USC FBE APPLIED ECONOMICS/CLEO WORKSHOP
presented by Michael Raith
FRIDAY, October 24, 2003
1:30 pm - 3:00 pm; **Room: HOH-601K**

Specific Knowledge and Input- vs. Output-Based Incentives

Michael Raith

University of Rochester and CEPR *

First version: May 2002
This version: October 2003

Preliminary and incomplete; comments welcome.

*Correspondence: Michael Raith, Simon School of Business, University of Rochester, Rochester, NY 14627, USA; email: raith@simon.rochester.edu. This paper was presented under the title "Implicit Delegation through Performance Pay" at the CEPR workshop "Incentives and Beyond-The Economics of Personnel and Organizations" in Stockholm in May 2002. I would like to thank Jim Brickley, Francine Lafontaine, Canice Prendergast, Andreas Roider, Lars Stole, and seminar participants in Stockholm for very helpful comments.

Specific Knowledge and Input- vs. Output-Based Incentives

Abstract

This paper examines optimal incentive contracts in which an agent's compensation can be based on both "input" measures closely related to an agent's effort, and "output" measures closely related to the principal's payoff. I argue that when the agent has specific knowledge (i.e. private information that is difficult to communicate) about how his actions contribute to the principal's payoff, output-based pay encourages the agent to use his knowledge while input-based pay does not. I show within a two-task agency model that partially output-based compensation is optimal even when the agent's effort on each task can be measured perfectly. I also study how changes in different parameters affect the optimal balance of input- vs. output-based pay. In particular, an increase in the variance of measured output that is unrelated to the agent's productivity leads to less output-based pay, consistent with standard agency models. In contrast, an increase in uncertainty about the productivities of the tasks leads to more output-based pay, similar to the predictions of recent models of delegation.

JEL codes:

Keywords: incentives, risk, multi-task agency model, input- vs. output-based pay, delegation, distortion, congruence

1 Introduction

One of the most important yet most difficult elements in designing effective organizations is the measurement of performance. Many organizations today recognize the usefulness of performance-based pay as a general principle, but in practice struggle to find the right measures on which to base compensation.

Principal-agent theory has identified two principles that should guide the choice of performance measures. The first and most obvious is that an agent's incentives should be aligned with the principal's objectives. The second is the Informativeness principle: when agents are risk-averse, performance measures should be chosen to measure an agent's effort as accurately as possible.

As is now well understood, these two principles may be difficult to reconcile. Using performance measures closely related to an agent's effort (which I will refer to as "input" measures) minimizes the agent's exposure to risk, but may also fail to measure some dimensions critical to the job and hence lead the agent to allocate his effort in unproductive ways. Measures more closely correlated with organizational goals ("output" measures), on the other hand, are often subject to influences beyond an employee's control; i.e. are risky.[1]

Yet while theory has emphasized measurement problems, firms have become quite adept at quantifying a large number of relevant dimensions of the performance of individuals and sub-units. Progress in implementing activity-based costing, balanced scorecards, and integrated information systems are only a few examples of this trend (for an overview of current trends, see Ittner and Larcker, 1998). Nowadays, while measurement problems of course remain, firms often appear to have a much greater range of performance measures at their hands than they know how to make use of.

The experience of many firms suggests that a central assumption of the theory is violated in practice: namely that firms know how their employees' activities contribute to the firm's bottom line. In contrast, the premise of this paper is that even if firms can accurately measure every dimension of performance they deem important, they may still

[1] Cf. Holmström and Milgrom (1991), Baker (1992), Feltham and Xie (1994), Baker (2000).

know only little about how these measures are related to organizational goals. Moreover, employees often have a better knowledge of how the different tasks they perform contribute to organizational goals than those who design their compensation plans. Their knowledge is often *specific* in the sense of Jensen and Meckling (1992) because it is difficult to communicate to others.

An example illustrates this assumption and its implications. In many firms, modern sales information systems combined with customer surveys enable firms to measure, at low cost, practically all relevant dimensions of a salesperson's performance: hours worked, number of customers contacted, new accounts created, quality of advice, etc. These (input) measures combined provide a much more accurate picture of a salesperson's effort than an output measure such as sales. But basing compensation on some formula combining these measures would direct the salesperson to allocate his effort in a very particular way, which would be unwise if the salesperson has a better understanding than the firm of how to allocate time and effort to promote the firm's goals. Basing compensation on output measures such as sales, in contrast, exposes the salesperson to greater income risk but also encourages him to allocate time and effort optimally.

The central tradeoff for the firm is thus between output measures that encourage the optimal use of specific knowledge but are risky, and input measures that minimize risk but discourage the use of specific knowledge. In general, the optimal compensation plan will be based partly on input, and partly on output measures.

I study the optimal choice between input- and output-based compensation in a two-task model in which a principal faces uncertainty about the productivities of different tasks, and in which an agent has private information about these productivities. The agent is risk-neutral but protected by limited liability, which implies that it is costly for the principal to expose the agent to risk. The principal can compensate the agent based on (verifiable) information about the agent's effort (i.e. his input), and on a noisy measure of realized output. Both real and measured output depend on the realized productivities of the tasks, the agent's effort devoted to each task, and some exogenous noise. I shall refer to the fluctuations in the productivities of tasks as *technological uncertainty*, and the noise in measured output that is unrelated to productivity as *environmental uncertainty* or sometimes *risk*.

With no technological uncertainty or private information, the principal always prefers input-based pay because output-based pay is risky. In contrast, if the agent has private information about the productivities of the tasks, it is optimal for the principal to let the agent decide how to allocate his effort across tasks, and to base his compensation on observed output as well. The optimal compensation function is in part input-, in part output-based. The optimal relative weight of output- vs. input-based pay depends on the parameters of the model as follows:

1. The better the agent's knowledge, the more compensation will be output-based. This intuitive prediction is in line with recent models of delegation, but relates to the choice of performance measures, *given* delegation. The result is quite different in spirit to that of Baker (1992), who rules out output-based compensation and finds that the better the agent is informed, the weaker incentives the principal will overall provide in order to prevent gaming.

2. The greater the environmental uncertainty, the more compensation will be input-based, i.e. the less the allocation of effort will be delegated to the agent. This negative relationship between risk and incentives is familiar from standard principal-agent theory, cf. Holmström (1979). The greater the technological uncertainty, on the other hand, the more compensation will be output-based. Since the agent has private information about the productivities of the tasks, a higher variance of the productivities implies a higher option value of the agent's information, as in Prendergast (2002). To benefit from more valuable information, it is optimal for the principal to put a greater weight on output-based pay. Thus both a standard negative relationship between uncertainty and incentives, and a positive relationship that is consistent with many empirical studies (cf. Prendergast, 2002), arise within the same model, depending on the nature of uncertainty that is varied.[2]

[2] The latter point is also emphasized in Baker and Jorgensen (2003); see Section 4.2 for a more detailed discussion.

3. A lower correlation between the tasks' productivities means that it is more likely that the agent should focus his effort on one task rather than both or none, and can therefore be interpreted as a job with more *task variety.* Since with greater task variety, the agent's information is more valuable for the principal, the principal prefers more output-based pay. Thus, it follows that the weight on output measures in compensation plans should be positively related to measures of job complexity.

4. In the model, all parameter changes that for given incentives increase expected output or total surplus, also lead to a greater weight on output and a lower weight on input in the agent's optimal contract. In this sense, a greater alignment of the agent's with the principal's objectives will lead to a greater reliance on output-based pay. Formally different but economically similar results hold in the models of Feltham and Xie (1994) and Baker (2000). In contrast, in recent formal models of delegation, there is no clear relationship between the alignment of the principal's and the agent's preferences and optimal delegation.

