak at best.

This mounting non-evidence has, in turn, inspired a theoretical re-examination of the relationship between incentives and uncertainty, where the recent papers have both enriched the basic moral hazard framework by introducing different "types" of uncertainty (Zabojnik 1996, Baker and Jorgensen 2003 and Shi 2007, for example) and added considerations of delegation (Prendergast 2002, Raith 2007 and Rantakari 2007, for example). The common modification applied in these papers is as follows: First, instead of trying to induce a given action, the principal's preferred action is unknown at the time of contracting. Second, after the signing of the contract but before his final action choice, the agent either passively receives or actively acquires information about the realization of that preferred action. The authors then make a distinction between uncertainty over the principal's preferred action and uncertainty that generates only performance measurement problems and show that while the equilibrium incentive strength is negatively related to the latter, it can be positively related to the former type of uncertainty.¹

While these papers highlight well the ideas that uncertainty can do much more than generate performance measurement problems and that one purpose of incentives is to manage the use and acquisition of information by the agent, any simple classification of the relationship between uncertainty and incentives that is based on the type of uncertainty quickly runs into difficulties when we consider a broader spectrum of agency problems. After all, any type of uncertainty, knowledge of the realization of which or simply the level of which

¹Various terms have been suggested to make this distinction. Zabojnik (1996) distinguishes between "ex ante" and "ex post" uncertainty. Baker and Jorgensen (2003) make the distinction between "volatility" and "noise," Raith (2007) calls them "technological" and "environmental" uncertainty, and Shi (2007) calls them "responsible" and "non-responsible" risk.

impacts either the action choice of the agent or the action that the principal would like to induce, will also impact the equilibrium strength of incentives and sometimes in a manner that generates a positive relationship between uncertainty and incentives.

To complement the existing literature, this paper proposes an alternative characterization of the link between uncertainty and incentives that is based on the impact that uncertainty has on the value and choice of action by the agent. With the help of a reduced-form agency model, we identify four basic channels through which an increase in uncertainty can lead to an increasing strength of incentives and then use this categorization to provide simple examples and discuss the existing literature.

The first channel is the *value effect* of uncertainty. If an increase in uncertainty increases the marginal value of effort, then the principal can find it beneficial to increase the strength of incentives. For example, if the task of the agent is to evaluate projects, then increasing the volatility of project returns can increase the marginal value of effort to limit the rejection of high-value projects and the acceptance of low-value projects. This channel is the driving force behind the model of Shi (2007), where the agent first exerts effort to acquire information about the marginal return to investment and then chooses an investment level given his information.

The second and most examined channel is the *interaction effect* of uncertainty. Even if the agent's action choice is independent of the *level* of uncertainty given the realized information, changing the amount of uncertainty changes the distribution of expected outcomes and, thus, the expected value of incentives. This channel is the driving force behind the models of random productivity, such as Zabojnik (1996), Baker and Jorgensen (2003), and Raith (2007). In these models, the state and the agent's effort are complements, so taking a mean-preserving spread of the productivity distribution leads to an increase in the expected value of effort. An examination of this channel also extends the results of these papers by providing a more general condition that the production technology has to satisfy in order to generate the possibility of a positive relationship between uncertainty and incentives.

The third channel is the *level effect* of uncertainty. If an increase in uncertainty causes the agent to reduce the amount of effort that he is devoting to productive activities, then the principal can find it beneficial to increase the strength of incentives to counter this diversion of effort. For example, in a random productivity model, where the agent has only imperfect information about the marginal productivity of his effort, an increase in uncertainty will directly reduce the effort put in by the agent for any given realization of the

signals if the agent is risk-averse.² This effect can also be present in settings that are not typically associated with the discussion of the relationship between risk and incentives, such as influence activities (Milgrom and Roberts 1988, for example) and career concerns (Gibbons and Murphy 1992, for example).

The fourth and final channel is the *sensitivity effect* of uncertainty. If an increase in uncertainty makes the agent more responsive to incentives, then the principal is able to get more "bang for a buck" from any given increase in incentives and can, thus, find it beneficial to increase the incentive strength. For example, an increase in uncertainty can lead to increased job flexibility granted to the agent, which in turn can make the agent more responsive to incentives. Such a setting is analyzed in Rantakari (2007).

Of course, in any given setting, a number of these channels can be present simultaneously, and they can either complement each other or work against each other. Further, any benefits from increasing the incentive strength need to be weighed against the resulting costs, such as the additional risk-compensation that needs to be paid to the agent (when the agent is risk-averse) or the increased risk of gaming of the contract by the agent (when the performance measures can be misaligned with the principal's objective). However, we hope that our simple classification can function as a guiding tool for thinking about the relationship between uncertainty and incentives in any given setting by providing intuitive conditions under which the relationship *can* be positive. Also, by highlighting the rich role that uncertainty can play in any given agency problem, we hope that this illustration encourages future work to pay more detailed attention to the treatment of uncertainty than is typically accorded to it.

²I would like to thank an anonymous referee for this observation.

2 The Framework

We consider a setting in which a risk-neutral principal hires a risk-averse (CARA) agent to undertake an action, where the efficient action can be uncertain at the time of contracting. The timing of the events is as follows: First, the principal offers a linear incentive contract. Second, the agent accepts or rejects the contract before any of the uncertainties are resolved. Third, the agent can observe and/or actively acquire information about the environment. Fourth, the agent makes his final action choice, and the payoffs are realized. In this section, we will derive the program that defines the optimal level of linear incentives. The steps parallel those of most agency models. The key differences are highlighted at the end of this section.

2.1 The Principal

Letting e denote the action of the agent, the payoff to the principal is given by

$$Y = \Pi(\theta, e, \sigma_\theta^2) + \eta,$$

where $\theta \sim F(\theta)$ indexes the productive conditions and $\eta \sim G(\eta)$, $E(\eta) = 0$ captures any exogenous shocks to the final payoff. We assume that both distributions are non-degenerate and continuous. We assume that $Y_e(\cdot) \geq 0$ and $Y_{ee}(\cdot) \leq 0$ while we will return to the other signs later in the analysis. While we are modeling effort as a one-dimensional variable, it is best viewed as a multidimensional variable, where some dimensions can be directed at information acquisition and some dimensions are directed at responding to that information. Some dimensions of effort can be fully non-productive.

