DISCUSSION PAPER SERIES

No. 4262

SPECIFIC KNOWLEDGE AND PERFORMANCE MEASUREMENT

Michael Raith

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP4262.asp

ISSN 0265-8003

SPECIFIC KNOWLEDGE AND PERFORMANCE MEASUREMENT

Michael Raith, University of Rochester and CEPR

Discussion Paper No. 4262
February 2004

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4262

February 2004

ABSTRACT

Specific Knowledge and Performance Measurement*

This Paper examines optimal incentives and performance measurement in a setting where an agent has specific knowledge about the consequences of their actions for the principal. I study incentive contracts in which the agent's compensation can be based on both 'input' measures closely related to the agent's actions, and 'output' measures closely related to the principal's pay-off. I argue that when the agent has specific knowledge (i.e. private information that is difficult to communicate) about how their actions contribute to the principal's pay-off, output-based pay encourages the agent to use their 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 actions on each task can be measured perfectly. Comparative statics results show how the optimal choice of performance measures and incentives depends on the agent's knowledge, environmental risk, technological uncertainty, and job complexity. The theory leads to several novel predictions, as well as new explanations for existing empirical findings.

JEL Classification: D82, J33 and M52
Keywords: distortion, incentives, input- vs. output-based pay, multitask agency model, performance measurement, risk and specific knowledge

Michael Raith
William E Simon Graduate
School of Business Administration
University of Rochester
Rochester, NY 14627
USA
Tel: (1 716) 275 8380
Email: raith@simon.rochester.edu

*I received many helpful suggestions from Jim Brickley, Luis Garicano, Francine Lafontaine, Kevin J Murphy, Oguzhan Ozbas, Canice Prendergast, Allison Raith, Andreas Roider, Joanna Wu, and Jan Zabojnik. Thanks also to seminar participants at the CEPR workshop ‘Incentives and Beyond-The Economics of Personnel and Organizations’ in Stockholm in May 2002, and at USC, where the paper was presented under different titles.

Submitted 08 November 2003

1 Introduction

Managers are expected to pursue activities and make decisions that enhance the value of their firms. They base their decisions on knowledge about market conditions and the firm that is difficult to communicate to others (termed "specific knowledge" by Jensen and Meckling, 1992), and thus usually work with little direct supervision. To ensure that managers' interests are aligned with their firms' objectives, it is necessary to provide them with appropriate incentives and to measure their performance. In spite of much research on incentive contracts, however, little is known about how to design incentives and performance measurement for agents with specific knowledge, largely because of the analytical difficulties involved. The purpose of this paper is to study this question.

The central premise of the paper is that in most principal-agent relationships, the agent has specific knowledge about how his actions contribute to the principal's objectives. This assumption contrasts sharply with the assumption made in most agency models that the principal cannot observe her agent's actions, but that both know equally well what actions the agent should be pursuing.

A principal can in general base the agent's compensation on two different kinds of performance measures, "inputs" and "output". Inputs are measures that are closely related to an agent's activities. For a salesperson, they might include hours worked, number of customers contacted, new accounts created, quality of advice, etc. Output, in contrast, measures the agent's contribution to the principal's objectives, such as sales or profit. It is usually influenced by factors outside the agent's control, and is therefore less closely correlated with the agent's actions than inputs. The distinction between inputs and output is familiar but takes on a novel meaning in this paper: in contrast with most moral-hazard models I assume the agent has specific knowledge of how the different inputs translate into output.

In designing optimal incentives, the principal faces a tradeoff. When all relevant actions can be measured without much noise (i.e. in the absence of multitask problems), the input measures combined provide a more accurate evaluation of the agent's actions than a noisy output measure, and thus minimize the exposure of the agent to risk. But basing compensation on some formula combining only input measures would not give the agent any incentive to use his specific knowledge. Instead, his actions would be entirely determined the compensation plan. With output-based compensation, in contrast, the agent is exposed to greater income risk but also has an incentive to use his knowledge in choosing his actions.

Formally, I study a two-task principal-agent model in which the principal faces *technological uncertainty* about the productivities of the tasks the agent works on. The agent receives private information about the productivities, which he is unable to communicate to the principal; his information is thus specific knowledge.[1] The agent then chooses his effort levels for both tasks based on his information. In a separate discussion, I argue that in this model, effort choice can be interpreted much more generally, e.g. as the choice between different investment projects.[2]

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. Measured output depends on the realized productivities of the tasks, the agent's effort devoted to each task, and *environmental uncertainty* (or *risk*) that is unrelated to productivity.

While technically a multitask model, the model abstracts from multitask measurement problems in the sense of Holmström and Milgrom (1991). Doing so simplifies the analysis but is also economically motivated. 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). Thus, firms often appear to have a much greater range of performance

[1] For models that allow for the upward communication of information, see e.g. Dessein (2002) or Marino and Matsusaka (2003).

[2] The latter argument is one reason for using a two-task framework, although most results of the paper can also be derived for a model with only one task. Another reason is to be able to derive predictions about job complexity (see Section 5.4).

measures at their disposal than they know how to make use of. To capture this situation in the simplest possible way, I assume that the principal can observe the agent's actions without noise but does not know how they contribute to her payoff. The main insights of the paper, however, remain valid even when multitask problems exist.

The model leads to a closed-form solution for the optimal compensation contract. When there is no technological uncertainty, or when the agent has no private information about it, the principal 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 base the agent's compensation at least partially on observed output, even if effort can be measured perfectly. This result helps to explain why firms tend to rely heavily on output performance measures even when a wealth of input measures are available.[3]

A comparative-statics analysis of the optimal contract leads to a numerous testable predictions. The optimal relative weight on output- vs. input-based pay will be larger...

1. ...the larger the value of managerial effort. Thus, for example, stock-based compensation should have a larger weight relative to some other forms of compensation the more competitive the product market.

2. ...the better the agent's knowledge. Since it is knowledge *relative* to the principal's that matters, the optimal choice of performance measures thus depends not only on the agent and the nature of the job, but also on who the principal is. The result leads to predictions about, for instance, how the performance measures used for CEO compensation depend on the firm's growth options and on its corporate governance, or about the use of subjective evaluations at different levels of a firm's hierarchy.

3. ...the lower the environmental uncertainty. This negative relationship between risk and incentives is familiar from standard principal-agent theory, cf. Holmström (1979). On the other hand, the relationship between (output-based) incentives and technological uncertainty is positive: the greater the variance of the productivities of the agent's tasks, the more valuable the agent's knowledge about them, and therefore the greater the optimal

[3] See Murphy (1999, p. 2521) for a similar observation and explanation.

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 considered.[4]

4. ... the lower the correlation between the tasks' productivities. The correlation can be interpreted as an inverse measure of *job complexity* since a lower correlation implies a greater probability that the agent should focus his effort on one task rather than both or none. Since this makes the agent's information more valuable for the principal, it is optimal to place a larger weight on output-based pay (unless the productivities are highly correlated). It follows that the weight on output measures in compensation plans should be positively related to measures of job complexity.

I consider three different variations of the basic model. In the first, I assume that input-based compensation is infeasible. This assumption reflects the fact that in many jobs, critical dimensions of the agent's actions cannot be measured. I show that in this case, the optimal output piece rate has exactly the same properties as the one in the basic model. Thus, the main insights of this paper do not hinge on the measurability of inputs (or the absence of multitask problems) but are general.