This paper takes well-known connections between specific knowledge, delegation and incentives one step further. The importance of dispersed information and its implications for economic organization were first emphasized by Hayek (1945). Applying this insight to the internal organization of firms, Jensen and Meckling (1992) argued that specific knowledge, i.e. knowledge located at lower levels of a firm that is difficult to communicate upwards, is a primary reason for decentralizing decisions. Brickley, Smith and Zimmerman (1995) have emphasized that when firms delegate decision rights to managers with specific knowledge, they must also measure the managers' performance and provide them with incentives to use their knowledge in the interests of the firm. For parallel arguments in the formal theoretical literature, see e.g. Holmström and Milgrom (1994) and Prendergast (2003).

So far, however, these insights provide little guidance on how to design performance measurement and incentives when there are many ways to measure performance. It may be obvious, for instance, that when salespeople have specific knowledge about customers, market conditions etc., it is optimal to let them decide how to go about their work, which in turn implies that they must be given incentives to put in effort and use their information optimally. The harder problem in practice is *how* to design optimal incentives.

The contribution of this paper is to show that an agent's specific knowledge, which may be a reason why the principal delegates decisions to him in the first place, also affects the optimal choice between input-and output-related compensation. Ignoring problems with measuring certain activities, input-based compensation is optimal when it is impractical to directly supervise the agent's activities, but where the principal knows how the agent *should* spend his time. Output-based compensation, in contrast, is optimal when the agent understands his job better than the principal does, and must be given incentives to use his knowledge optimally. In general, the optimal incentive contract depends on both kinds of measures.

The arguments of this paper directly relate to the question of how firms should use of modern tools for performance measurement and management such as the balanced scorecard and EVA. The balanced scorecard provides a wide array of financial and non-financial performance measures, whereas EVA provides a single overall measure of value creation. Thus, in the terminology of this paper, the balanced scorecard provides input measures while EVA is an output measure. According to existing theory, compensation should be (partially) based on EVA if using the scorecard alone would lead to dysfunctional behavior. This could occur either because critical dimensions of performance cannot be measured (Holmström and Milgrom, 1991), or because the individual performance measures are not sufficiently aligned with the firm's objectives (Baker, 2000).

Practitioners, however, emphasize a slightly different problem. As Stern Stewart, the advocates of EVA, point out, whatever weights may be attached to the measures of the scorecard are arbitrary. As a consequence, "tying rewards only to the scorecard metrics exacerbates the problem employees have in making multiple tradeoffs without having a definition of what better is" (Stern Stewart, 2003). In other words, if firms knew how the individual scores on the scorecard translate into value, EVA would be a redundant tool as long as the measures on the scorecard provide a sufficiently comprehensive picture of a manager's activities. But since they do not, rewards must be tied in part to an overall performance measure such as EVA to motivate managers to use their information at hand, no matter how sophisticated a scorecard is used. This is the problem emphasized in this paper.

I present a formal model in Section 2, and compare its assumptions with those of the most closely related papers. The equilibrium contract is derived in Section 3, and comparative statics results and empirical implications are derived in Section 4. In Section 5, I discuss in greater detail the relationship of my model to multitask models, and to the models of Feltham and Xie (1994) and Baker (2000). Section 6 concludes.

2 Model

A principal hires an agent to produce some output by exerting effort on two tasks.

Production: Output, denoted Y, is stochastic and can take the values 0 or 1. The probability that $Y = 1$ is realized depends on the effort $\mathbf{a} = (a_1, a_2)$ that the agent exerts on the two tasks, and the productivities $\boldsymbol{\theta} = (\theta_1, \theta_2)$ associated with the two tasks. Specifically, $\Pr(Y = 1)$ is given by $\max\{a_1\theta_1 + a_2\theta_2, 1\}$.

Technological uncertainty: The productivity θ_i of each task i is either high or low. It is given by $\theta_i = \bar{\theta}(1 - t\tau_i)$, where $\bar{\theta}$ is expected productivity and τ_i takes the values 1 or -1 with equal probability. The parameter $t \in [0, 1]$ determines the variance of θ_i and thus measures the degree of technological uncertainty.

The τ_i, and thus the θ_i, may be correlated; specifically, the probability that $\tau_1 = \tau_2$ is given by $(1 + \rho)/2$ for $\rho \in [-1, 1]$. When $\rho = 1$, the model in effect collapses to a one-task model. At the other extreme, when $\rho = -1$, the agent's first-best and individually optimal total effort is constant across all states of nature, and the agent only decides how to allocate total effort between the two tasks.

Information about θ: The principal is assumed to know the expected value $\bar{\theta}$, but not the realization θ_i, of each task. The agent receives a private signal $\mathbf{s} = (s_1, s_2)$ about $\boldsymbol{\tau}$, where $s_i \in \{-1, 1\}$. The probability that $s_i = \tau_i$ is given by $(1 + k)/2$, where $k \in [0, 1]$

captures the quality of the agent's knowledge. The conditional random variables $s_i|\tau_i$ are independent, so s_1 and s_2 are correlated only indirectly through the correlation between τ_1 and τ_2. The agent cannot communicate $\mathbf{s}$ to the principal; his signal is thus *specific knowledge* in the sense of Jensen and Meckling (1992).

Agent's utility: The agent is risk-neutral; his utility is given by $w - d(\mathbf{a})$, where $d(\mathbf{a})$ is the disutility of exerting effort. The disutility function has the form

$$d(\mathbf{a}) = \frac{d}{1+\phi}(a_1^2 + a_2^2 + 2\phi a_1 a_2),$$

for $\phi \in [-1, 1]$. If $\phi = 1$, the tasks are perfect substitutes for A in the sense that $d(\mathbf{a})$ reduces to $d(a_1+a_2)^2/2$. If $a_1 = a_2 = a$, then $d(\mathbf{a})$ reduces to $2da^2$. Thus, scaling disutility by $1/(1+\phi)$ ensures that changes in ϕ affect the interaction $\partial^2 d(\mathbf{a})/(\partial a_1 \partial a_2)$ but not the level of disutility for equal levels of effort on each task.

Performance measurement: The principal can measure the agent's effort $\mathbf{a} = (a_1, a_2)$ without noise, and $\mathbf{a}$ is contractible. Output, however, can be measured only imperfectly by a contractible variable y that takes the values 0 or 1. The extent to which y measures the true output Y is parameterized by a measure of *environmental uncertainty* or *risk*, denoted by $e \in [0, 1]$. Specifically, the probability that $y = Y$ is given by $(2-e)/2$. At the extremes, if $e = 0$, y measures Y perfectly, whereas if $e = 1$, y is entirely uninformative. For $\theta_1 a_1 + \theta_2 a_2 \in (0, 1)$, the expected value of y conditional on $\boldsymbol{\theta}$ and $\mathbf{a}$ is given by

$$\mathrm{E}(y|\boldsymbol{\theta}, \boldsymbol{a}) = \frac{e}{2} + (1-e)(\theta_1 a_1 + \theta_2 a_2) \tag{1}$$

Note from (1) that while environmental risk does not affect the true productivity of the agent's effort, it does reduce the responsiveness of *measured* performance to the agent's effort. This is a standard feature of any moral-hazard model with only two possible outcomes (see e.g. Laffont and Martimort (2001), but stands in contrast to e.g. the model of Holmström and Milgrom (1987, 1991).

Compensation: I restrict feasible contracts to those that are linear in both a_i and y; that is, the agent's total compensation is given by

$$w = \alpha + \beta(a_1 + a_2) + \gamma y.$$

The restriction to linear contracts is standard in the literature and motivated by their use in practice. On the other hand, the particular assumption made here that the principal pays piece rates for effort although effort can be measured perfectly, is a further simplification. Linear contracts are most realistic when performance measures are noisy measures of effort, whereas with verifiable effort and a known technology (known θ), a forcing contract would be superior. However, since the role of measurement errors for the design of optimal contracts is well known from standard agency theory[3], here I simplify the analysis by ignoring measurement errors while retaining the assumption of linear incentive contracts. As we will see, this simplification also has the virtue of highlighting one of the central points of this paper: *even if* effort can be perfectly measured, basing compensation on risky output measures is generally desirable when the agent has specific knowledge.