The principal's true payoff is not contractible. Instead, there exists a performance measure

$$X = \pi(\mu, \theta, e, \sigma_\theta^2, \sigma_\mu^2) + \varepsilon,$$

where $\mu \sim V(\mu)$ and $\varepsilon \sim W(\varepsilon)$, $E(\varepsilon) = 0$ are additional shocks that enter the performance measure. We assume that $X_e(\cdot) \geq 0$ and $X_{ee}(\cdot) \leq 0$ and that $\text{sign}(X_{e\sigma_\theta^2}) = \text{sign}(Y_{e\sigma_\theta^2})$. We also assume that $\text{Cov}(Y_e, X_e) > 0$, so that the performance measure is informative of the value of the action to the principal, and that $\partial \text{Cov}(Y_e, X_e) / \partial \sigma_\mu^2 \leq 0$, so that the informativeness

of the performance measure is decreasing in the amount of additional uncertainty. Following Baker (1992), we also assume that the performance measure is normalized so that $EX_e = EY_e$.

The principal will then use the performance measure to incentivize the agent to improve the principal's payoff. In particular, before any of the shocks are realized, the principal offers a linear incentive contract

$$T = \alpha X + \beta$$

to the agent, where β is a fixed transfer used to satisfy the agent's participation constraint and α determines the level of incentives.³ To set up some terminology, we will refer to σ_θ^2 as *productive uncertainty* and the remainder of the shocks collectively as *noise*. This distinction is based on the observation that only the realization of θ (or the level σ_θ^2) influences the action that the principal would like to induce from the agent.⁴ However, as will become clear below, this distinction is somewhat arbitrary because any of the shocks (or the variance of which) can influence the action choice of the agent, given the level of incentives, and so lead to changes in the optimal strength of incentives. Any reference simply to *uncertainty* is applicable to any of the four shocks, and we will use $\boldsymbol{\sigma}^2$ to denote the vector $(\sigma_\theta^2, \sigma_\mu^2, \sigma_\varepsilon^2, \sigma_\eta^2)$.⁵ Finally, we assume that the principal has no direct access to information about the specific realization of these shocks at any time. However, the principal might attempt to make inferences about the shocks based on the realization of outcomes.

³We are assuming linear incentives purely for tractability of the framework and have no claims on the general optimality of linear incentives. Indeed, in the examples we will consider later, there are often explicit reasons why linear incentives will *not* be optimal. However, the forces that impact the shape of incentives are usually relatively unrelated to the forces that drive the expected value of incentives, for which the linear setting provides a simple measure. For example, in models of random productivity, non-linear incentives can be used to induce more efficient self-selection by the agent. In this setting, the shape would be driven by the incentive-compatibility constraint of the agent coupled with the informativeness of the performance measure about the underlying state, as in any screening setting. However, the task of optimally satisfying the incentive compatibility constraint for self-selection is relatively distinct from choosing the overall level of effort provided.

⁴In alternative terms, any shock (or the variance of which) that impacts the payoff-maximizing (*not* performance-measure maximizing) action choice is called productive uncertainty.

⁵Of course, σ_η^2 plays only a peripheral role in the analysis because it does not impact the preferred action of either party and it does not directly enter the agent's compensation.

2.2 The Agent

The agent has CARA preferences and faces a monetary cost of effort of

$$C = C(e, \sigma^2).$$

Thus, we assume that the cost function itself does not exhibit any uncertainty and we also assume that $C_e \geq 0$ and $C_{ee} > 0$, as in most agency models. However, we do allow the levels of uncertainty over the different parameters to impact the shape of the cost function, which could be the result of job design, for example.

To allow the framework also to contain models of signal jamming (or separating equilibria of signaling models), such as influence activities and career concerns, in an internally consistent way, we will add a *perceived benefit of effort* function

$$B = B(e, \sigma^2, E_P),$$

where E_P stands for the expectations of the principal and the signs of the derivatives will be application-dependent.

Given the contract $T = \alpha X + \beta$, the agent's utility function has the certainty-equivalence representation of

$$U = E_A((\alpha X + B) | \omega) + \beta - C - R(e, \alpha, \sigma^2),$$

where ω stands for the information about the underlying parameters that is available to the agent at the time when the utility is evaluated and $R(\cdot)$ is the utility cost of the net compensation risk (risk premium) for the agent at the time when utility is evaluated:

$$R(\cdot) = U(E_A(\alpha X + \beta + B - C | \omega)) - E_A(U(\alpha X + \beta + B - C) | \omega).$$

Because this risk premium is dependent on the information held by the agent at the time the utility function is evaluated, we will use $R(\cdot)$ to denote the risk premium at the time of contracting (when the agent knows only the distributions of uncertainty), and $r(\cdot)$ to denote the interim risk premium. If the agent learns no information, then $r(\cdot) = R(\cdot)$, while, if the agent learns perfectly the realization of all the shocks, then $r(\cdot) = 0$.

If we had assumed that the agent had simple mean-variance preferences with risk-aversion coefficient r , then the risk premium would be simply the

variance of the net compensation: $R(\cdot) = r\text{Var}(\alpha X + B - C)$. However, for the CARA family, while the constant β drops out because of the absence of income effects, the risk premium does not generally simplify any further.⁶

Having accepted the contract $T = \alpha X + \beta$ offered by the principal and given the information ω available to him when choosing his action, the agent then solves

$$\max_e E_A(\alpha X + B|\omega) + \beta - C - r(e, \alpha, \sigma^2),$$

which implies that the choice of action $e^*(\cdot, \omega)$ is implicitly defined (assuming a unique interior solution) by

$$E_A(\alpha X_e + B_e|\omega) - r_e(e, \alpha, \sigma^2) = C_e,$$

which simply states that at the time of the action, given his information, the agent equates the expected marginal benefit of an action with the marginal cost of that action. Note that the action choice of the agent can depend on both his ex ante information about the levels of uncertainty σ^2 and any additional information ω learned before the action choice about the realized shocks.

With respect to the information that is available to the agent, we will assume that both η and ε are never observed by the agent. However, since the realization of neither of these shocks would impact the action choice of the agent, this assumption is largely without loss of generality. Knowledge of θ and μ , on the other hand, can impact the agent's action choice, and we will consider different situations in which the agent's information varies from nothing to perfect information.