In the second variation, two agents work on one task each to produce a common output. Each agent's compensation can be based on the common output and on the agent's effort. Now, even if the output can be measured without risk, purely output-based compensation is costly as it leads to a free-rider problem in the agents' provision of effort. Nevertheless, if the agents have sufficiently important specific knowledge about how their tasks contribute to the firm's output, the optimal contract is again partly input- and partly output-based.

In the third variation, I assume the principal can base compensation on an additional signal of the tasks' productivities, which may also take the form of relative performance

[4] This point is also emphasized in papers by Zabojnik (1996) and Baker and Jorgensen (2003). The model studied here, however, is the first that leads to a closed-form solution for the contract as well as unambiguous comparative-statics results without requiring additional simplifying assumptions. See Sections 3.3 and 5.3 for a more detailed discussion.

evaluation. According to the informativeness principle (Holmström 1979), it is optimal to filter out influences on measured performance outside the agent's control, as long as these influences are irrelevant for the agent's actions. In contrast, the scope for filtering out changes in performance caused by events outside the agent's control is very limited when the agent must respond to those events (i.e. when the uncertainty is technological rather than environmental). This result may explain, for example, Bertrand and Mullainathan's (2001) finding that managers' compensation responds to changes in performance even when the changes are driven by observable events outside the managers' control.

To my knowledge, the model studied in this paper is the first to offer a tractable analysis of the role of specific knowledge for the design of optimal incentives. In contrast, other moral-hazard models allowing for hidden information are either analytically intractable or cumbersome, or focus on private information about measured performance rather than about the true value of the agent's actions. See Section 3.3 for a discussion of the related literature.

The paper thus closes a gap in the theory that has been pointed out by a number of observers (Murphy, 1999; Lambert, 2001). It also helps to explain the somewhat limited empirical success of standard agency models (Prendergast, 1999; Bushman and Smith, 2001). Finally, it bridges the gap between two strands of principal-agent theory that to date largely lie side by side, namely moral-hazard models in the tradition of Holmström (1979),[5] and more recent models of delegation whose emphasis on the role of private information goes back to Hayek (1945).[6] On this point, see also Section 10.

The paper is organized as follows. Section 2 discusses several applications of the distinction between inputs and outputs to which I refer in subsequent sections. In Section 3, I describe the formal model and compare it with the most closely related papers (see also Section 9). The equilibrium contract is derived in Section 4, and comparative statics results and empirical implications are derived in Section 5. In Section 6, I show that all

[5] See e.g. Holmström and Milgrom (1987, 1991), Feltham and Xie (1994), Baker (2000, 2002).

[6] See e.g. Aghion and Tirole (1997), Dessein (2002), Prendergast (2002), Stein (2002), Marino and Matsusaka (2003).

results about the optimal output piece rate derived in Section 5 remain valid when inputs cannot be measured. Section 7 modifies the basic model to one with two agents who work on one task each. Section 8 revisits the informativeness principle in the context of an agent with specific knowledge. In Section 9, I discuss how the paper relates to multitask models. Section 10 contains concluding remarks.

2 Inputs and Outputs: Applications

The distinction between "input"- and "output"-based pay is central to this paper. To many personnel economists, it is a quite natural one, cf. Lazear (1995, p. 20). In contrast, Baker (2002) has argued that any performance measure is either more or less noisy, and is either more or less well aligned with the principal's objectives. Performance measures are therefore more usefully thought of in terms of this two-dimensional continuum rather than in terms of a (one-dimensional) input-output distinction.

The distinction between inputs and outputs is much more meaningful when an agent has specific knowledge relevant to his actions. When the principal does not know how the agent's actions translate into output, precise measures of the agent's actions are *necessarily* less closely correlated with the principal's objectives. Conversely, measures closely aligned with the principal's objectives will be noisier measures of the agent's effort. The model in this paper applies to any situation in which available performance measures can be distinguished (along a continuous scale) in this way. Consider the following applications:

CEOs: In the case of CEOs, one can think of stock price as output, and accounting and non-financial measures as inputs.[7] Stock price is usually the best available measure of firm value, but it is also influenced by many factors outside a CEO's control. Accounting measures and non-financial (e.g. operational) measures, on the other hand, are less closely related to firm value but are more under the CEO's control because "managers understand

[7] Clearly, not all non-financial measures are inputs. For example, market share is a non-financial measure, but would hardly be considered an "input" relative to (output) measures such as share price or earnings.

and can ‘see’ how their day-to-day actions affect year-end profitability” (Murphy 1999, p. 2506).[8] If CEOs have specific knowledge of how their actions contribute to firm value, then shareholders face a tradeoff between undistorted but risky stock-based compensation, and compensation based on other measures that is less risky but also gives CEOs less incentives to use their knowledge to maximize firm value.

It is important to observe that whether a performance measure should be regarded as an input or output depends on what alternative measures are being considered. For instance, although Economic Value Added (EVA) is an accounting measure, it is a measure of value (output) creation when compared to the wide array of financial and non-financial performance measures provided by a “balanced scorecard”. Treating balanced-scorecard measures as inputs and EVA as an output is also consistent with practitioners’ views on the merits of using one or the other in compensation plans. Stern Stewart, the advocates of EVA, point out that 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, 1999). When firms do not know how scorecard measures translate into firm value, using a more aggregate measure such as EVA is a better way to motivate CEOs to use their specific knowledge.

Non-CEO managers: The “output” produced by division- or lower managers is often best captured by (accounting) measures of division profit. Relative to division profit (a financial measure), “inputs” may include a variety of operational measures.[9] Similarly,

[8] See also Sloan (1993), who argues that firms’ use of accounting measures in addition to stock price can be understood as a way to shield CEOs from market risk. Other possible reasons for using accounting measures are that they simply provide additional information (Lambert and Larcker 1987), or help balance a manager’s incentives across different activities (Feltham and Xie 1994). Watts (2003) argues that firms use accounting-based compensation to ensure that payments to managers or shareholders are made based on realized performance rather than the expected future performance captured by stock prices.

[9] Ittner and Larcker (1998) document an increase in the use of non-financial in addition to the more traditional financial performance measures over time. Ittner, Larcker and Rajan (1997) study the determinants of the use of financial and non-financial measures, testing hypotheses derived from the informa-

retail chains face a choice between compensating the managers of company-owned outlets based on sales figures or instead operational measures (such as "quality, service, cleanliness" in the classic McDonalds case, see Sasser and Pettway, 1974). Salespeople can be compensated based on the sales or profit they generated (output), but also on very detailed input measures such as new accounts created, hours worked, number of items sold in certain product categories, etc.

Input measures may also consist of subjective performance evaluations. In the theoretical literature, subjective evaluations are treated either as noisy signals of effort (see Prendergast 1999), or as noiseless output measures that are noncontractible and can therefore be used only as part of a relational contract (Baker, Gibbons, Murphy 1994). In some situations, however, it is more appropriate to think of subjective evaluations as input measures. That is trivially the case whenever measures that are obviously inputs are based on subjective evaluations, such as in the McDonald's case mentioned above. A more general point is that whenever managers have specific knowledge about how to do their job, then their superiors' subjective evaluations are likely to be based more on observable actions than on the managers' true contributions to the firm. Under this interpretation, subjective evaluation shields a manager from the noise contained in his objective measure of performance ("your unit had a bad year but we know you worked hard").[10] On the other hand, in the presence of specific knowledge the firm would not want to rely on subjective evaluation alone, which explains the prevalence of objective measures in executive compensation plans.[11]

3. Individual vs. group performance: Many managers' compensation is based

tiveness principle.

[10] See Gibbs et al. (2003a) for evidence that firms are more likely to pay bonuses based on subjective evaluation in periods with losses.