As mentioned, the agent is risk-neutral. However, he is protected by limited liability in the sense that his compensation must always be non-negative for any effort level and any realized output (given the disutility of effort, his realized utility may be negative, however). Since both the a_i and y can be zero, this means that the salary α must be nonnegative. Moreover, the agent will agree to a contract only if his resulting expected utility (if he chooses his effort optimally) exceeds his reservation utility, which may require paying the agent a positive salary. However, I assume that this constraint is not binding even when $\alpha = 0$. That is, I assume that the agent's reservation utility is smaller than his expected utility from working for the principal even when $\alpha = 0$, so that there is no need to pay him a salary.

Timing:

1. The principal offers a contract (β, γ); the agent accepts or rejects.

2. The technology shock $\boldsymbol{\tau}$ is realized.

3. The agent receives a signal $\mathbf{s}$ about $\boldsymbol{\tau}$.

[3] see e.g. Prendergast (1999) and Lafontaine and Slade (2000).

4. The agent chooses an effort vector **a**.

5. The output Y and the measured performance y are realized, and the agent is compensated accordingly.

To simplify the calculations, I restrict the analysis to parameters which lead to interior equilibrium solutions for all endogenous variables. There are three relevant constraints: in equilibrium, the optimal γ must be nonnegative, the agent's optimal effort levels for each task must be nonnegative, and the probability that $Y = 1$, namely $a_1\theta_1 + a_2\theta_2$, must be less than 1. To make the corresponding conditions on the exogenous parameters explicit, define

$$\eta = \frac{2(1-\phi)\rho(1-k^2\rho) + (1-\rho)^2(1+k^2\phi\rho)}{1-k^4\rho^2}. \tag{2}$$

Then the following conditions are necessary and sufficient for the existence of an interior solution for the principal's optimal contract:

$(A1)$ $\quad 2(1-e)k^2t^2\eta\bar{\theta}^2 > (1-\phi)de$

$(A2)$ $\quad (1-\phi^2)(1-\rho)de + (1-e)\eta kt\bar{\theta}^2[(1-\phi)(1-k^2\rho) - 2(1+\phi)(1-\rho)kt] > 0$

$(A3)$ $\quad 2(1-e)(1+k^2\rho)\eta kt[2d -$

$\quad (1+t)\bar{\theta}^2] - (1+t)(1+\rho)[2(1-e)\eta k^2t^2\bar{\theta}^2 - de(1-\phi)] > 0$

Assumption 1 ensures that the optimal γ in the contract is positive, and states that $\bar{\theta}$ must be sufficiently large relative to d. Assumption 2 ensures that for any realization of **s**, the agent chooses positive levels of effort for each task. It states that d must be sufficiently large relative to $\bar{\theta}$. Assumption (A3), finally, ensures that $a_1\theta_1 + a_2\theta_2$ never exceeds 1. None of the conditions (A1)-(A3) are essential for the main results of this paper. If they do not hold, however, the expressions used here are no longer valid, and tedious case distinctions become necessary.

The model is closely related to a number of other agency models. The explicit multitask framework adopted here is similar to that of Holmström and Milgrom (1991). The main difference is that here, the agent is assumed to have private knowledge about the productivities of tasks. Private information, in turn, is an important ingredient of the models of Baker (1992), Baker, Gibbons and Murphy (1994), and Lafontaine and Bhattacharyya (1995). While neither of these papers considers multiple tasks, other differences are more important: in Baker (1992), contracts can be based only on one performance measure. Lafontaine and Bhattacharyya (1995) do not allow for contracts based on an input measure; moreover, their model requires the use of simulations in place of mathematical analysis to obtain comparative statics results. Baker, Gibbons and Murphy (1994) consider compensation based on both an input and an output measure. They assume that the agent is risk-neutral but do not impose limited liability as I do here. However, in their model, the output measure is subjective and noncontractible, and therefore purely output-based compensation may not feasible because it would give the principal a too strong incentive to renege on promised bonus payments.

Baker (2000) allows for multiple performance measures, but treats the (mis)alignment between the principal's objectives and the agent's measured performance as parametric, whereas here it is derived from the agent's private information. The relationship between the present model and Baker (2000) is discussed in more detail in Section 5.2.

Finally, the most closely related model is that of Baker and Jorgensen (2003), which distinguishes between two types of uncertainty and allows for input- and output-related pay. Their model, which is an extension of the continuous (static) model of Holmström and Milgrom (1987), allows for a comparative-statics analysis but does not lead to a closed-form solution when the agent has private information before choosing his effort. More importantly, their main result, namely a positive relationship between volatility (what I call technological uncertainty) and incentives, is an existence result, whereas here, it holds generally. Baker and Jorgensen also show that (as here) optimal compensation is partially based on output even when effort is verifiable, but do not study the determinants of input- vs. output-related pay.

3 Optimal Contract

The principal maximizes her expected net profit $E(Y) - E(w)$ with respect to the contract parameters, subject to the agent's optimal choice of $\mathbf{a}$ which depends on his private information $\mathbf{s}$. I therefore first derive the agent's optimal choice, and then derive the optimal contract.

3.1 Agent's Beliefs and Optimal Choice of Effort

Upon observing $\mathbf{s}$, the signal about the productivities $\boldsymbol{\theta}$, the agent forms expectations about $\boldsymbol{\theta}$, denoted by $\hat{\boldsymbol{\theta}}(\mathbf{s})$, and subsequently chooses his optimal effort based on his expectations.

Lemma 1 *The expected value of θ_i conditional on $\boldsymbol{s}$ is given by*

$$\hat{\theta}_i(\mathbf{s}) = \left[1 + t\frac{k(s_i + \rho s_j)}{1 + k^2 \rho s_1 s_2}\right]\bar{\theta} \tag{3}$$

for $i = 1, 2$ and $j \neq i$.

Proof: see the Appendix.

In particular, for $k = 0$ the agent's expectation is $\hat{\theta}_i = \bar{\theta}$, whereas for $k = 1$, $\hat{\theta}_i$ equals θ_i irrespective of ρ and s_j (this follows because $s_i \in \{-1, 1\}$). The agent's expected utility is

$$\begin{aligned} & E_{\theta,y}(w(\mathbf{a})|\mathbf{s}) - d(\mathbf{a}) \\ = \; & \beta(a_1 + a_2) + \gamma E(y(\boldsymbol{\theta}, \mathbf{a})|\mathbf{s}) - d(\mathbf{a}). \end{aligned} \tag{4}$$

Assuming that

$$a_1\theta_1 + a_2\theta_2 \leq 1 \quad \text{for all possible } \boldsymbol{\tau}, \tag{5}$$

(which I will later show is satisfied when assumption (A3) holds), we have $\mathrm{E}(Y|\boldsymbol{\theta}) = \widehat{\theta}_1 a_1 + \widehat{\theta}_2 a_2$, and together with (1), (4) expands to

$$\beta(a_1 + a_2) + \gamma\{e/2 + (1 - e)[\widehat{\theta}_1 a_1 + \widehat{\theta}_2 a_2]\} - \frac{d}{2}\left(a_1^2 + a_2^2 + 2\phi a_1 a_2\right). \tag{6}$$

This expression is strictly concave in $\mathbf{a}$, and the first-order conditions for maximization with respect to $\mathbf{a}$ are

$$\beta + \gamma(1 - e)\widehat{\theta}_i - da_i - \phi a_j = 0, \tag{7}$$

for $i = 1, 2; j \neq i$. The solution of the system (7) for a_1 and a_2 is

$$a_i^*(\hat{\boldsymbol{\theta}}) = \frac{(1-\phi)\beta + (1-e)(\widehat{\theta}_i - \phi\widehat{\theta}_j)\gamma}{(1-\phi^2)d} \tag{8}$$

for $i = 1, 2; j \neq i$, which is an affine function of $\widehat{\boldsymbol{\theta}}$. I will for now assume both that (5) holds (so that (8) is valid), and that (8) leads to an interior (i.e. $a_i \geq 0$) solution for all realizations of $\boldsymbol{\theta}$ and $\mathbf{s}$, and will later turn to the necessary and sufficient parameter conditions for this to be true.