⁶As shown in Holmström and Milgrom (1991), the mean-variance representation is also valid in a setting with exponential utility and additive normal shocks. Since our shocks need not be additive, the mean-variance representation of the agent's utility is no longer generally valid even if the shocks were otherwise normal and the agent's utility function was exponential.

2.3 The ex ante problem

In the contracting stage, after substituting out the participation constraint of the agent, we can write the principal's problem simply as

$$\begin{aligned} \max_{\alpha} & E(Y(\theta, e^*, \sigma_{\theta}^2) - C(e^*, \sigma^2) - R(e^*, \alpha, \sigma^2)) \\ \text{s.t.} & \quad E_A(\alpha X_e + B_e | \omega) - r_e(e^*, \alpha, \sigma^2) = C_e. \end{aligned}$$

The objective function of the principal has thus been decomposed into two components: a *productive* component $E(Y(\theta, e^*, \sigma_{\theta}^2) - C(e^*, \sigma^2))$, which measures the monetary surplus generated in expectation to the principal given $e^*(\alpha, \cdot)$, and a *risk-compensation* component $R(e^*, \alpha, \sigma^2)$, which measures how much of that monetary surplus is dissipated due to the risk-aversion of the agent. Most importantly, note that while $B_e \neq 0$ if signal-jamming considerations are present, $E(B(e, \sigma^2, E_P)) = 0$ because the principal is not fooled in equilibrium. In consequence, it cannot be used to provide an alternative means of compensation to the agent (which, in turn, would make C and B equivalent) and so does not enter the surplus maximization problem except through its effect on the agent's action choice.

The first-order condition is then

$$FOC_P : E(Y_e(\theta, e^*, \sigma_{\theta}^2) \frac{\partial e^*}{\partial \alpha} - C_e(e^*, \sigma^2) \frac{\partial e^*}{\partial \alpha}) = \frac{d}{d\alpha} R(e^*, \alpha, \sigma^2).$$

That is, the expected marginal benefit of incentives has to equal the expected marginal monetary and risk cost of incentives. We assume that the solution is unique and has $e^*(\alpha^*, \cdot) < e^{FB}(\cdot)$, so that the first-best is not achieved. Let α^* denote the solution.

2.3.1 Principal differences to the standard principal-agent model

While the setup of the model parallels closely that of a standard principal-agent model, there are two main differences that generate the ambiguous relationship between risk and incentives. First, the framework allows the agent to learn decision-relevant information (information about the realized θ and μ) between the signing of the contract and the action choice. When such learning takes place, the incentive contract needs to manage the agent's response to that information across all the possible outcomes. This aspect is different from the

standard agency setting, in which the incentive contract is typically intended to induce a given action that is known at the time of contracting.

Second, even when the goal of the contract is to induce a particular action (as is the case when the action in question is taken before the agent learns any additional information), uncertainty can influence the desired choice of that action through considerations over and above the risk-aversion of the agent. For example, if the purpose of the action is to acquire information in the first place, then the level of underlying uncertainty is going to influence the value of that information and, thus, the desired level of information acquisition. Also, the levels of uncertainty influence the implicit incentives provided through, for example, signal-jamming considerations, which are rarely embedded in models of formal incentive provision.

3 The Impact of Uncertainty on the Incentive Strength

Having reviewed the principal's problem, we can now look at what happens when we alter the volatility of any of the underlying parameters. That is, we are interested in the sign of $\partial\alpha^*/\partial\sigma_i^2$, where $i \in \{\theta, \mu, \varepsilon, \eta\}$. Now, by the assumption of global concavity of the program and the resulting unique α^* , we know from the implicit function theorem that

$$\text{sign}(\partial\alpha^*/\partial\sigma_i^2) = \text{sign}(dFOC_P(\alpha)/d\sigma_i^2).$$

Thus, the basic relationship between the strength of incentives and the level of uncertainty is simple: if the marginal benefit of incentives increases more with the level of uncertainty than the marginal cost of incentives does, then the relationship will be positive and vice versa. In other terms

$$\frac{\partial\alpha^*}{\partial\sigma_i^2} > 0 \quad \text{iff}$$

$$\frac{d}{d\sigma_i^2} [E[Y_e(e^*, \theta, \sigma_\theta^2)] - [C_e(e^*, \sigma^2)]] \frac{\partial e^*}{\partial\alpha} - \frac{d}{d\alpha} R(e^*, \alpha, \sigma^2) > 0.$$

Of course, this condition is not conveying much new information. However, it highlights the point that, in the end, the impact of uncertainty on the incentive strength has to depend on how changes in uncertainty influence the marginal cost and the marginal benefit of incentives. Consequently, thinking

more carefully about how uncertainty can impact the marginal benefit of incentives can be very helpful in illustrating the various settings in which the relationship between uncertainty and incentives can be positive.

Performing the differentiation gives us the expanded condition of

$$\begin{aligned}
& E \left[Y_{e\sigma_i^2} (e^*, \theta, \sigma_\theta^2) \frac{\partial e^*}{\partial \alpha} \right] - E \left[C_{e\sigma_i^2} (e^*, \sigma^2) \frac{\partial e^*}{\partial \alpha} \right] \\
& + \int_{\bar{i}}^{\underline{i}} \left[E_{\sigma^2 \setminus \sigma_i^2} [[Y_e (e^*, \theta, \sigma_\theta^2)] - [C_e (e^*, \sigma^2)]] \frac{\partial e^*}{\partial \alpha} \right] \frac{\partial f(i)}{\partial \sigma_i^2} \\
& + E \left[[Y_{ee} (e^*, \theta, \sigma_\theta^2) - C_{ee} (e^*, \sigma^2)] \frac{\partial e^*}{\partial \alpha} \frac{\partial e^*}{\partial \sigma_i^2} \right] \\
& + E \left[[Y_e (e^*, \theta, \sigma_\theta^2) - C_e (e^*, \sigma^2)] \frac{\partial^2 e^*}{\partial \alpha \partial \sigma_i^2} \right] \\
& - \frac{d}{d\alpha d\sigma_i^2} R (\alpha, e^*, \sigma^2) > 0.
\end{aligned}$$