[11] That is to say, although subjective measures can be used only as part of relational contracts, as Baker, Gibbons, and Murphy (1994) have pointed out, one would expect subjective evaluations to play a much greater role in executive compensation than they do, if that were the only problem. Chances are that subjective measures are of only limited use when evaluators are not as well informed as those they evaluate.

on both measures of their individual performance (or that of the unit they manage) as well as on the performance of a group of units encompassing the manager's (where the group may be the entire firm), cf. Murphy (1999). Even when reliable measures of individual performance are available, it may still be optimal to use group measures if a manager's actions affect the group's performance in ways not already measured by individual performance (see Bushman, Indjejikian and Smith 1995, 1996; and Keating, 1997). A different rationale is explored in this paper: even when group performance does not provide additional information about effort, the firm may still want to use it in the manager's compensation plan if the manager has specific knowledge about how his actions affect the group. This idea is discussed in more detail in Section 7, where I extend the basic model to one with two agents.

To conclude, thinking in terms of inputs and outputs as defined in this paper is a novel way of looking at what determines the use of different performance measures. Most empirical research derives its hypotheses from either single-effort or multitask models. In single-effort models, the informativeness principle implies that two performance measures X and Y should both be used in a compensation contract to the extent that both are informative signals of the agent's effort (see e.g. Prendergast 1999). A much smaller literature focuses on multitask problems and investigates to what extent X and Y are used to induce effort on different tasks, or to what extent a performance measure Y is used to compensate for distortions caused by compensating for X (see Gibbs et al. 2003b). The examples above suggest a third way to think of two measures X and Y, namely as input and output measures; the empirical predictions follow from the results derived below.

3 Model and Related Literature

In this section, I first describe the model (Section 3.1). I then argue that the agent's actions can interchangeably be interpreted as as one of effort choice or project selection (Section 3.2). A discussion of the related literature concludes the section (Section 3.3).

3.1 Model

A principal hires an agent to produce an 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 $\min\{a_1\theta_1 + a_2\theta_2, 1\}$.

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. (Many of the model's parameters are normalized between 0 and 1 or -1 and 1, which significantly simplifies the subsequent calculations.)

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 total effort is constant across all states of nature, and the agent only decides how to allocate total effort between the two tasks. I argue in Section 5.4 that ρ can be interpreted as an inverse measure of job complexity.

Information about θ: The principal is assumed to know the expected productivity $\bar{\theta}$, 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.

The productivity vector $\boldsymbol{\theta}$ and the agent's specific knowledge about it should be interpreted very broadly: $\boldsymbol{\theta}$ may represent any kind of information about the firm's production technology, market conditions, or optimal responses to changes in those conditions. All that matters is that the agent has specific knowledge relevant to his (first-best) optimal level and allocation of effort, which is what the multiplicative terms $a_i\theta_i$ are meant to capture.

Agent's utility: The agent is risk-neutral; his utility is given by $w - d(\mathbf{a})$, where w is his total compensation (described below) and $d(\mathbf{a})$ 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 the agent 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 captured by $e \in [0, 1]$, a measure of *environmental uncertainty* or *risk*. 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, Chapter 4.3), but stands in contrast to e.g. the model of Holmström and Milgrom (1987, 1991).[12]

[12] Another source of uncertainty in addition to technological uncertainty and environmental risk is of course the randomness of Y itself. This randomness, however, is irrelevant for the results because both principal and agent are risk-neutral and therefore only care about the expected value of Y. The randomness of Y is merely a convenient way to map continuous actions into a set of two outcomes.

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, and is motivated by their use in practice. The assumption 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. But since the role of measurement errors for the design of optimal contracts is well known,[13] I ignore errors in measuring effort 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.

The agent is risk-neutral, but is also protected by limited liability. As is well known but will be further explained in subsequent sections, the combination of risk neutrality and limited liability is economically very similar to assuming that the agent is risk-averse. The agent's 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. As is standard in models with limited liability, I also assume that the agent's participation constraint is not binding; therefore, we can without loss of generality set $\alpha = 0$.

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 $\boldsymbol{\tau}$.

4. The agent chooses an effort vector $\mathbf{a}$.

[13] See e.g. the discussions in Prendergast (1999) and Lafontaine and Slade (2000).

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

To simplify the calculations and avoid tedious case distinctions, I restrict the analysis to parameters which lead to interior equilibrium solutions for all endogenous variables. 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:

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

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

$$(A3) \qquad 2(1-e)(1+k^2\rho)\eta kt[2d - (1+t)\bar{\theta}^2] - (1+t)(1+\rho)[2(1-e)\eta k^2t^2\bar{\theta}^2 - de(1-\phi)] > 0$$

Assumption (A1) ensures that the optimal γ in the contract is positive, and states that $\bar{\theta}$ must be sufficiently large relative to d. For given $\bar{\theta}$ and d, it states a lower bound to kt. Assumption (A2) ensures that for any realization of $\mathbf{s}$, 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.[14]

3.2 Effort Exertion vs. Project Choice

Many observers have noted the limitations of hidden-action models, with their focus on "effort" provision, for the question of how to induce managers to "make the right decisions" (Holmström 1992, Murphy 1999, Lambert 2001). On one hand, contract theorists have always understood that "effort" can be interpreted far more broadly than in terms of hours worked or calories burned. On the other hand, it is also clear that many managerial decisions, such as the choice of investment projects, have implications beyond their

[14] That parameters satisfying (A1)-(A3) exist can be proven by example: take $\bar{\theta} = 1$, $e = .1$, $t = k = .5$, $d = 1$, $\rho = 0$, and $\phi = .5$, in which case $\eta = 1$.

expected effect on profit, such as on risk taking. It may then be difficult or impossible to capture all relevant dimensions of decisions in a scalar or vector termed "effort".

Sometimes, however, a distinction between effort exertion and project choice is made largely based on the presumption that specific knowledge plays a role for decision making but not for effort choice. That is a feature of, for instance, the model of Athey and Roberts (2001), who argue that "motivating effort is done best by rewarding agents on precise measures of their effort... At the same time, ... getting the right investment choices may require that the decision-maker's rewards be tied to total value created" (p. 200).

But it is not immanent to the economic nature of principal-agent relationships that effort and project choice ought to be distinguished because of a difference in the role of the agent's knowledge. Instead, *both* minimizing risk and maximizing alignment are objectives for the design of *both* incentives to exert effort and incentives to choose the right projects. When agents are risk-averse, they should be shielded from risk they cannot respond to, irrespective of whether that risk pertains to measures of their effort or of their project choices. Similarly, their incentives should be aligned with the principal's objectives irrespective of whether their choice is about effort or projects. Thus, the choice of investment projects is meaningfully different from "effort exertion" only when there are implications beyond expected profit to consider.

The symmetry between effort choice and project selection can be seen most clearly in the model above by setting $\rho = -1$. This situation can be interpreted as one in which the agent must choose between two projects, where exactly one project is profitable. With perfect knowledge ($k = 1$), the agent will then choose one of two (symmetric) effort allocations, which can be reinterpreted as the choice of pursuing one of the projects. With less than perfect knowledge, the agent will generally hedge his bets and invest in both projects (depending on the information he receives), but that does not affect the interpretation. Thus, while most principal-agent models are concerned with either effort choice (see footnote 5) or project selection (see footnote 6), the model studied here fits both interpretations. Although most results derived below also hold in a single-task model, the explicit multitask framework used here better illustrates the equivalence of effort and project choice.