3.2 Optimal Input- vs. Output-Related Incentives

The principal chooses the contract parameters β and γ to maximize the expected value of her net profit $\pi = Y - w$, subject to the agent's choice of $\mathbf{a}$ given by (8).

Proposition 1 *Under assumptions (A1)-(A3), the optimal contract parameters are given by*

$$\beta^* = \frac{(1-\phi)de}{4(1-e)k^2t^2\eta}\bar{\theta} \quad \textit{and} \tag{9}$$

$$\gamma^* = \frac{1}{2(1-e)} - \frac{(1-\phi)de}{4(1-e)^2k^2t^2\eta\bar{\theta}^2} \tag{10}$$

Proof: see the Appendix.

At one extreme, when pay is only input-based ($\gamma = 0$), then it is clear from inspection of (8) that the agent's effort depends only on β and not on his private information $\hat{\boldsymbol{\theta}}(\mathbf{s})$.

At the other extreme, purely output-based pay ($\beta = 0$) is costly for the principal whenever measuring output is subject to errors ($e > 0$). This result is standard when the agent is risk-averse, but holds for similar reasons when the agent is risk-neutral but protected by limited liability. More precisely, when the agent is risk-averse and when measured performance is additive in effort and noise (as in the Holmström-Milgrom (1987) framework), then the variance of the noise term affects the agent's participation constraint but not his incentive constraint, i.e. his optimal effort. More generally, however, and in particular in models with only two realizations of performance, the variance of performance affects the incentive constraint as well as the participation constraint (see Laffont and Martimort, 2001).

With a risk-neutral agent and limited liability, finally, risk affects only the incentive constraint: the more measured performance is affected by noise, the less it is affected by effort, cf. (1). Inducing the agent to exert effort therefore requires the principal to pay a higher reward for a good outcome. Since the reward for a bad outcome is already bounded from below at zero, greater risk is costly for the principal (cf. Laffont and Martimort, 2001).

Thus, the key tradeoff in this model is that input-based pay is riskless but fails to make use of the agent's information, whereas output-based pay encourages the optimal use of information but is costly for the principal because output is measured with some noise. The optimal incentive contract therefore is partly input- and partly output-based.

4 Comparative Statics

In this section, I study the effects of changes in the model parameters on the optimal contract parameters β and γ as given in Proposition 1. Typically, any parameter change will lead to a change in β and γ in opposite directions. I first discuss the comparative-statics results individually, and then (in Section 4.4) present a summarizing result.

Recall that η defined in (2) does not depend on $\bar{\theta}$ or d. By inspection of (9) and (10) in Proposition 1, we immediately have:

Proposition 2 *γ^* is increasing in $\bar{\theta}$ and decreasing in d. The derivatives of β^* in these parameters have the opposite sign.*

In other words, an increase of the expected productivity of effort lead to more output-based pay, an increase in the disutility of effort to more input-based pay.

4.1 Specific Knowledge and Implicit Delegation

It is a well-known idea that a principal may want to delegate decisions to an agent if the agent has superior information pertaining to those decisions compared with the principal. This idea has been emphasized, and is also a key element of recent models of delegation, such as Aghion and Tirole (1997), Dessein (2002) and Prendergast (2002).

Here, in contrast, it is assumed that the agent is given nominal freedom in how to choose his effort; that is, delegation is already given. This assumption applies perhaps best to executives and other employees who generally work without much direct supervision. Nevertheless, the choice of performance measures indirectly influences how much freedom the agent is given in choosing how to allocate effort. The greater the weight on output-based pay, the more the agent is encouraged to use his knowledge about where his effort will be most productive. Thus, given formal authority (in the sense of Aghion and Tirole, 1997) over how to choose effort, the relative weight of output- vs. input-related pay defines a degree of *implicit delegation*. An intuitive result is:

Proposition 3 *γ^* is increasing in k; β^* is decreasing in k.*

Proof: see the Appendix.

Thus, the greater the agent's information advantage over the principal, the more the principal will want to base compensation on output. Notice that k measures the agent's information advantage *relative* to the principal. It follows that the agent's optimal compensation plan depends on how much the principal knows about the agent's job. For example, the optimal compensation plan for an engineer heading a manufacturing plant depends on whether or not his boss, the division manager, is an engineer herself. Other things equal, the optimal compensation plan for the plant manager is more likely to include non-financial (input) measures when the division manager also is an engineer, than when her background is different.

Some elements of the agent's knowledge will always be specific to the agent as an individual and difficult to measure; for instance, the agent's "personal working style". Knowledge about the agent's job, on the other hand, may be more, or less, specific to the agent, allowing for empirical investigation. Building on the previous example of the plant manager, one could e.g. investigate how the mix between input- and output-based compensation depends on whether agent and principal (or rather, employees and those those designing their compensation plans) have the same educational background.

4.2 Uncertainty and Incentives

Recently, much attention has been devoted to the apparent discrepancy between theoretical predictions and empirical findings on the relationship between risk and incentives. The standard principal-agent model due to Holmström (1979) and Holmström and Milgrom (1987) predicts that the noisier a performance measure, the lower the optimal piece rate on that measure will be when the agent is risk-averse; i.e. it predicts a negative relationship between risk and incentives. Empirical studies, however, tend to find the opposite, namely a positive correlation between various measures of risk and the strength of incentives provided to agents, cf. Prendergast (2002) and Lafontaine and Slade (2000).

One source of the discrepancy may be that the value of managerial effort varies in unaccounted ways across observations, and the empirical literature continues to search for appropriate proxies for the value of effort. Another reason for the discrepancy may be an argument due to Demsetz and Lehn (1985) for why the relationship between risk and incentives might be positive: "In less predictable environments, however, managerial behavior ... figures more prominently in a firms fortunes... Hence, noisier environments should give rise to more concentrated ownership structures."

Both ideas are investigated in recent theoretical work. Prendergast (2002) presents a variety of scenarios that lead to a positive relationship. In one of these, a principal may want to delegate the choice among different available projects to an agent if the agent has better information about the profitability of the projects. In line with Demsetz and Lehn's argument, Prendergast shows that an increase in the riskiness of the payoffs associated with the projects increases the value of the agent's information and thus makes delegation more likely. In Raith (2003), I argue that both the value of managerial effort and empirical measures of firm risk endogenously depend on the degree of product market competition. In the model, differences in the degree of competition induce a positive correlation between incentives and the variance of firms' profits without any direct causal link between the two.

The model studied here demonstrates that whether one should expect to see a positive or a negative relationship between incentives and risk depends on the source of uncertainty.

What I termed environmental risk, captured by e, plays a very similar role as risk in the standard principal-agent model does. Environmental risk has no effect on the productivity of effort, but reduces the effect of effort on measured performance, and thus leads to less effort (as is also standard in two-outcome models with risk aversion). Technological uncertainty (t), on the other hand, captures the variance of the contribution of each task to output. Changes in e and t have very different implications:

Proposition 4 *(a) If $\partial E(y)/\partial e \geq 0$, then γ^* is decreasing in e; β^* is increasing in e for all e.*

(b) γ^* is increasing in t, β^* is decreasing in t.

Proof: see the Appendix.

To understand part (a), notice that in any agency model with only two possible outcomes, changing the variance of the outcome by changing the probabilities of the realizations will also change the expected value. Consequently, in this model, a change in e affects not only the variance of y, but its expected value as well. It is obvious that the principal would want to decrease γ^* if $\mathrm{E}(y)$ increases and vice versa. The more interesting question is how γ^* is affected when $\partial \mathrm{E}(y)/\partial e = 0$. As part (a) of Proposition 4 states, an increase in e unambiguously leads to a decrease in γ. This effect is reinforced for higher values of e, where $\mathrm{E}(y)$ is increasing in e, whereas for smaller values it may be outweighed by the upward adjustment of e because of a decrease in $\mathrm{E}(y)$.