Each line now captures in simple terms the different ways that uncertainty can impact the cost and benefit of incentives. The first line gives the *value effect* of uncertainty, where changes in the level of uncertainty directly alter either the marginal return or marginal cost of effort. If the return to effort is increasing in uncertainty, then the benefit from increasing incentives is naturally positive as well. The second line gives the *interaction effect* of uncertainty. Even if the *level* of uncertainty does not impact the action choice of the agent given his information, knowledge of the *realizations* of that uncertainty can alter the action choice of the agent. The distribution of uncertainty impacts the distribution of possible responses by the agent and, therefore, the expected value of incentives. The third line gives the *level effect* of uncertainty, where changes in uncertainty alter the expected level of effort provided by the agent. If an increase in uncertainty leads the agent to divert effort away from productive activities, then the benefit from increasing incentives is positive, other things constant. Finally, the fourth line gives the *sensitivity effect* of uncertainty, where uncertainty alters the responsiveness of the agent to incentives. If the agent becomes more responsive to incentives as uncertainty increases, then the benefit from increasing incentives is positive as well. If any of these components is positive, the relationship between uncertainty and incentives *can* be positive; we will discuss each channel in more detail below. However, note that the relationship *will* be positive if and only if the net effect of all four lines is sufficiently positive to overcome the increased cost of risk given by the fifth line.

3.1 The value effect

The first step of the differentiation is simply the direct impact that uncertainty has on the marginal value (or marginal cost) of effort:

$$E \left(Y_{e\sigma_i^2} (e^*, \theta, \sigma_\theta^2) - C_{e\sigma_i^2} (e^*, \sigma^2) \right) \frac{\partial e^*}{\partial \alpha}.$$

For all practical purposes, $\frac{\partial e^*}{\partial \alpha} > 0$. In consequence, the component can be positive if either $Y_{e\sigma_i^2}(\cdot) > 0$ or $C_{e\sigma_i^2}(\cdot) < 0$. In simple terms, the marginal benefit of incentives is increasing in uncertainty if either the marginal return to effort is increasing in uncertainty or the marginal cost of effort is decreasing in uncertainty. However, note that part of the value effect is directly internalized through the response of the agent. In particular, while the margin on the principal's payoff is only partially internalized because of agency problems ($\alpha^* < 1$), the agent fully internalizes the margin on marginal cost. Hence, while increases in the marginal productivity of effort can lead to increases in the strength of incentives, changes in the marginal cost of effort have a much smaller net impact and are more likely to be overturned by risk-compensation considerations. The value effect is thus primarily associated with productive uncertainty.

3.1.1 Examples

A natural application of the value effect is provided by the problem of motivating information acquisition. Of course, in this setting the information acquired has to have some impact on some decision made by either the principal or the agent, but it is possible to interpret the payoff functions as the reduced-form setting where the equilibrium decision rule has been already substituted in.⁷ In this case, $C_{e\sigma_i^2} = 0$ but, depending on the specific setting, $Y_{e\sigma_i^2}(e^*, \theta, \sigma_\theta^2)$ can be positive. (Of course, by definition, no other uncertainty enters the production function.)

A simple example is provided by a setting in which the agent is tasked to screen the quality of projects. Consider the following setting: A risk-neutral

⁷Of course, there is nothing that says that the value of the agent's effort or the cost of providing that effort could not be environment-specific. Indeed, casual empiricism suggests that how much a manager is able to contribute to the performance of a firm can be highly environment-specific. However, controlling for such heterogeneity is primarily an empirical problem.

principal hires a risk-averse agent with mean-variance preferences (with risk aversion r) to acquire information about the quality of a project. The expected return can be either y or $-y$, and both are ex ante equally probable. The ex ante variance of the expected return (σ_θ^2) is, thus, proportional to y^2 . In addition to the underlying quality of the project, the return is subject to an exogenous shock $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. After investing $C(p)$ in information acquisition, the agent receives a signal s , which is correct with probability $p \geq 1/2$, about the quality of the project. The principal retains the right to choose when to implement the project and, thus, implements the project whenever the signal indicates a high-return project. Finally, we assume that, while the principal can interpret the signal, it is not verifiable by an outside court, and, therefore, the agent has to be incentivized through tying his compensation on the realized payoff.

The expected return to the principal is then given by

$$py + (1 - p)(-y) = (2p - 1)y,$$

which is increasing in the variance of project quality y . Thus, the relationship between incentives and y can be positive.⁸

The same logic continues to hold in more elaborate settings of information acquisition. For example, Shi (2007) provides a model in which the productivity of investment is ex ante uncertain and the level of investment is conditioned on the information learned in the acquisition stage. Again, even if the agent is risk-averse, the relationship between uncertainty over the productivity of investment and the strength of incentives is positive for a range of parameter values. Rantakari (2007) discusses decision-making in the context of quadratic loss functions with similar conclusions. Finally, while the value effect tends to be positive in models of information acquisition, this need not always be so because the value of information is dependent on how it is used. For example,

⁸Given a share α of the project return, the agent's utility is then given by

$$\alpha(2p - 1)y - C(p) - r\alpha^2 \left(p(1 - p)(2y)^2 + \sigma_\varepsilon^2 \right),$$

which gives p^* as the solution to $2\alpha y + r\alpha^2(2p - 1)(2y)^2 = C'(p)$. (Note that, because the agent is risk-averse, the presence of uncertainty provides implicit incentives to the agent because information reduces the variance of the realized return.) The surplus maximization problem, in turn, is then

$$\max_{\alpha} (2p(\alpha) - 1)y - C(p(\alpha)) - r\alpha^2 \left(p(\alpha)(1 - p(\alpha))(2y)^2 + \sigma_\varepsilon^2 \right).$$

If σ_ε^2 is large, then α will be small, and $\partial\alpha^*/\partial y > 0$.

if the agent is risk-averse, he can become more conservative in his use of information as the level of productive uncertainty goes up, resulting in a decreased marginal value of effort.⁹

3.2 The interaction effect

The second step of the differentiation is the interaction effect. Even if, given his information, the agent's action choice is independent of the levels of uncertainty, as long as the agent's action choice is a function of his information, changes in the amount of uncertainty change the distribution of outcomes and, thus, the expected value of incentives. This effect is captured by the second line

$$\int_{\underline{i}}^{\bar{i}} E_{\sigma^2 \setminus \sigma_i^2} \left[[Y_e(e^*, \theta, \sigma_\theta^2) - C_e(e^*, \sigma^2)] \frac{\partial e^*}{\partial \alpha} \right] \frac{\partial f(i)}{\partial \sigma_i^2}.$$

By definition, this effect exists only if the agent's action choice is a function of the acquired information, so it can be present only for μ and θ .