3.3 Related Literature

Multitask agency problems are studied in Holmström and Milgrom (1991), Baker (1992), Feltham and Xie (1994) and Baker (2000, 2002), and Datar, Kulp and Lambert (2001). Although technically a multitask model, my model differs from these papers not only in its emphasis on specific knowledge, but also in that it abstracts from any measurement problems, whereas they are the focus of papers cited. These differences are discussed in greater detail in Section 9.

Other models in which an agent has private information before choosing his action include Baker (1992), Baker, Gibbons and Murphy (1994), Baiman, Larcker and Rajan (1995), Lafontaine and Bhattacharyya (1995), Zabojnik (1996), Baker and Jorgensen (2003), and Shi (2003). In Baker (1992) and Baker, Gibbons and Murphy (1994), the agent learns information about how his effort is related to measured performance (and, in Baker (1992), to actual output). Both papers emphasize the distortionary consequences of private information, see also the discussion in Section 9.[15]

In the other papers mentioned, the agent obtains private information about the true productivity before choosing his action. Except for Baiman, Larcker and Rajan (1995), all are extensions of the continuous (static) model of Holmström and Milgrom (1987), and all focus on the relationship between risk and incentives rather than the role of specific knowledge. Solving models of this kind is notoriously difficult because the agent's income risk is a function of his effort, which in turn depends on private information that is learned only after the contract is written. Consequently, Baiman, Larcker and Rajan (1995), Lafontaine and Bhattacharyya (1995) and Shi (2003) rely on numerical simulations

[15] 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.

to obtain their predictions.

Only Zabojnik (1996) and Baker and Jorgensen (2003) obtain analytical results. Zabojnik (1996) obtains a closed-form solution for the optimal contract, but makes a simplifying assumption that the agent is risk-averse with respect to (environmental) risk but risk-neutral with respect to technological uncertainty. Baker and Jorgensen (2003) obtain formal comparative-statics results, although a closed-form solution for the optimal contract does not exist. They also distinguish between two types of uncertainty and allow for input- and output-related pay as I do here. Their main result is that incentives can be positively related to volatility (what I call technological uncertainty), whereas here they always are. 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.

4 Optimal Contract

The principal chooses the contract parameters to maximize her expected net profit $\mathrm{E}[Y - w]$, subject to the agent's participation constraint and his optimal choice of $\mathbf{a}$ which depends on his private information $\mathbf{s}$. I first derive the agent's optimal choice of effort, and then derive the optimal contract.

4.1 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 $\widehat{\boldsymbol{\theta}}(\mathbf{s})$, and subsequently chooses his optimal effort at stage 4 of the game, based on his expectations.

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

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

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

Proof: see the Appendix.

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

$$\mathrm{E}_\theta[w(\mathbf{a})|\mathbf{s}]-d(\mathbf{a})=\beta(a_1+a_2)+\gamma\mathrm{E}_\theta[y(\boldsymbol{\theta},\mathbf{a})|\mathbf{s}]-d(\mathbf{a}).^{16} \tag{4}$$

Assuming that

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

(which I will later show is satisfied when assumption (A3) holds), we have $\mathrm{E}_\theta[Y(\boldsymbol{\theta},\mathbf{a})]=\widehat{\theta}_1a_1+\widehat{\theta}_2a_2$, and together with (1), (4) expands to

$$\beta(a_1+a_2)+\gamma\{e/2+(1-e)[\widehat{\theta}_1a_1+\widehat{\theta}_2a_2]\}-\frac{d}{1+\phi}\left(a_1^2+a_2^2+2\phi a_1a_2\right). \tag{6}$$

This expression is strictly concave in $\mathbf{a}$, and maximization with respect to $\mathbf{a}$ leads to

$$a_i^*(\widehat{\boldsymbol{\theta}})=\frac{\beta}{2d}+\frac{(1-e)(\widehat{\theta}_i-\phi\widehat{\theta}_j)\gamma}{2d(1-\phi)} \tag{7}$$

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 (7) is valid), and that (7) leads to an interior solution for all realizations of $\boldsymbol{\theta}$ and $\mathbf{s}$. It is shown in the proof of Proposition 1 that (A2)-(A3) are the necessary and sufficient parameter conditions for this to be true.

4.2 Optimal Input- vs. Output-Related Incentives

At stage 1 of the game, 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 (7).

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

$$\beta^*=\frac{(1-\phi)de}{4(1-e)k^2t^2\eta\bar{\theta}} \quad \textit{and} \quad \gamma^*=\frac{1}{2(1-e)}-\frac{(1-\phi)de}{4(1-e)^2k^2t^2\eta\bar{\theta}^2}. \tag{8}$$

[16] All expected values without a subscript (as in (1)) are taken over the random variables that determine Y and y. Expected values with a subscript are taken over these variables as well as those specified in the subscript.

Proof: see the Appendix.

At one extreme, when pay is only input-based ($\beta > 0$, $\gamma = 0$), then it is clear from inspection of (7) that the agent's effort depends only on β and not on his private information $\widehat{\boldsymbol{\theta}}(\mathbf{s})$. This means that with input-based pay, the agent has no incentive to use his specific knowledge. Even though the agent is nominally free to choose his actions, they are in equilibrium fully determined.

At the other extreme, purely output-based pay ($\beta = 0$, $\gamma > 0$) is costly for the principal whenever measuring output is subject to errors ($e > 0$). This result is familiar when the agent is risk-averse, but holds for similar reasons when the agent is risk-neutral but protected by limited liability (on the following arguments, see also Laffont and Martimort, 2001, Chapter 4). 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.

With a risk-neutral agent and limited liability, 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.

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. When technological uncertainty and the agent's specific knowledge are sufficiently important, the optimal incentive contract is partly input- and partly output-based.

5 Comparative Statics and Predictions

I now study the effects of parameter changes on the optimal contract parameters β^* and γ^*, and discuss their empirical implications. As it turns out, most parameter changes lead to changes in β^* and γ^* in opposite directions. Thus, changes in the parameters affect not only the overall incentives provided to the agent, but also the optimal relative weights placed on input and output performance measures.

The results below also characterize how the agent's expected total wage changes with each parameter, allowing for adjustments of the contract parameters. The corresponding changes in expected output or total surplus are much more difficult to characterize and are therefore not considered. Numerical simulations suggest, however, that expected wage, output and surplus typically all move in the same direction with changes in the exogenous parameters.

5.1 Value of Managerial Effort

A change in the expected value of the productivity of effort, $\bar{\theta}$, has the same effect on γ^* as it has on the piece rate in a single-effort, single-measure model, but has the opposite effect on β^*:

Proposition 2 *(a) γ^* is increasing in $\bar{\theta}$; β^* is decreasing in $\bar{\theta}$. (b) Allowing for adjustments in γ^* and β^*, the agent's expected wage is increasing in $\bar{\theta}$.*

Proof: see the Appendix.

In other words, an increase in the expected productivity of effort leads to more output-based pay, as well as to a higher level of pay. In the literature, the value of managerial effort has been linked to, for instance, firm size and the degree of product market competition.

First, it is well documented that the level of executive compensation is positively related to firm size (Murphy 1999). Less clear is whether the pay-to-performance sensitivity is larger in large or in small firms; for an analysis, see Baker and Hall (2002). Irrespective of the details of appropriate measurement, Proposition 2 implies that whenever a change in firm size is predicted to lead to a greater pay-to-performance sensitivity, it should also lead to a larger relative weight on output-based pay. Empirically, this could take the form of a larger weight on stock-based incentives relative to accounting-based incentives, individual performance measures, or subjective performance evaluation (cf. Section 2).