An increase in t, in the other hand, unambiguously induces the principal to increase γ and decrease β; i.e. leads to more delegation. With greater technological uncertainty, the principal faces greater uncertainty about which tasks the agent should focus his effort on. As in Prendergast (2002), this implies a larger difference between the expected output if the agent uses his private information, and the expected output if he does not. The principal therefore increases γ^* and decreases β^* to induce the agent to make better use of his information. A similar result is derived, as a possibility result, in Baker and Jorgensen (2003).

Proposition 4 still does not quite pin down the precise reason for why environmental risk and technological uncertainty have opposite effects on incentives. Several features of the ways in which uncertainty enters the model are worth pointing out: 1. $\boldsymbol{\theta}$ is information the agent learns *before* choosing effort, whereas the realization of y conditional on $\boldsymbol{\theta}$ and $\mathbf{a}$ is not. 2. $\boldsymbol{\theta}$ affects the agent's first-best action through the interaction of $\boldsymbol{\theta}$ and $\mathbf{a}$ in the determination of Y. 3. $\boldsymbol{\theta}$ affects the agent's individually rational effort through the effect of $\boldsymbol{\theta}$ and $\mathbf{a}$ on y.

As it turns out, among these features, 1. and 2. are the ones that lead to a positive relationship between (output) incentives and technological uncertainty. First, for reasons already explained, exposing the agent to risk is always costly for the principal and hence leads to lower optimal incentives, *unless* the realized state of nature influences the agent's actions. Thus, knowledge of signals about the state of the world, and their relevance for the agent's chosen actions, is necessary to obtain a positive relationship between uncertainty and incentives, cf. also Baker and Jorgensen (2003) on this point.

Next, however, it also matters what kind of information the agent obtains. In this model as well as in Baker and Jorgensen's, the agent's information pertains to the true productivity of effort. The agent's information is valuable in the sense that acting upon it is optimal from the principal's point of view. While the agent cares about measured performance y, knowledge of $\boldsymbol{\theta}$ affects y only through its influence on Y. Thus, here the agent's information is "good" information.

On the other hand, it is conceivable that the agent's information pertains only to measured performance but not to true productivity. In this case, the agent's information may affect his individually optimal action but not his first-best action. For instance, if the agent receives information that measured performance will be low, the agent has little incentive to exert effort even if the effect of effort on the principal's payoff may be high. In this sense, the agent's information is "bad" information. This is the situation modeled by Baker, Gibbons and Murphy (1994), who show that an increase in the variance of the objective performance measure (but not true output) will lead to a lower optimal weight on the performance measure. An intermediate case is considered by Baker (1992), where the agent's private information about the state of nature pertains to both true output and measured performance, with a constant correlation between the two. In Baker's model, too, an increase in the variance of the state of nature leads to a lower optimal weight on measurable performance. More generally, whether an agent's information is "good" or "bad" will depend on the precise correlation structure between the agent's information, the productivity of effort, and measured performance.[4]

4.3 Task Variety

The smaller ρ, i.e. the correlation of the θ_i, the more likely the θ_i, and hence the efficient levels of effort for each task, will be different. In this sense, one can think of ρ is a (negative) measure of (objective) *task variety*: a job with greater task variety (smaller ρ) is a job in which the tasks the agent should work on are more likely to differ from one point in time to the next. Put differently, greater task variety means that the optimal allocation of total effort across the two tasks is more difficult to specify in advance. This measure of task variety is "objective" in the sense that it is a characteristic of the production function, i.e. agent's job. In contrast, the parameter ϕ in the agent's utility function measures the substitutability or complementarity of tasks from the agent's subjective perspective.

The term "task variety" is borrowed from Hackman and Oldham (1974).[5] In this influential contribution to the Human Relations school of personnel research, the authors argue that workers' productivity depends on their "intrinsic motivation". Among the ways in firms can increase intrinsic motivation is to give workers jobs with a greater range of tasks (termed "job enrichment") and to give workers greater autonomy. Such ideas appear to contrast with the mainstream approach of personnel economics, which focuses

[4] Baker and Jorgensen (2003) argue that a positive relationship can exist between volatility and incentives, where volatility is defined as "uncertainty that does affect the agent's optimal agent, and which the agent is able to react to." While this is correct, it needs to be emphasized that "optimal action" should mean first-best optimal *and* individually optimal. The statement is incorrect when "optimal action" means only individually optimal, as the agent might respond to information pertaining to measured performance but not true output.

[5] It would be inappropriate to speak of substitutability or complementarity of tasks in the context of the parameter ρ because, for reasons of tractability, the cross-partial derivatives of a_1 and a_2 in the expected value of Y are zero, meaning that the tasks are independent in this sense.

on explicit incentives as the drivers of workers' behavior and tends to ignore other sources of motivation (see Gibbs and Levenson, 2000).

But since the present model includes a parameter that can be interpreted as a measure of "task variety", we can investigate elements of the Hackman-Oldham theory more formally. The first observation is that optimal effort as given by (8) does not depend on ρ. In particular, an increase in task variety (decrease in ρ) has no effect on effort unless the principal also changes the agent's explicit incentives β and γ. But the principal will in fact want adjust the incentives:

Proposition 5 *γ^* is decreasing in ρ, and β^* is increasing in ρ, if and only if*

$$k^2\phi(1-\rho^2)(1-k^4\rho^2)+2(1-k^2)[\phi(1-k^2\rho^2)-\rho(1-k^2)]\geq 0. \tag{11}$$

Condition (11) holds for all $\rho \leq \tilde{\rho}$ for some $\tilde{\rho} \in (0,1)$ and is violated for all $\rho > \tilde{\rho}$. In particular, it also holds if $a_1^*(\boldsymbol{\theta})$ is larger when the agent obtains a bad signal on task 2 than when he obtains a good signal.

Proof: see the Appendix.

Proposition 5 states that if the productivities of the tasks are not already highly correlated, an decrease in ρ will induce the principal to increase γ and decrease β, which may be interpreted as giving the agent greater "autonomy". The intuition is similar as for increases in technological uncertainty: the smaller ρ, the more likely it is that the agent should focus his effort on one task rather than both. It is then optimal to make better use of the agent's private information about which task is the most productive, by basing compensation more on output.[6]

Recall that Propositions 3 and 4 state that compensation will be more output-based the greater technological uncertainty (t) or the agent's knowledge advantage (k). Proposition 5, in addition states that compensation will be more output-based the greater the

[6] This effect is only partially offset by the fact that with lower ρ, the agent's effort levels a_1 and a_2 are less correlated anyway even for given β and γ. This can be seen by evaluating $a_1 - a_2$ using (8) and (3). The derivative of this expression with respect to ρ turns out to have the sign of $s_2 - s_1$, which means that a decrease in ρ will lead to an increase in $a_1 - a_2$ if $s_1 = 1$ and $s_2 = -1$.

variety of tasks an agent needs to perform. This leads to the empirical prediction that the weight on output-based incentives should vary positively with measures of job complexity (for an example of the use of such measures, see e.g. John and Weitz, 1989). Because of the different consequences of environmental risk and technological uncertainty, it may be difficult to find a strong correlation between incentives and measures of risk. In contrast, although empirical measures of risk and complexity are likely to be correlated, one would expect to find a clearer positive relationship between job complexity and incentives.

Changes in ϕ, i.e. the degree of subjective substitutability between tasks, have, in a sense, the opposite effect on optimal incentives:

Proposition 6 *γ^* is increasing in ϕ, and β^* is decreasing in ϕ.*

Proof: see the Appendix.

A lower ϕ means that the agent perceives the two tasks as less similar and prefers to work on both tasks rather than one. In response, his optimal allocation of effort depends to a lesser extent on his private information about the productivities of the tasks. This can be seen by evaluating $a_1 - a_2$ using (8) and (3). The derivative of this expression with respect to ϕ has the sign of $s_1 - s_2$, which is positive if e.g. $s_1 = 1$ and $s_2 = -1$. The principal's optimal response to the agent's lower responsiveness to his information is to decrease γ and increase β. Unfortunately, characteristics of individuals' utility functions are difficult to measure, and therefore Proposition 6 seems difficult to test empirically.