The logic behind this component is as follows. Because the agent responds to information, the gap between the marginal return and the marginal cost of effort is going to vary with the specific realizations of the shocks. Weighing each realization by the responsiveness of the agent to incentives $\partial e^* / \partial \alpha$ gives the state-contingent marginal benefit of incentives, and taking the expectation over the states yields the expected marginal benefit of incentives. Changing the underlying distributions changes the weight placed on each conditional outcome and, therefore, the expected marginal benefit of incentives. If the expected marginal benefit increases with uncertainty, then the relationship between uncertainty and incentives can be positive.

While somewhat cumbersome as an expression, the basic logic is directly related to Jensen's inequality. If the state-contingent marginal benefit is convex in the state, then the expected benefit from increasing incentives is increasing in uncertainty. Conversely, if the state-contingent marginal benefit is concave in the state, then the expected benefit from increasing incentives is decreasing in uncertainty. In most cases, this condition is negative for σ_μ^2 but usually positive for σ_θ^2 .

⁹Similarly, if an organization responds to increased uncertainty by choosing a production technology that is more robust to "mistakes," the value of information can go down.

3.2.1 Examples

There are two slightly different effects that can be at work here. The first effect is the performance measurement effect analyzed in Baker (1992) and relates to how the correlation between the action induced by the performance measure and the principal's preferred action varies with uncertainty. To briefly summarize the result, recall that, if the agent has perfect interim information and there are no additional considerations, we can approximate the relevant condition by¹⁰

$$\frac{\partial}{\partial \sigma_i^2} [Cov(Y_e, X_e) - \alpha Var(X_e)].$$

Thus, as σ_μ^2 , which enters only X_e , increases, the misalignment problem becomes more severe, and the marginal benefit of incentives goes down. Conversely, if an increase in σ_θ^2 makes the performance measure relatively more aligned with the principal's preferences, then the marginal benefit of incentives goes up.

If the agent has imperfect information about the environment, the distinction need not be as clear cut. This result follows because, when forming expectations about X_e , the agent utilizes his information about both θ and μ . If an increase in σ_μ^2 makes the agent more reliant on his information about θ , then the agent can become more aligned with the principal's preferences, and the marginal benefit of incentives goes up. In summary, the conclusion from the performance measurement perspective is that $\partial\alpha^*/\partial\sigma_\theta^2$ can be positive (and will typically be if the agent is risk-neutral) while $\partial\alpha^*/\partial\sigma_\mu^2$ is typically negative, although this need not always be the case.

The second effect is the variance in the action choice of the agent caused by θ and arises even in the absence of σ_μ^2 .¹¹ It is this effect that is the driving force behind the recent models of random productivity that have largely abstracted away from misaligned performance measures ($X = Y$) and focused on analyzing the risk-incentives trade-off. Zabochnik (1996) and Baker and Jorgensen (2003) consider a production function given by $Y = \theta e + \varepsilon$ but treat the agent's risk aversion differently. Raith (2007) considers an analogous setting but with risk-neutrality, limited liability, two effort dimensions, and binomial output.¹²

¹⁰We have dropped the constants that are independent of σ_i^2

¹¹This effect is inherently embedded in the performance measurement problem and, therefore, also present in Baker (1992) but not explicitly analyzed.

¹²The setting in Shi (2007) is somewhat different. The output is given by $y = \theta I - I^2/2$ and the agent bears no direct cost in choosing the investment level. Given his quality of

Following the above, assume that the output is given by $Y = \theta e + \varepsilon$ and is measurable, the cost of effort is $e^2/2$, and the agent learns θ perfectly before his choice of effort. Then, the agent will choose $e = \alpha\theta$, and the state-contingent marginal benefit $[Y_e(e^*, \theta, \sigma_\theta^2) - C_e(e^*, \sigma^2)] \frac{\partial e^*}{\partial \alpha}$ is simply $(1 - \alpha)\theta^2$, which is convex in the state and so suggests that the marginal benefit from increasing incentives is positive. The standard intuition is straightforward: the higher the productive uncertainty, the higher the value of inducing the agent to respond to the realizations of that uncertainty.

Note that while the explicit models to date have focused on situations where $Y_{e\theta} > 0$, satisfaction of this condition is by no means necessary for the value of incentives to increase in productive uncertainty. As shown above, the relevant condition is that the state-contingent marginal benefit is convex in the state. To illustrate with a simple example, let $X = Y = \Pi(e, \theta)$ and $B = 0$. Then, $\frac{\partial e^*}{\partial \alpha} = \frac{Y_e}{C_{ee} - \alpha Y_{ee}}$ and the condition can be written as

$$(1 - \alpha) \int_{\underline{\theta}}^{\bar{\theta}} \frac{Y_e^2}{C_{ee} - \alpha Y_{ee}} \frac{\partial f(\theta)}{\partial \sigma_\theta^2}.$$

This condition can equally well be satisfied when $Y_{e\theta} < 0$, so that the effort and the state are substitutes.¹³ For example, a high-quality project can be very successful but hard for the agent to improve upon, while a bad project might have a low baseline return but improve quickly through the effort of the agent. An example of such a setting could be given by a production function

$$Y = x\theta + (\gamma - \theta)e + \varepsilon,$$

where $\bar{\theta} \leq \gamma$. Now, if we assume that the agent's effort cost is $e^2/2$, then $C_{ee} = 1$, $Y_{ee} = 0$ and $Y_e^2 = (\gamma - \theta)^2$. Thus,

$$\frac{Y_e^2}{C_{ee} - \alpha Y_{ee}} = (\gamma - \theta)^2$$

is convex, and the marginal benefit of incentives is increasing in uncertainty. Finally, it should be possible to construct examples in which the condition is

information, first-best investment level could be achieved by simply letting $\alpha \rightarrow 0$. However, $\alpha > 0$ is needed to motivate information acquisition as discussed above. This is also the reason why Shi (2007) is more an example of the value effect, since the decision-making problem itself would be trivial if information was exogenously given.

¹³Of course, this is just a re-labeling of the state but works to highlight that the effect is not inherently tied to the specific examples.

violated so that an increase in productive uncertainty would lead to a decreasing marginal benefit of incentives. Such a setting could be a situation in which the return to both high- and low-quality projects is largely beyond the control of the agent, but intermediate projects can be improved upon relatively easily.