Second, in Raith (2003) I argue that greater product market competition implies a higher marginal value for firms of improving their competitive position (relative to others) through cost reductions or quality improvements. Since managerial effort is required to implement such changes, greater competition therefore implies a higher value of managerial effort. Proposition 2 then predicts that, other things equal, greater product market competition should be associated with more output-based pay. Evidence consistent with this prediction is provided by Kole and Lehn (1999), who find that deregulation in the U.S. airline industry was followed by substantial increases in CEO pay as well as a shift towards more stock-based compensation.

5.2 Specific Knowledge

It is well known that specific knowledge is an important reason to delegate decision rights to agents. However, conditional on delegation, specific knowledge also affects the optimal incentive contract. It is intuitive from the tradeoff between input- vs. output-based pay (cf. Section 4.2) that the greater the agent's information advantage, the more the principal will want to base compensation on output:

Proposition 3 *(a) γ^* is increasing in k; β^* is decreasing in k. (b) Allowing for adjustments in γ^* and β^*, the agent's expected wage is increasing in k.*

Proof: see the Appendix.

It is worth pointing out that the agent's knowledge in this model is defined as knowledge *relative* to that of the principal. Thus, Proposition 3 implies that the relative weight on output-based pay is positively related to the *gap* between agent and principal in knowledge pertaining to the agent's job, and therefore also depends on characteristics of the principal.[17] Several testable predictions follow:

[17] It is more difficult to make predictions about the overall strength of incentives because that will depend on the principal's and agent's absolute, not just relative, levels of knowledge. For example,

First, the use of performance measures in an agent's compensation plan should be expected to depend on the education backgrounds of both the agent and the principal. For example, the optimal compensation plan for an engineer heading a manufacturing plant is more likely to include non-financial (input) measures when his boss, the division manager, also is an engineer, than when her background is different.

Second, Smith and Watts (1992) have suggested that both the level of compensation and the incentives provided to executives should be positively related to a firm's growth opportunities. They argue that the greater the growth options, the less shareholders are able to evaluate a CEO's decisions, and hence the more they will want to tie compensation directly to firm value by placing a larger weight on stock-based pay. Since k in the model measures the relative knowledge gap, better growth options are, according to this argument, associated with a larger value of k. We then obtain the prediction that firms with better growth options place a relatively larger weight on stock-based pay, consistent with the evidence of Smith and Watts.[18] Conversely, older firms and firms in more mature industries should be expected to place relatively less emphasis on stock-based pay due to a smaller knowledge gap between board and CEO.

Third, it is reasonable to assume that boards of corporations are less able to evaluate a CEO's performance (implying a larger k) the worse they are governed. This leads to the prediction that the relative weight on output- (e.g. stock-) based pay in CEO compensation should be inversely related to measures of the quality of corporate governance.

---

an increase in the principal's knowledge should in the context of the present model be interpreted as a *decrease* in k, leading to shift towards input-based pay. On the other hand, a better informed principal will normally also want to provide stronger overall incentives to the agent, just like she would if her monitoring costs decrease. This argument explains the apparent contradiction between Proposition 3 and the result of Baiman, Larcker and Rajan (1995) that "the compensation risk imposed on the business manager generally increases with the principal's relative expertise" (p.206).

[18]Bushman, Indjejikian and Smith (1996) find that the weight on individual performance measures increases with a firm's growth options, seemingly in contrast to the above prediction. However, the increase appears to occur at the expense of accounting-based rather than stock-based compensation. A possible explanation is that the better a firm's growth options, the less informative accounting measures are, implying a shift towards both stock- and individual-based compensation.

While I am not aware of any detailed study of this relationship, the prediction is consistent with two broad trends that have not previously been connected: a trend towards better corporate governance, brought about for instance by greater outsider presence on boards (cf. Hermalin 2003, Murphy and Zabojnik 2003); and a trend towards greater use of non-financial performance measures (cf. Ittner and Larcker 1998). The link between governance and k may not very tight, though: a greater presence of outsiders on the board is normally regarded as an improvement in governance, but may also imply an *increase* in the knowledge gap between board and CEO.

Finally, Murphy and Oyer (2003) find that executives at lower ranks in a firm are more likely to receive discretionary bonuses based on subjective evaluation than are CEOs. Murphy and Oyer explain this result by arguing that for those executives, it is difficult to find reliable (financial) measures of output, which is why available objective measures are often augmented by subjective performance evaluation. An alternative explanation, however, is that the knowledge gap between CEOs and boards is likely to be greater than that between lower-level executives and their superiors who evaluate them. If subjective evaluations can be interpreted as input measures (compared with objective financial measures, cf. Section 2), then Proposition 3 predicts that they are less likely to be used for CEOs than for lower-level executives, consistent with the evidence.

5.3 Uncertainty

A central prediction of principal-agent theory is that a risk-neutral principal provides weaker incentives to a risk-averse agent the noisier the measure on which the agent's compensation is based. However, empirical evidence of an inverse relationship between risk and incentives is rather scarce. Although studies supporting this prediction (e.g. Aggarwal and Samwick 1999) are well known, Prendergast (2002) concludes that they are greatly outnumbered by studies finding a positive or insignificant relationship between risk and incentives. This discrepancy between theory and evidence, also observed earlier by Zabojnik (1996) and Lafontaine and Slade (2000), has attracted much interest in recent research.

One cause 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. 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.

Another possible cause of the discrepancy can be found in an argument due to Demsetz and Lehn (1985): "In less predictable environments, however, managerial behavior ... figures more prominently in a firm's fortunes... Hence, noisier environments should give rise to more concentrated ownership structures." Prendergast (2002) studies a model in which 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.

The present model makes clear 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 plays a role very similar to that of risk in standard agency models. It reduces the effect of effort on measured performance, and thus leads to less effort. Technological uncertainty, on the other hand, captures the variance of the contribution of each task to output. Changes in e and t have opposite effects on optimal incentives:[19]

Proposition 4 *(a) γ^* is decreasing in e if $\partial E(y)/\partial e \geq 0$; β^* is increasing in e for all e. (b) γ^* is increasing in t, β^* is decreasing in t. (c) Allowing for adjustments in γ^* and β^*, the agent's expected wage is decreasing in e and increasing in t.*

Proof: see the Appendix.

[19] Similar results appear in Zabojnik (1996) and Baker and Jorgensen (2003). See Section 3.3 for a discussion of these papers.

To understand part (a), notice that in any moral-hazard model with only two possible outcomes, changing the variance of the outcome by changing the probabilities of the realizations will generally 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, other things equal, the principal would want to decrease γ^* if $\mathrm{E}(y)$ increases, and vice versa. The more interesting question is how γ^* changes with e when $\partial \mathrm{E}(y)/\partial e = 0$. Part (a) of Proposition 4 states that in this case, 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, on the other hand, unambiguously induces the principal to increase γ and decrease β. 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.

Two factors are responsible for the positive relationship between technological uncertainty and (output-based) incentives: first, the realization of $\boldsymbol{\theta}$ affects the agent's (individually optimal) effort because he obtains information about $\boldsymbol{\theta}$ *before* choosing effort, and because knowledge of the productivity of his effort is *relevant* for his effort choice (Shi (2003) uses the term "respondable risk"). This is a necessary condition because exposing the agent to risk is always costly for the principal and implies lower optimal incentives *unless* the realized state of nature influences the agent's actions.