4.4 Alignment

A key ingredient in recent models of moral hazard and delegation is the assumption that the agent's incentives are imperfectly aligned with the principal's. In the moral-hazard models of Feltham and Xie (1994) and Baker (2000), for instance, it is assumed that available performance measures are imperfectly correlated with the principal's payoff. In these models, the correlation structure is an exogenous part of the model and thus defines how well (in terms of the model parameters, not the optimal contract) the agent's and the principal's objectives are aligned (see also Section 5.2).

Similarly, in models of delegation, a principal's decision whether to delegate authority to an agent depends on how closely the agent's and principal's preferences are aligned. Aghion and Tirole (1997) call this degree of alignment "congruence" and measure it by the probability that the principal's and the agent's favorite projects are the same.

In models of delegation, greater alignment benefits the principal both with and without delegation. With delegation, the agent is more likely to make decisions in the interest of the principal. But without delegation, the principal benefits too because the agent is less likely to distort private information that he reports to the principal. It is therefore a priori clear whether an increase in alignment will lead to more or less delegation. The results in the delegation literature reflect this ambiguity. Aghion an Tirole (1997) do not study the effect of a change in congruence on the principal's decision whether to grant formal authority to the agent. In Dessein's (2002) model, greater congruence makes delegation more likely. In Prendergast (2002), in contrast, a higher correlation between the principal's and the agent's preferred projects makes delegation less likely.

In the present model, there is no parameter of alignment; the agent always dislikes effort but beyond that does not receive any private benefits from pursuing particular actions. Building on the above observations about delegation models, however, one way to think about alignment is to consider any parameter change that, for a given contract, leads to a higher payoff for the principal, as an increase in the alignment of the principal's and the agent's objectives.

A comparative-statics analysis for expected output $E(Y)$ and expected surplus $\mathrm{E}(Y) - E(d(\mathbf{a}))$, combined with Propositions 2 through 6, leads to the following result:

Proposition 7 *Any change in of one the parameters $\bar{\theta}$, d, t, k, ρ and ϕ leads to an increase in γ^* and a decrease in β^* if and only if such a change leads to a higher expected output $E(Y)$, for β and γ held fixed, which occurs if and only if the same change increases the expected total surplus $E(Y) - E(d(\mathbf{a}))$, for β and γ held fixed. The same holds for a change in e whenever $\partial E(y)/\partial e \geq 0$.*

Proof: see the Appendix.

Proposition 7 thus states that any parameter change that leads to a better alignment (in the sense suggested above) of the principal's and the agent's objectives also leads to more output-based pay, i.e. a greater degree of implicit delegation.

A word of caution about the generality of Proposition 7 is due, however: first, the proposition summarizes rather than generalizes the previous results in the sense that it is proven for each parameter separately rather than in a general way. Second, while the result holds for all parameters of the model, one can envision a case where it would fail to hold: suppose the agent's effort could be measured only with some measurement error. Then a decrease in the error would lead to an increase in surplus, but also a shift toward input-based pay. Thus Proposition 7 would not hold for all parameters in this extended model.

5 Discussion

This section discusses the relationship between the perspective on incentives offered in this paper, and two related but distinct contributions to the theory of incentives that have emerged over the past decade, namely multi-task models and insights on the tradeoff between noise and distortion.

5.1 Relationship With Multi-Task Models

The model studied here is formally closely related to previous multi-task models. Its results, however, differ in spirit from the multi-task literature and more closely resemble those of delegation models.

The central insight of the multi-task literature is that if what the principal can measure does not perfectly coincide with what she seeks to maximize, incentive pay may fail because the agent is induced to "game the system". Holmström and Milgrom (1991), for example, show that if performance on one task is easy to measure and performance on the other task difficult to measure, then it may be optimal to offer low-powered incentives on both tasks to prevent the agent from focusing his effort too heavily on the task that is easier to measure. Similarly, Baker (1992) shows that if the agent has private information about how strongly his effort is related to measured performance, the principal may want to offer lower-powered incentives to counteract the agent's tendency to distort the performance measure.

In spite of formal similarities, the flavor of the results of this paper is quite different. First, while multi-task papers emphasize the dysfunctional consequences of incentive pay when performance measures and objectives are not aligned, the emphasis here is on how to make the best use of the agent's private information when inducing him to exert effort. More formally, an agent's private information about the relationship between effort and measured performance (as in Baker 1992) gives rise to distortion and hence reduces the principal's payoff. In contrast, an agent's private information about the relationship between effort and output (as assumed here) is valuable for the principal, and the optimal contract is designed to encourage rather than suppress the use of private information.

Second, when it is feasible to base contracts on measured output and not only on inputs, then greater technological uncertainty or better private information not simply leads to lower incentives as e.g. in Baker (1992), but leads to a shift from input- to output-based incentives. Which prediction is more realistic for any particular situation depends on whether suitable measures of output are available. Standard multi-task models explain why explicit pay for performance is rare in jobs in which some relevant tasks are difficult to measure and for which "output" is not well-defined. On the other hand, in some of the most complex jobs, for instance those of executives, incentives are high-powered but based on more aggregate measures even though more accurate indicators of "effort" are often available. For those jobs, the theory presented here seems better suited than multi-task models.

5.2 Noise, Distortion and Interference

Feltham and Xie (1994) and Baker (2000) have argued that in choosing the optimal set of performance measures for an agent, a principal faces a tradeoff between exposing the agent to risk and encouraging him to distort the performance measures. That is, basing compensation on indicators that measure the agent's input most accurately also tends to encourage the agent to allocate his effort in unproductive ways. Feltham and Xie (1994) and Baker (2000) thus argue that the optimal choice of performance measures depends on both the noisiness of performance measures as well as their alignment with the principal's objectives. In contrast, the tradeoff emphasized here is not that between noise and distortion but between noise and interference.

The problem of alignment essentially already appears in the work of Holmström and Milgrom (1991) and Baker (1992). Performance indicators are less than perfectly aligned with the principal's objectives if either they fail to measure some aspects relevant to the agent's job or if the agent has private information on how to manipulate the indicator. To use one of Baker's (2000) examples, banks typically compensate loan officers based on loans originated rather than loans successfully completed. The performance measure is imperfect because it fails to measure the officer's effort spent on assessing the quality of a loan applicant.

Suppose in contrast that, thanks to balanced scorecards and other modern management tools, all activities relevant to the agent's job could be measured and the contributions of those activities to the principal's objectives were known. Then there would not be any problem of alignment, since the performance measure could simply be chosen to equal the deterministic component of the agent's overall contribution to profit. In this scenario, the tradeoff emphasized by Baker (2000) would not exist. Nevertheless, as pointed out in this paper, the firm might want to use aggregate financial measures anyway if the agent knows better than the firm how the inputs contribute to profit.

Thus, in the framework considered here the principal faces a tradeoff between noise and *interference* rather than between noise and distortion. Basing compensation on performance measures closely related to the agent's activities may lead to distortions if some activities critical to the job are difficult to measure, as pointed out by the multi-task literature. But even if all relevant activities can be measured, the agent may know better than the principal what actions to choose to maximize the principal's profit. Basing compensation on inputs would then amount to telling the agent how to spend his time, and thus could amount to too much interference.

At a different level, the present model can also be regarded as an endogenization of the tradeoff between distortion and noise. Baker (2000) discusses reasons for why a tradeoff between precise but distortionary measures and well-aligned but noisy measures may exist, but does not model the tradeoff. In particular, in the models of Feltham and Xie (1994) and Baker (2000), the distortion associated with any performance measure is parametric, i.e. exogenous. Here, the tradeoff between distortion and noise arises endogenously because of technological uncertainty and specific knowledge: Output-based compensation is risky, but aligns the agent's objectives with the principal's. Input-based compensation is safe, but distortionary because the tasks the principal induces the agent to pursue may not be much related to the firm's output if the principal does not know the productivity of each task.