3.3 The level effect

The third step of the differentiation is to account for any changes in the level of effort provided by the agent and its consequences. This effect is given by

$$E \left[[Y_{ee}(e^*, \theta, \sigma_\theta^2) - C_{ee}(e^*, \boldsymbol{\sigma}^2)] \frac{\partial e^*}{\partial \alpha} \frac{\partial e^*}{\partial \sigma_i^2} \right].$$

By assumption, $Y_{ee}(e^*, \theta, \sigma_\theta^2) \leq 0$ and $C_{ee}(e^*, \boldsymbol{\sigma}^2) > 0$. Thus, the first component is negative. Similarly, $\partial e^*/\partial \alpha > 0$. Thus, this effect is positive if $\partial e^*/\partial \sigma_i^2 < 0$. In other words, the value to increasing incentives can be positive if an increase in uncertainty leads the agent to divert effort away from productive activities. Intuitively, if the agent cuts back his equilibrium effort, the principal can find it valuable to compensate for this diversion by increasing the strength of incentives.

Note that this channel is conceptually different from the value effect. As discussed above, when the value effect is positive, the agent naturally internalizes a part of that change through the incentive contract. However, the effect is in the opposite direction ($\partial e^*/\partial \sigma_i^2 > 0$), and because the internalization is not perfect, the impact on incentive strength can remain positive. In contrast, in the present setting an increase in uncertainty needs to *decrease* the level of effort provided.

Recall that the agent's first-order condition was given by

$$E_A(\alpha X_e + B_e | \omega) - r_e(e^*, \alpha, \boldsymbol{\sigma}^2) = C_e.$$

Thus, there are three broad categories of effects that can lead the agent to divert effort away from productive activities: (i) the perceived marginal benefit of productive effort decreases in uncertainty, (ii) an increase in uncertainty causes the agent to reduce productive effort because of its impact on risk, and (iii) the performance measure becomes increasingly manipulable when uncertainty increases. The last avenue, however, cannot be fixed by increasing the weight on the performance measure because such a change would only exacerbate the problem. Similarly, the second avenue is somewhat tenuous absent other constraints because, as with the variation in marginal cost, the agent

internalizes the full margin on risk. (However, when discussing the sensitivity effect below, we do give an example of a setting in which the interaction of level and sensitivity effects can lead to a positive relationship between uncertainty and incentives.) Below we will examine the first avenue.

3.3.1 Examples

The most natural setting for considering a situation in which the perceived benefit of effort is impacted by uncertainty over and above its direct impact on the productive environment (Y, X, C, R) is provided by settings in which the agency problem is coupled with an inference problem by the principal, thereby adding signal-jamming (or signaling, as long as information is learned after the signing of the contract) considerations to the picture. An example of the logic of this setting is provided by the career concerns model analyzed in Gibbons and Murphy (1992). Consider a two-period model in which the per-period output by the agent is given by

$$y_t = \theta + e_t + \varepsilon_t,$$

where θ is the agent's ability (most analogous to productive uncertainty), e_t is the agent's effort choice in period t , and ε_t is a stochastic noise term. Both the ability and the shock are normally distributed, with mean ability θ . The labor markets are competitive, the employers are risk-neutral, and the agents are risk-averse with exponential utility. Both sides of the market have initially no information about θ .

In period 2, competition in the labor market will drive the compensation of the agent to move one-to-one with the market's expectation over the agent's ability θ . This expectation is given by

$$E(\theta|y_1) = \frac{\sigma_\varepsilon^2 \tilde{\theta} + \sigma_\theta^2 (y_1 - E e_1)}{\sigma_\varepsilon^2 + \sigma_\theta^2}.$$

To form expectations over θ , the market needs to form expectations over e_1 . If $E e_1$ is "too low," then the agent has incentives to exceed these expectations to fool the market and to increase his second-period wage. (Of course, in equilibrium this doesn't work.) Note that the higher σ_ε^2 , the less effective the agent's effort is in impacting the market's beliefs, and similarly, the higher σ_θ^2 , the more responsive the market is. Thus, an increase in σ_θ^2 will increase the implicit incentives provided by the market, and the strength of explicit incentives can be decreased. Conversely, an increase in σ_ε^2 will decrease the

implicit incentives and can lead to a positive relationship between uncertainty and explicit incentives. The conclusion is, therefore, roughly the converse of our earlier discussion.

As shown in Gibbons and Murphy (1992), the first-period strength of incentives (with a discount factor δ) is given by

$$\alpha_1^* = \frac{1}{1+r(\sigma_\theta^2+\sigma_\varepsilon^2)C''(e_1(\alpha_1^*))} - \delta(1-\alpha_2^*)\frac{\sigma_\theta^2}{(\sigma_\varepsilon^2+\sigma_\theta^2)} - \frac{r\delta\alpha_2^*\sigma_\varepsilon^2C''(e_1(\alpha_1^*))}{(1+r(\sigma_\theta^2+\sigma_\varepsilon^2))C''(e_1(\alpha_1^*))},$$

where α_2^* is the second-period strength of incentives. While, in this particular setting, risk-aversion considerations tend to dominate the choice of incentives, it is possible to find some configurations where the decrease in implicit incentives (given by the second component) dominates and $\partial\alpha_1^*/\partial\sigma_\varepsilon^2 > 0$, while $\partial\alpha_1^*/\partial\sigma_\theta^2 < 0$ is always the case.¹⁴

While this is the only setting that we are aware of in which the strength of incentives has been explicitly derived, the signal-jamming or signaling logic is directly applicable to many other settings, for example, the general logic of influence activities (Milgrom and Roberts 1988). Given the available signals, the principal can try to make decisions about capital allocation, promotion decisions, and the like. If the signals can be influenced by the agent, their presence provides implicit incentives to alter the levels of effort. However, the specifics are going to be highly context-specific. For example, if the principal is determining the capital allocation based on the realized X and Y , then the analysis would be similar to the above, and the presence of influence activities would be generally beneficial. If, on the other hand, there are signals that are manipulable but orthogonal to immediate production, then effort would be diverted away, and influence activities would be bad for the overall performance. Finally, as already mentioned above, if the performance measure itself is manipulable and increased uncertainty leads to increased manipulation instead of productive effort, then increasing the strength of incentives is, of course, not helpful in restoring the balance.