Second, the agent's information pertains to the true productivity of his effort and thus his first-best optimal effort. While in choosing effort, the agent cares about measured performance y and not actual output Y, $\boldsymbol{\theta}$ affects y only through its effect on Y. Thus, the agent's use of his private information is always to the benefit of the principal. In this sense, the agent's information is "good" information.

In contrast, it is conceivable that the agent's information pertains only to measured performance and not to true productivity. The agent's information may then 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 on the principal's payoff is large. 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 may lead 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.

Empirically, it may be difficult to determine to what extent a particular source of risk is "respondable", especially since technological and environmental uncertainty are likely to be correlated. It may also be difficult to determine whether a manager's knowledge is good or bad in the sense used above, although one might expect that information about market conditions is "good" knowledge. Shi (2003) argues that firm- and industry-specific risk is more likely to be respondable than market-wide risk because CEOs are likely to have specific knowledge (relative to e.g. shareholders and analysts) about their firm and the industry, whereas they have little or no information advantage with regard to market risk. Consistent with the conjecture, Shi finds that CEO incentives are negatively related to market risk but positively to firm- and industry-specific risk. For a similar conjecture and results on managerial ownership, see Demsetz and Lehn (1985).

Part (b) of Proposition 4 ought to be distinguished from Prendergast's (2002) result that greater risk may lead to the simultaneous delegation of authority and the provision of incentives. That result leads to the prediction that risk and incentives are likely to be positively related if delegation decisions are not controlled for (e.g. if company-owned and franchised outlets of a chain are pooled), whereas controlling for delegation, they should be negatively related for the usual reasons. Here, in contrast, delegation of authority is already implicitly assumed, which is an appropriate assumption for most managerial occupations.

5.4 Job Complexity

The parameter ρ (the correlation of the θ_i) can be interpreted as an inverse measure of *job complexity*. The smaller ρ, the more likely the θ_i will be different, and hence the efficient levels of effort for each task as well. Intuitively, a more complex job (smaller ρ) is one in which the tasks the agent should work on are more likely to differ from one point in time to the next, implying that the optimal allocation of total effort across the two tasks is more difficult to specify in advance.[20]

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

$$k^2\phi(1-\rho^2)(1-k^4\rho^2)+2(1-k^2)[\phi(1-k^2\rho^2)-\rho(1-k^2)]\geq 0. \tag{9}$$

There exists $\widetilde{\rho}\in(0,1)$ such that (9) holds if and only if $\rho\leq\widetilde{\rho}$. A sufficient condition for (9) to hold is that $a_1^*(\boldsymbol{\theta})$ is larger when the agent obtains a bad signal on task 2 than when he obtains a good signal.

(b) Allowing for adjustments in γ^* and β^*, the agent's expected wage is decreasing in ρ under the conditions stated in part (a).

Proof: see the Appendix.

Part (a) of Proposition 5 states that if the productivities of the tasks are not already highly correlated, a decrease in ρ will lead to an increase in γ^* and a decrease in β^*. The intuition is similar to that 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

[20] It would be inappropriate to speak of ρ as a complementarity parameter because the tasks are independent in the sense that the cross-partial derivatives of a_1 and a_2 in the expected value of Y are zero.

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.[21]

Proposition 5 leads to the prediction that the weight on output-based incentives should vary positively with measures of job complexity. For papers employing such measures, see e.g. John and Weitz (1989), Van Ophem, Hartog and Vijverberg (1993), Drago and Garvey (1998), and Ortega (2003). 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, one would expect to find a clearer positive relationship between job complexity and incentives.

The formal similarity between the comparative statics with respect to knowledge and with respect to job complexity has a parallel in empirical measurement. When direct information about the knowledge gap between an agent and his principal is not available, it is reasonable to use measures of job complexity as proxies for the importance of specific knowledge. Fortunately, this approach causes no problems as far as the theory developed here is concerned, as the effects of increases in k and decreases in ρ are very similar.

Regarding the relationship between job complexity and wages, there are two conflicting views. One view is that a more complex job requires greater effort or in other ways increases a worker's disutility. The worker will then want to be compensated accordingly, leading to the prediction that wages are positively related to job complexity. Van Ophem, Hartog and Vijverberg (1993) find evidence supporting this prediction. A very different view is expressed by Hackman and Oldham (1974) in their influential contribution to the Human Relations school of personnel research. Hackman and Oldham argue that workers' productivity depends on their "intrinsic motivation". Among the ways in which firms can increase intrinsic motivation is to give workers a greater variety of tasks. This view suggests that job complexity increases workers' utility and therefore should be negatively

[21] The incentive to increase the weight on γ^* is only partially offset by the fact that with lower ρ, the agent's equilibrium effort levels a_1 and a_2 are already less correlated even for given β and γ. To see that they are, evaluate $a_1 - a_2$ using (7) and (3). The derivative of this expression with respect to ρ turns has 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 $s_2 = -1$ (note that $a_1 - a_2 = 0$ if $s_1 = s_2$).

related to wages.

Part (b) of Proposition 5 predicts a positive relationship between job complexity and wages. This prediction is consistent with the evidence of Van Ophem et al. (1993), but its interpretation is quite different from the argument of Van Ophem et al. Instead, greater job complexity implies a greater reliance on output-based pay, which is generally associated with a higher expected wage (cf. Propositions 2 through 5).[22]

5.5 Managerial Preferences

The last part of this section looks at the effects of changes in the agent's utility function on optimal incentives. As characteristics of preferences are difficult to measure, however, the results lend themselves to testable predictions less readily than the previous ones.

As one might expect, the effects on optimal incentives of changes in the disutility parameter d are the opposite to those of changes in $\bar{\theta}$, cf. Proposition 2:

Proposition 6 *(a) γ^* is decreasing in d, β^* is increasing in d. (b) Allowing for adjustments in γ^* and β^*, the agent's expected wage is decreasing in d.*

Proof: Straightforward by differentiating β^* and γ^* in (8) and $E(W)$ in (21) with respect to d. ■

More interesting is the effect of changes in ϕ. This parameter measures how substitutable the two tasks are from the agent's point of view. For $\phi \to 1$, the agent views both tasks as perfect substitutes and is likely to devote all of his effort to the task he believes to be the most productive, based on his private information. The lower ϕ, the more the agent would prefer to work on both tasks rather than just one.

Thus, while ρ is an inverse measure of *objective* job complexity (or task variety, in the language of Hackman and Oldham), ϕ can be interpreted as an inverse measure of

[22] Numerical simulations suggest that the expected equilibrium disutility also increases with job complexity, consistent with Van Ophem et al. However, this effect is largely driven by the change in optimal incentives. The direct effect (i.e. before adjusting the incentives) of job complexity on disutility is only very small, and the direct effect on expected effort is zero.

subjective task variety.[23] The effects on optimal incentives of changes in ϕ are the opposite to those of changes in ρ:

Proposition 7 *(a) γ^* is increasing in ϕ, and β^* is decreasing in ϕ. (b) Allowing for adjustments in γ^* and β^*, the agent's expected wage is increasing in ϕ.*

Proof: see the Appendix.

Intuitively, the lower ϕ, the less the agent's optimal allocation of effort depends on his private information about the productivities of the tasks. This can be seen by evaluating $a_1 - a_2$ using (7) 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.

Part (b) of Proposition 7 states that the expected wage is decreasing in *subjective* job complexity, which is consistent with the prediction derived from Hackman and Oldham's (1974) arguments. Propositions 5 and 7 thus offer a potential resolution to the two conflicting views on the relationship between job complexity and incentives discussed in Section 5.5: in this model, the expected wage is positively related to objective job complexity as measured by ρ, but negatively related to subjective job complexity as measured by ϕ. In both cases, however, changes in job complexity affect the agent's effort, disutility and expected wage primary through the adjustment of optimal incentives, not directly.