6 Concluding Remarks

(To be completed.)

Appendix: Proofs

Proof of Lemma 1: For $\tau_i \in \{-1, 1\}$, we have $\tau_1\tau_2 = 1$ if $\tau_1 = \tau_2$ and $\tau_1\tau_2 = -1$ otherwise. Since $\Pr(\tau_1 = \tau_2) = (1 + \rho)/2$, it follows that for $\boldsymbol{\tau} \in \{-1, 1\}^2$, we have $\Pr(\boldsymbol{\tau}) = (1 + \rho\tau_1\tau_2)/4$. Moreover, since $\Pr(s_i = \tau_i) = (1 + k)/2$ and since the conditional distributions $s_i|\tau_i$ are independent, it follows that for any $\boldsymbol{\tau}, \boldsymbol{s} \in \{-1, 1\}^4$, the probability of $\boldsymbol{s}$ conditional on $\boldsymbol{\tau}$ is given by $\Pr(\boldsymbol{s}|\boldsymbol{\tau}) = (1 + ks_1\tau_1)(1 + ks_2\tau_2)/4$.

The unconditional probability of $\boldsymbol{s} \in \{-1, 1\}^2$ is then given by

$$\Pr(\boldsymbol{s}) = \sum_{\boldsymbol{\tau}\in\{-1,1\}^2} \Pr(\boldsymbol{s}|\boldsymbol{\tau}) \Pr(\boldsymbol{\tau}),$$

which simplifies to $(1 + k^2\rho s_1 s_2)/4$.

The expected value of $\boldsymbol{\tau}$ conditional on $\boldsymbol{s}$ is therefore given by

$$\mathrm{E}(\boldsymbol{\tau}|\boldsymbol{s}) = \sum \boldsymbol{\tau} \in \{-1, 1\}^2 \boldsymbol{\tau} \Pr(\boldsymbol{s}|\boldsymbol{\tau}),$$

which reduces to

$$\mathrm{E}(\tau_i|\boldsymbol{s}) = \frac{k(s_i + \rho s_j)}{1 + k^2\rho s_1 s_2}$$

for $i = 1, 2$ and $j \neq i$, leading to (3). ■

Proof of Proposition 1: Using (8), expected output conditional on $(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}})$ is

$$\begin{aligned} Y(\boldsymbol{\theta}, \widehat{\boldsymbol{\theta}}) &= a_1\theta_1 + a_2\theta_2 & (12)\\ &= \frac{\beta}{d(1+\phi)}(\theta_1 + \theta_2) + \frac{(1-e)\gamma}{d(1-\phi^2)}[\theta_1(\hat{\theta}_1 - \phi\hat{\theta}_2) + \theta_2(\hat{\theta}_2 - \phi\hat{\theta}_1)] & (13) \end{aligned}$$

To obtain the expected value of (13) over $(\boldsymbol{\theta}, \mathbf{s})$, observe that the expected value of $\theta_1 + \theta_2$ in the first term is just $2\hat{\theta}$. The expected value of the term in [] brackets in the second term can be computed as $2[1 - \phi + k^2t^2\eta]$ Thus, we have

$$2\frac{\bar{\theta}\beta}{(d(1+\phi)} + 2\frac{(1-e)\gamma\bar{\theta}^2}{d(1-\phi^2)}(1 - \phi + k^2t^2\eta) \tag{14}$$

The principal's expected profit conditional on $(\boldsymbol{\theta}, \mathbf{s})$ is

$$\begin{aligned} \pi &= Y(\boldsymbol{\theta}, \mathbf{s}) - w & (15)\\ &= Y(\boldsymbol{\theta}, \mathbf{s}) - \beta(a_1 + a_2) - \gamma\left[\frac{e}{2} + (1-e)Y(\boldsymbol{\theta}, \mathbf{s})\right] \\ &= [1 - \gamma(1-e)]Y(\boldsymbol{\theta}, \mathbf{s}) - \beta(a_1 + a_2) - \frac{\gamma e}{2} & (16) \end{aligned}$$

The expected value of $a_1 + a_2$ in (16) over $(\boldsymbol{\theta}, \mathbf{s})$ is

$$2\frac{\beta + \gamma(1-e)\bar{\theta}}{d(1+\phi)}$$

Substitute this expression and (14) into (16) to obtain the principal's expected profit:

$$[1-\gamma(1-e)\left[\frac{\bar{\theta}\beta}{(d(1+\phi)} + \frac{(1-e)\gamma\bar{\theta}^2}{d(1-\phi^2)}(1-\phi+k^2t^2\eta)\right] - \beta\frac{\beta+\gamma(1-e)\bar{\theta}}{d(1+\phi)} - \frac{\gamma e}{2}. \quad (17)$$

Differentiating (17) with respect to β and γ leads to the first-order conditions

$$\frac{2(1-2(1-e)\gamma)\bar{\theta} - 4\beta}{d(1+\phi)} = 0 \quad \text{and} \quad (18)$$

$$-\frac{e}{2} + \frac{2(1-e)\bar{\theta}(1-2(1-e)\gamma)(1-\phi+k^2t^2\eta)\bar{\theta} - 2(1-\phi)\beta}{d(1-\phi^2)} = 0 \quad (19)$$

The solution of (18) and (19) for β and γ is stated in the proposition. Since

$$\frac{\partial^2 \mathrm{E}(\pi)}{\partial\beta^2} = -\frac{4}{d(1+\phi)}, \quad \frac{\partial^2 \mathrm{E}(\pi)}{\partial\gamma^2} = -\frac{4(1-e)\bar{\theta}}{d(1+\phi)}, \quad \text{and}$$

$$\frac{\partial^2 \mathrm{E}(\pi)}{\partial\beta^2}\frac{\partial^2 \mathrm{E}(\pi)}{\partial\gamma^2} - \left(\frac{\partial^2 \mathrm{E}(\pi)}{\partial\beta\partial\gamma}\right)^2 = \frac{16(1-e)\eta k^2t^2\bar{\theta}^2}{d^2(1-\phi)(1+\phi)^2},$$

$\mathrm{E}(\pi)$ is strictly concave in β and γ and so the solution of (18) and (19) is indeed a maximum.

What remains to be shown is that under Assumptions (A1)-(A3), the interior solution obtained here is indeed valid. Three conditions must be satisfied: (1) the expression in (10) must be positive (otherwise the optimum is a corner solution with $\gamma = 0$); (2) The agent's effort must be positive for each realization of $\mathbf{s}$ for (8) to be valid; and (3) $\theta_1 a_1 + \theta_2 a_2$ must never exceed one for the calculation of $\mathrm{E}(Y)$ and hence $E(\pi)$ to be valid. I examine each constraint in turn.