¹⁴Also, increasing the number of periods naturally tends to load more weight on the career concerns component.

3.4 The sensitivity effect

The fourth and final line of the differentiation is to account for any changes in the responsiveness of the agent to incentives. This effect was given by

$$E \left[[Y_e(e^*, \theta, \sigma_\theta^2) - C_e(e^*, \sigma^2)] \frac{\partial^2 e^*}{\partial \alpha \partial \sigma_i^2} \right].$$

By construction, the expected gap between the marginal productivity and the marginal cost of effort is predominantly positive.¹⁵ As a result, the component is positive if $\frac{\partial^2 e^*}{\partial \alpha \partial \sigma_\theta^2} > 0$. That is, the relationship between uncertainty and incentives can be positive if the responsiveness of the agent to incentives increases with uncertainty. Intuitively, the more responsive the agent is to incentives, the more "bang for a buck" the principal is able to obtain from a given increase in incentives.

As an illustration, letting $X = Y$ and $B = 0$, the agent's sensitivity to incentives is given by

$$\frac{\partial e^*}{\partial \alpha} = \frac{EY_e - r_{e\alpha}}{C_{ee} + r_{ee} - \alpha EY_{ee}}.$$

Thus, changes in the marginal value of effort naturally lead to changes in responsiveness, but this change is linked to the value effect. Thus, the primary determinants of sensitivity are the concavity of the production function and the convexity of the cost function, coupled with any changes in the agent's response to risk. While such functions can be highly environment-specific (in particular, the concavity of the production function resonates well with how much the agent is able to influence the environment), little to no formal theory has examined this relationship. However, considerations of job design give rise to endogenous variations in C_{ee} , and the sensitivity effect can also arise in conjunction with the other effects, either complementing or countering the initial effect.

¹⁵When $\sigma_\mu^2 > 0$, then knowledge of μ and the resulting response can sometimes lead to $Y_e < C_e$. However, α is then adjusted accordingly to limit such gaming of the contract, as given by the interaction effect.

3.4.1 Examples

If an increase in uncertainty increases the value of the *responsiveness* of the agent (as when the interaction effect is present), then job design can be used to complement any changes in the preferred strength of incentives. A simple example of this joint effect is given in Rantakari (2007). Consider a situation where the principal hires an agent to make a decision, d . The payoff to the principal is given by $K - (\theta - d)^2$, while the cost of the decision to the agent is d^2 and the agent is risk-neutral. The principal's payoff, however, is not directly contractible. Instead, there exists a performance measure $s = \theta + \mu$ that imperfectly tracks the preferences of the principal and is perfectly observed by the agent before making his decision. Both θ and μ are normally distributed, with $E\theta = \tilde{\theta}$ and $E\mu = 0$. Then, as shown in the paper, second-best decisions can be implemented through a linear contract

$$T = K - \alpha_P (s - d)^2 - \alpha_d (\bar{d} - d)^2,$$

where α_P is the strength of incentives, \bar{d} is a default decision specified by the contract and α_d is the penalty rate for deviating from this decision (and, thus, the inverse of job flexibility). The solution to the optimization problem gives

$$\alpha_P = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_\mu^2}, \quad \alpha_d = \frac{\sigma_\mu^2}{\sigma_\theta^2 + \sigma_\mu^2} \quad \text{and} \quad \bar{d} = \tilde{\theta}.$$

Thus, the more volatile the environment relative to the risk of gaming by the agent, the stronger the incentives to encourage responsiveness by the agent to incoming information. At the same time, as the value of responsiveness increases, the agent is given greater freedom to deviate from the default decision. Relatedly, Holmström and Milgrom (1991) show that job flexibility is decreasing in σ_ε^2 . In general, the value of responsiveness is decreasing in σ_μ^2 and in σ_ε^2 because less flexibility limits the opportunities for gaming and can help to shield the agent from income risk, but it can be increasing in σ_θ^2 to support flexible reactions to information about θ .

As a final example, we will consider a setting in which the interaction between the level and sensitivity effects can lead to a positive relationship between uncertainty and incentives, even in the absence of external considerations (B_e) and direct value effects. The output of the principal is given by $Y = \gamma e_1$, where γ is a fixed constant. The productivity could be random, but because such randomness is not needed to generate the desired result, we assume that it is not. The performance measure is given by $X = \mu e_1$,

where μ is normally distributed and $E\mu = \gamma$. The agent receives no information about μ before making his effort choice, faces a cost $e_i^2/2$ of effort, and is risk-averse with mean-variance preferences and a risk-aversion coefficient $r/2$. In addition to increasing the output, the agent can invest in reducing the noisiness of the performance measure. In particular, the residual noise is given by $(e_1 - e_2)^2 \sigma_\mu^2$. While productive effort amplifies uncertainty, the second dimension can be used to control it.

Given the strength of incentives α , the agent solves

$$\min_{e_1, e_2} \alpha\gamma e_1 - \frac{e_1^2}{2} - \frac{e_2^2}{2} - \frac{r}{2}\alpha^2 (e_1 - e_2)^2 \sigma_\mu^2,$$

the solution to which is given by

$$e_1 = \frac{\alpha\gamma(1+r\alpha^2\sigma_\mu^2)}{1+2r\alpha^2\sigma_\mu^2} \quad \text{and} \quad e_2 = \frac{\gamma r\alpha^3\sigma_\mu^2}{(1+2r\alpha^2\sigma_\mu^2)}.$$

An increase in σ_μ^2 leads to a decrease in e_1 (because effort amplifies return volatility) and an increase in e_2 (to counter the increase in volatility). Effort has, thus, been diverted away from production. Even if C_{ee} is constant in our setting, the response through e_2 that results from the risk-aversion of the agent leads to variations in the responsiveness of the agent to incentives. In particular, we have that

$$\frac{\partial^2 e_1^*}{\partial\alpha\partial\sigma_\mu^2} > 0 \quad \text{iff} \quad 2r\alpha^2\sigma_\mu^2 > 3,$$

suggesting that the agent becomes more responsive to incentives only when the initial noise is sufficiently high (because that is when the response through e_2 is large). Thus, even if incentives are now costlier to provide because of the higher risk, the agent is now also more responsive to incentives. Indeed, while solving the model analytically is infeasible, a simple numerical solution shows that $\partial\alpha^*/\partial\sigma_\mu^2 > 0$ if $\sigma_\mu^2 \geq \bar{\sigma}_\mu^2(r, \gamma)$.