6 Unobservable Inputs

In many jobs, input-based compensation is infeasible because important dimensions of the agent's job cannot be measured. Even if effort exerted on some tasks can be measured, basing compensation on those measures may still be undesirable if doing so would lead to too little effort spent on unmeasurable but important tasks, cf. Holmström and Milgrom

[23] The two parameters are only "in a way" counterparts because ϕ is a measure of substitutability of task, whereas ρ is a measure of the correlation of the productivites, cf. footnote 20.

(1991). This section shows that even when input-based compensation is infeasible, all results about γ^* derived in Section 5 remain valid.

When the agent's compensation contract is restricted to the form $w = \gamma y$, the conditions for the existence and uniqueness of an interior solution to the principal's contracting problem change to

$$
\begin{aligned}
&(A1') \quad && 2(1-e)(1-\phi+\eta k^2 t^2)\bar{\theta}^2 > (1-\phi)de,\\
&(A2') \quad && (1-\phi)(1-k^2\rho) > (1+\phi)(1-\rho)kt,\\
&(A3') \quad && 4d(1-e)(1+k^2\rho)(1-\phi+k^2t^2\eta)\\
& && -(1+t)(1+k^2\rho+(1+\rho)kt)(2(1-e)(1-\phi+\eta k^2t^2)\bar{\theta}^2 - de(1-\phi))
\end{aligned}
$$

We then have:

Proposition 8 *Under assumptions (A1')-(A3'), the principal's optimal value of γ is given by*

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

Proof: see the Appendix.

It can then be shown that all properties of γ^* stated in Propositions 2-7 carry over to the case of only-output-based compensation:

Proposition 9 *γ^* as given by Proposition 8 is increasing in $\bar{\theta}$, k and t, decreasing in e whenever $\partial E(y)/\partial e \geq 0$, decreasing in ρ if and only if (9) holds, decreasing in d and increasing in ϕ.*

Proof: see the Appendix.

The effects of changes in the value of effort, the agent's disutility, and the level of risk on the optimal piece rate are familiar from standard agency models. In addition, Proposition 9 states that optimal incentives are stronger the greater the technological uncertainty and the better the agent's knowledge about it. Also, optimal incentives are stronger the greater the complexity of the job and the more the agent regards the two tasks as substitutes.

7 Two Agents

The performance of any organization is the outcome of the actions of its members. When the performance of agents working in an organization cannot be measured individually, it may be necessary to provide agents with incentives based on the performance of the group (which may be the entire organization) they are part of. But even when measures of individual performance are available, firms use both measures of individual and group performance to compensate executives (cf. Ittner and Larcker, 1998; Murphy, 1999). What principles should guide the choice of weights placed on those measures?

When specific knowledge plays no role, providing incentives based on group performance in addition to individual performance may be desirable for two reasons. First, group performance may be informative about an agent's (scalar) effort beyond the information provided by individual performance measures. Second, when there are multiple tasks, it may be necessary to give an agent incentives to direct his effort to tasks that enhance group performance but not necessarily individual performance. Both arguments suggest that group performance is likely to receive a greater weight in a compensation plan when there are greater interdependencies between the agents or subunits of the group. For example, firm performance should have a larger weight in the compensation of division managers the greater the interdependencies between different divisions of the firm. This prediction is tested, and supported, by Bushman, Indjejikian and Smith (1995) and Keating (1997).

Here, I show that when agents have specific knowledge about how their units contribute to the performance of the group, then group incentives may be desirable even if there are no externalities, and if individual performance can be measured without noise. To see this, consider the following modification of basic model.

Instead of one agent who works on two tasks, suppose there are now two agents, 1 and 2, who work on one task each. Each agent i receives a signal s_i about the productivity of task i, and thereupon chooses his effort a_i. The disutility of choosing a_i is da_i^2; there is no longer an interaction parameter ϕ. Agent i's compensation is given by $w_i = \frac{\alpha}{2} + \beta a_i + \frac{\gamma}{2} y$. The sum of wages for both agents thus equals $\alpha + \beta(a_1 + a_2) + \gamma y$, as in the original model. All other assumptions of the model are the same as before. The conditions for the existence and uniqueness of an interior solution to the principal's contracting problem change to

$$(A1'') \qquad (1-e)(2k^2t^2-1)\bar{\theta}^2 - 2de > 0,$$

$$(A3'') \qquad \{8k^2t^2 + [2k(1+t) - 8k^2t - 1]et - 1\}d > (1-e)kt(1+t)(3kt + 2k^2t^2 - 1)\bar{\theta}^2$$

(a counterpart to (A2) is not needed). We can then show:

Proposition 10 *Under assumptions (A1") and (A3"), the optimal contract parameters are given by*

$$\beta^* = \frac{3de + (1-e)(1+k^2t^2)\bar{\theta}^2}{(1-e)(8k^2t^2-1)\bar{\theta}} \quad \textit{and} \quad \gamma^* = 2\frac{(1-e)(2k^2t^2-1) - 2de}{(1-e)^2(8k^2t^2-1)\bar{\theta}^2}. \tag{11}$$

Proof: see the Appendix.

As is usual with team production, compensating agents based on group performance leads to a free-rider problem. In contrast to the single-agent model above, therefore, it is optimal to base compensation in part on input ($\beta^* > 0$) even if output can be measured without noise ($e = 0$) and there is thus no risk for the agents. Without specific knowledge, there would also be no reason to use output-based incentives at all because by assumption a_i can already be perfectly measured; i.e. none of the reasons discussed above for using group-based incentives apply. But if the agents have sufficiently important specific knowledge about how their individual performance contributes to the performance of the group, then it is optimal to use group-based incentives as well (a necessary condition is $2k^2t^2 > 1$). Most results of Propositions 2-7 carry over to the two-agent case:

Proposition 11 *γ^* as given by Proposition 10 is increasing in $\bar{\theta}$, k and t, and decreasing in e and d. β is increasing in $\bar{\theta}$, decreasing in k and t, and increasing in e and d.*

Proof: see the Appendix.

The only differences with Propositions 2-7 are that γ^* is now always decreasing in e, and that β^* is increasing in $\bar{\theta}$. Two predictions that can be derived from the effect of k on β^* and γ^* are the following:

First, interpreting measures of job complexity as measures of k,[24] we obtain the prediction that the weight on group-based incentives should be increasing in job complexity, while the weight on individual incentives should be decreasing. That is precisely what Ortega (2003) finds in his study of the determinants of incentive compensation from survey data on European employees (as Ortega points out, the first result is consistent with a standard agency model, but the second is not).[25]

Second, Proposition 11 suggests a complementarity between the use of group incentives and the adoption of "open-book management". When employees have little knowledge of how their activities contribute to the performance of their team or firm, there is little to gain from tying their compensation to group performance. Open-book management helps employees to understand the processes in their group and their own contribution to it, and thus implies a larger value of k. Proposition 11 thus leads to the prediction that firms or teams with open-book management are more likely to provide group-based incentives.

8 Informativeness Revisited

The informativeness principle (Holmström 1979) implies that optimal contracts should make use of all information about the agent's effort not already contained in other performance measures that are used. In particular, influences on measured performance that are unrelated to the agent's effort should be filtered out. It is straightforward that the informativeness principle continues to apply to the present model as far as information about environmental risk is concerned. The following discussion turns to the more interesting question of how optimal incentives are affected when the principal has information about the uncertain production technology $\boldsymbol{\theta}$ which interacts with the agent's effort.