(To be completed)

■

Proof of Proposition 3: From inspection of (9) and (10), the result holds if and only if $k^2\eta$ is increasing in k. Substituting from (2), one can compute

$$\frac{\partial(k^2\eta)}{\partial k} = \frac{2k\left[(1-\rho)^2(1+k^2\rho)^2 + 2\rho(1-\phi)(1-k^2)(1-k^2\rho^2)\right]}{(1-k^4\rho^2)^2} > 0.$$

■

Proof of Proposition 4: Part (b) follows from inspection of (9) and (10). For part (a), it follows from inspection of (9) that β is increasing in e. The derivative of γ with respect to e is

$$\frac{1}{2(1-e)^2} - \frac{d(1+e)(1-\phi^2)}{8(1-e)^3 k^2 t^2 \eta \bar{\theta}^2},$$

which is positive if and only if

$$d(1+e)(1-\phi^2) - 4(1-e)k^2 t^2 \eta \bar{\theta}^2 > 0. \tag{20}$$

Next, use (1) and (14) to obtain

$$\mathrm{E}(y) = \frac{e}{2} + \frac{2(1-e)\bar{\theta}}{d(1+\phi)} \left[\beta + (1-e)\gamma\bar{\theta}\frac{1-\phi+k^2t^2\eta}{1-\phi}.\right] \tag{21}$$

Differentiate (21) with respect to e and substitute β and γ from Proposition 1. Set the resulting expression to zero and solve for e to obtain

$$\tilde{e} = \frac{2\eta k^2 t^2 \left[4(1-\phi+k^2t^2\eta)\bar{\theta}^2 - d(1-\phi^2)\right]}{d(1-\phi)^2(1+\phi) + 8\eta k^2 t^2 (1-\phi+k^2t^2\eta)\bar{\theta}^2} \tag{22}$$

The expression in (22)is positive if and only if

$$\bar{\theta}^2 \geq \frac{d(1-\phi^2)}{4(1-\phi+k^2t^2\eta)} \tag{23}$$

Substitute $\tilde{e}$ for e into (20) to obtain:

$$\frac{d(1-\phi^2)\left[4k^2t^2\eta(3(1-\phi)+2k^2t^2\eta)\bar{\theta}^2 + d(1-phi^2)(1-\phi-2k^2t^2\eta)\right]}{d(1-\phi)^2(1+\phi) + 8k^2t^2\eta(1-\phi+k^2t^2\eta)\bar{\theta}^2} \tag{24}$$

The expression (24) has the same sign as the term in []-brackets, which means that (24) is positive if $\bar{\theta}^2$ exceeds some lower bound. Evaluate (24) at the smallest value for $\bar{\theta}^2$ for which $\tilde{e}$ is positive; that is, substitute (23) into (24), which yields

$$\frac{d(1-\phi)^2(1+\phi)}{1-\phi+k^2t^2\eta}. \tag{25}$$

Since (25) is positive, we have thus shown that (24) holds for $e = \tilde{e}$ whenever $\tilde{e}$ is positive. Since moreover the l.h.s. of (20) is increasing in e, it follows that γ is decreasing in e for any $e \geq \tilde{e}$. ■

Proof of Proposition 5: Inspection of (9) and (10) shows that β is increasing and γ decreasing in ρ if and only if η is decreasing in ρ. The derivative of η with respect to ρ is

$$-\frac{k^2\phi(1-\rho^2)(1-k^4\rho^2)+2(1-k^2)[\phi(1-k^2\rho^2)-\rho(1-k^2)]}{(1-k^4\rho^2)^2},$$

which is negative if the numerator (i.e. the expression in (11)), is positive. The l.h.s. of (11) equals $phi(2-k^2)>0$ for $\rho=0$, equals $-2(1-\phi)(1-k^2)^2$ for $\rho=1$, and its derivative of the l.h.s. of (11) with respect to ρ is

$$-2[(1-k^2)^2+\phi\rho k^2(3-2k^2+k^4(1-2\rho^2))],$$

which is always negative. This establishes the claim that (11) holds if and only if ρ is not too large. (To be completed)

Proof of Proposition 6: Inspection of (9) and (10) shows that β is decreasing and γ increasing in ϕ if and only if $(1-\phi)/\eta$ is decreasing in ρ. This is the case: the derivative of $(1-\phi)/\eta$ with respect to ρ is

$$-\frac{1}{\eta}-\frac{1-\phi}{\eta^2}\frac{\mathrm{d}\eta}{\mathrm{d}\phi}=-\frac{1}{\eta}+\frac{(1-\phi)\rho[2-(1+\rho^2)k^2]}{(1-k^4\rho^2)\eta^2},$$

which substituting for η from (2) reduces to

$$-(1-\rho)(1+k^2\rho)<0.$$

■

Proof of Proposition 7: By inspection of (14), $\mathrm{E}(Y)$ is increasing in $\bar{\theta}$, decreasing in d and e, and increasing in t, which in combination with Propositions 2 and 4 proves the result for these parameters. Also by inspection, $\mathrm{E}(Y)$ is increasing in k if and only if ηk^2 is, which was already shown to be the case in Proposition 3. Also, $\mathrm{E}(Y)$ is increasing in ρ if and only if η is increasing in ρ, which is the same condition for γ to be increasing in ρ, cf. Proposition 5. Finally, $\mathrm{E}(Y)$ is increasing in ϕ if $(1-\phi+k^2t^2\eta)/(1-\phi)=1+k^2t^2\eta/(1-\phi)$ is, or equivalently, if $(1-\phi)/\eta$ is decreasing in ϕ, which was already shown in Proposition 6. ■

References

Aghion, Philippe; and Jean Tirole (1997): Formal and Real Authority in Organizations, *Journal of Political Economy* **105** (1), 1–29

Baker, George (1992): Incentive Contracts and Performance Measurement, *Journal of Political Economy* **100**, 598–614

Baker, George (2000): Distortion and Risk in Optimal Incentive Contracts, *American Economic Review*, or working paper version?

Baker, George; Robert Gibbons and Kevin J. Murphy (1994): Subjective Performance Measures in Optimal Incentive Contracts, *Quarterly Journal of Economics* 1125-1156

Baker, George; and Bjorn Jorgensen (2003): Volatility, Noise and Incentives, Working paper

Brickley, James; Clifford Smith and Jerold Zimmerman (1995): The Economics of Organizational Architecture, *Journal of Applied Corporate Finance*, Vol. 8, No. 2, 19-31 (Summer 1995).

Demsetz, Harold; and Kenneth Lehn (1985): The Structure of Corporate Ownership: Causes and Consequences, *Journal of Political Economy* **93** (6) 1155–77

Dessein, Wouter (2002): Authority and Communication in Organizations, *Review of Economic Studies*

Feltham, Gerald A.; and Jim Xie (1994): Performance Measure Congruity and Diversity in Multi-Task Principal/Agent Relations, *Accounting Review* **69** (3) 429–453

Gibbs, Michael; and Alec Levenson (1999), The Economic Approach to Personnel Research, in: S. Grossbard-Shechtman (ed.), *The Expansion of Economics*, forthcoming

Hackman, J. Richard; and Greg R. Oldham (1980): *Work Redesign*, Addison-Wesley

Hayek, F.A. (1945): The Use of Knowledge in Society, *American Economic Review*

Holmström, Bengt (1979): Moral Hazard and Observability, *Bell Journal of Economics* **10** (1), 74-91

Holmström, Bengt; and Paul Milgrom (1987): Aggregation and Linearity in the Provision of Intertemporal Incentives, *Econometrica* **55** (2), 303–328

Holmström, Bengt; and Paul Milgrom (1991): Multi-task Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design, *Journal of Law, Economics and Organization* **7** (0), 24-52

Holmström, Bengt; and Paul Milgrom (1994): *American Economic Review*

Ittner and Larcker (1998): Innovations in Performance Measurement: Trends and Research Implications, *Journal of Management Accounting Research*

Jensen, Michael; and William Meckling (1992): Specific and General Knowledge, and Organizational Structure, in: L. Werin and H. Wijkander (eds.), *Contract Economics*, Oxford (Blackwell), 251–74

John, George; and Barton Weitz (1989): Salesforce Compensation: An Empirical Investigation of Factors Related to the Use of Salary Versus Incentive Compensation, *Journal of Marketing Research*

Laffont, Jean-Jacques; and David Martimort (2001): *The Theory of Incentives - The Principal-Agent Model*. Princeton and Oxford: Princeton University Press

Lafontaine, Francine; and Sugato Bhattacharyya (1995): The Role of Risk in Franchising, *Journal of Corporate Finance* **2**, 39–74

Lafontaine, Francine; and Margaret Slade (2000): Incentive Contracting and the Franchise Decision, in: K. Chatterjee and W. Samuelson (eds.), *Advances in Business Applications of Game Theory*, Kluwer Academic Press

Prendergast, Canice (2002): The Tenuous Tradeoff Between Risk and Incentives, *Journal of Political Economy*

Raith, Michael (2003): Competition, Risk and Managerial Incentives, *American Economic Review*