4 Summary of the effects and "types" of uncertainty

This paper has illustrated the various channels through which uncertainty can impact the equilibrium strength of incentives. By using a simple agency model as the guiding tool, we categorized these channels based on the impact that uncertainty has on the action choice preferred by the principal and the action taken by the agent. We then used these categories to derive conditions under which the relationship between uncertainty and incentives can be positive and to discuss the existing literature.

The channels and conditions were (i) the value effect, which was positive whenever the marginal value of the agent's effort increased with uncertainty; (ii) the interaction effect, which was positive when an increase in uncertainty altered the distribution of final outcomes in a way that increased the expected marginal benefit of incentives increased; (iii) the level effect, which was positive when an increase in uncertainty decreased the level of productive effort provided by the agent; and (iv) the sensitivity effect, which was positive when an increase in uncertainty made the agent more responsive to incentives. Each condition could lead to a positive relationship between uncertainty and incentives, as illustrated by existing models and simple examples.

We refrained from categorizing uncertainty based on its "type," with the exception of isolating productive uncertainty from the other categories. The reason for doing so was the fact that, given the right setting, the relationship between any type of uncertainty and the strength of incentives could be positive. However, there are clear associations between the effects discussed here and "types" of uncertainty, as determined by where they enter the contracting environment and whether they are observable by the agent.

Consider first the additive shocks η and ε . Because neither their realization nor their level impacts the action desired by the principal or the action taken by the agent, an increase in σ_η^2 or σ_ε^2 can lead to an increase in incentives only when the basic agency setting is coupled with an inference problem by the principal, such as career concerns or influence activities. However, to the extent that σ_η^2 or σ_ε^2 weakens the implicit incentives, the strength of explicit incentives can be used to counter this effect. Thus, these shocks are only associated with the level effect on effort. Otherwise, σ_η^2 plays no role in the analysis as it does not enter the agent's direct compensation while σ_ε^2 only increases the cost of incentives through the risk-aversion of the agent.

The roles of θ and μ , on the other hand, are somewhat more complex. Consider first the role of σ_μ^2 . If μ is perfectly observable before the action

choice, then an increase in σ_μ^2 leads to weaker incentives because it simply leads to increased gaming of the contract by the agent. However, if μ is only partially observed, then there are at least two possible reasons for the relationship between σ_μ^2 and incentives to be positive. First, an increase in σ_μ^2 , by interacting with the agent's action choice, can lead the agent to divert effort away from productive activities, thus generating a positive level effect of uncertainty. Second, an increase in σ_μ^2 can make the agent more reliant on his information about θ with respect to the marginal return to his action. Such a response can increase the alignment of his action with the preferences of the principal, generating a positive interaction effect.

The most prominent component in the analysis is naturally θ , which can play a role in any of the four effects. First, whenever the agent learns information about θ , increasing incentives can be valuable in increasing the agent's response to that information when the interaction effect is positive. This is the effect that is the driving force behind the various recent random productivity models. Second, as long as some information is learned by the agent, increased responsiveness can be encouraged through the sensitivity effect by increasing the job flexibility of the agent. Third, given that information about θ is imperfect, σ_θ^2 can function through the level effect much like σ_μ^2 . Finally, the value of the agent's effort can be directly increasing in σ_θ^2 , leading to a positive value effect. For example, even if solving the problem of using given information were trivial, incentives could be needed to motivate the generation of that information.

We hope that our analysis has helped to clarify the basic relationship between uncertainty and incentives. Two dimensions, however, remain relatively unexamined. First, beyond simple examples, the joint problem of motivating the acquisition and use of information has received only limited attention in the literature to date. Second, when and how much the expected outcome can be improved upon with the help of non-linear incentives remains an open question. Both dimensions appear promising avenues for future research.

References

- [1] Baker, G. (1992), "Incentive Contracts and Performance Measurement," *Journal of Political Economy* 100:598-614
- [2] Baker, G. (2002), "Distortion and Risk in Optimal Incentive Contracts," *Journal of Human Resources* 37(4):728-51
- [3] Baker, G. and B. Jorgensen (2003), "Volatility, Noise and Incentives," mimeo, Harvard University and Columbia University
- [4] Gibbons, R. and K. Murphy (1992), "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, 100(3):468-505
- [5] Holmström, B. (1979), "Moral Hazard and Observability," *Bell Journal of Economics*, 10:74-91
- [6] Holmström, B. and P. Milgrom (1991), "Multi-task Principal-Agent Analyses: Incentive Contracts, Asset Ownership and Job Design," *Journal of Law, Economics and Organization* 7(0):24-52
- [7] Holmström, B. and P. Milgrom (1994), "The Firm as an Incentive System," *American Economic Review*, 84(4):972-991
- [8] Jensen, M. and W. Meckling (1976), "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure," *Journal of Financial Economics*, 3:305-360
- [9] Jensen, M. and W. Meckling (1992), "Specific and General Knowledge and Organizational Structure," in L. Werin and H. Wijkander (eds.) *Contract Economics*, Oxford: Blackwell
- [10] Milgrom, P. and J. Roberts (1988), "An Economic Approach to Influence Activities in Organizations," *American Journal of Sociology*, 94:154-179
- [11] Prendergast, C. (1999), "Provision of Incentives in Firms," *Journal of Economic Literature*, 38:7-63
- [12] Prendergast, C. (2002), "The Tenuous Trade-off Between Risk and Incentives," *Journal of Political Economy*, 110(5):1071-1102
- [13] Raith, M. (2007), "Specific Knowledge and Performance Measurement," mimeo, University of Rochester

- [14] Rantakari, H (2007), "Uncertainty, Delegation and Incentives," mimeo, University of Southern California
- [15] Shi, L. (2007), "Responsible Risk and Incentives for CEOs: the Role of Information-collection and Decision-making," mimeo, University of Washington
- [16] Simon, H. (1951), "A Formal Theory of Employment Relationship," *Econometrica*, 19(3):293-305
- [17] Zabochnik, J. (1996), "Pay-performance sensitivity and production uncertainty," *Economics Letters*, 53(3):291-296