[24] See Section 5.2. Note that although ρ is still part of the model, it is irrelevant for the agents' effort choice, and therefore also irrelevant for optimal incentives.

[25] Some caution is due in interpreting Ortega's evidence as supportive of the present model, because Ortega studies the simultaneous choice of delegation and incentives, whereas the model here takes delegation as given.

To simplify the analysis and facilitate comparison with standard agency models, I will assume that compensation can be conditioned only on output, cf. Section 6. In some cases considered below, I will also assume that there is only one task with productivity θ, which is equivalent to the case $\rho = 1$. As will become clear, how information about $\boldsymbol{\theta}$ affects optimal incentives depends on two factors, how exactly effort and $\boldsymbol{\theta}$ interact, and much the principal knows about the interaction.

8.1 Three Scenarios

Suppose there is only one task, and expected output is given by $\mathrm{E}[Y] = a\theta$, with $\theta \in \{\theta_H, \theta_L\}$. For the case in which the principal does not know θ but the agent does, the optimal piece rate γ^* is given by Proposition 9 for $\rho = k = 1$. If, in contrast, the principal can condition the contract on information about θ, she will choose a piece rate γ_H if $\theta = \theta_H$, and γ_L if $\theta = \theta_L$, with $\gamma_H > \gamma^* > \gamma_L$. Thus, compared to the case of an uninformed principal, the agent will face stronger incentives in good states of the world and weaker incentives in bad states. These incentives give rise to a convex pay-performance relationship, consistent with the evidence of Garvey and Milbourn (2003).

Ideally, the principal would still like to insure the agent against income risk resulting from fluctuations in θ, since θ is outside the agent's control. Thus, the principal might want to include a term $-\delta\theta$ with $\delta > 0$ in the agent's wage function. With limited liability, however, that is not possible since the agent's main bonus γy can always end up being zero because of bad luck. In a model with risk aversion instead of risk neutrality with limited liability, the latter conclusion would not be as stark; i.e. some amount of filtering would likely be optimal. Nevertheless, the overall effect on the sensitivity of pay to measured performance is likely to be similar as with risk neutrality and limited liability.

Now suppose there is still only one task, but that the productivity of effort is *negatively* related to θ. Of course, in the model it is assumed that optimal effort and θ are positively related, which is without loss of generality since values of θ can always be appropriately relabeled. For any *given* measure of θ, however, the relationship is not so obvious. For instance, if θ is a measure of market demand, it will depend on the particular circumstances whether the value of effort is highest when demand is high, or when demand is low. If optimal effort is negatively related to θ, then the conclusions are the opposite of those above: incentives are stronger in the bad state of the world than the good one, giving rise to a concave pay-performance relationship.

The perhaps most realistic case, however, is one where different states of the world call for different optimal actions, but in a way that cannot be associated with "high" or "low" effort. Specifically, suppose that expected output is given by $a_1\theta_1 + a_2\theta_2$ as in the basic model, with $\rho = -1$. In this case, (first-best-)optimal effort takes only two values, depending on whether $\boldsymbol{\tau} = (1, -1)$ or $\boldsymbol{\tau} = (-1, 1)$. Now, there is no scope for using information about $\boldsymbol{\theta}$ in the incentive contract, unless compensation can be conditioned on input measures as well!

8.2 Rewarding Agents for "Luck"

The most important conclusion from the above discussion is that observable shocks outside the agent's control will not be filtered out of the performance measure if the agent is expected to respond to them. In other words, to use the language of Bertrand and Mullainathan (2001), compensation will generally depend as much on a "lucky dollar" (a change in performance due to observable shocks outside the agent's control) as on a "general dollar". The model thus provides an alternative explanation for Bertrand and Mullainathan's finding that CEO compensation in oil firms responds as much to changes in firm performance due to oil price shocks as to general changes in firm performance.[26]

Bertrand and Mullainathan also find that firms are more likely to reward a "lucky dollar" the worse they are governed. They interpret this finding as evidence of the idea that CEOs in poorly governed firms are able to capture the pay-setting process. An alternative

[26] Bertrand and Mullainathan consider this explanation but ultimately reject it, arguing that if specific knowledge were the explanation, one would expect to see only unusually successful CEOs rewarded for luck, contrary to the data. However, this counterargument does not hold in the present model, since even if the average CEO is no better able than others to forecast changes in market conditions, he is still likely to have specific knowledge of how best to respond to them.

explanation derived from present model, however, is that in more poorly governed firms there is a larger knowledge gap between the board and the CEO, in which case it is optimal to place a relatively larger weight on firm performance in the CEO's compensation plan.

While I assumed in the model that technological uncertainty and environmental risk are independent, in reality they are likely to be correlated. In other words, many events outside the agent's control require the agent to respond, but also affect measured performance for reasons unrelated to effort. To what extent changes in performance due to observable shocks should be filtered out then depends on the relative importance of technological uncertainty and risk. For example, although oil price shocks require CEOs of oil firms to respond, they arguably also affect firm performance for reasons entirely unrelated to CEOs' actions. Therefore, while it is not surprising that CEOs are rewarded for luck to some extent, it is surprising that they are (as Bertrand and Mullainathan find) rewarded as much as for unexplained changes in firm performance. Thus, the present model may not be able to fully explain Bertrand and Mullainathan's finding.

8.3 Relative Performance Evaluation

When one agent's performance is correlated with that of another, one way of filtering out influences on measured performance that are unrelated to the agent's effort is to use relative performance evaluation (RPE), cf. Holmström (1982). When agents have specific knowledge, however, peer performance inevitably depends on both technological uncertainty and environmental risk, making it difficult to infer information even when the two sources of uncertainty are uncorrelated. The following discussion is therefore somewhat more tentative than the previous arguments.

In the case of pure technological risk, the conclusions of Section 8.1 apply without change. When the performance of both firm A and firm B is positively related to θ, then B's performance can be used as a proxy for θ. When in addition optimal effort in both firms is positively related to θ, then firm A's optimal contract is one in which *the incentives* for A's agent (not just the level) vary positively with B's performance. The reverse holds when optimal effort in both firms is negatively related to θ. Finally, when optimal actions depend on $\boldsymbol{\theta}$ but overall effort and performance do not, then information about B's performance is of no use to A. Overall, in the case of pure technological risk, the scope for RPE seems very limited.

More realistically, events outside the agent's control both require the agent to respond, and affect measured performance for reasons unrelated to effort. Here, the main conclusion of Section 8.2 still holds: the desirability of using RPE depends on the importance of the environmental-risk component of events outside the agent's control, relative to the importance of the agent's responses to those events.

On the other hand, the usefulness of RPE will also depend on the correlation between the performances of firms A and B. The correlation is higher the greater the correlation of the firms' environmental risk, but also the greater the correlation of the firms' technological risk. This observation implies that in identifying peers for the purposes of RPE, a firm (and researchers) might want to look for other firms that are similar both in terms of the exogenous shocks they face and in their ability to respond to them. For a similar argument and evidence consistent with it, see Albuquerque (2003).

9 Relationship With Multitask Models

The formalization of multitask problems has been one of the most important extensions of classical principal-agent theory, cf. Gibbons (1998). The central insight of the multitask literature is that incentive pay may lead to dysfunctional responses by agents if available performance measures are not perfectly aligned with a principal's objectives. This may occur if an agent has different tasks to perform and if there are differences in how easily performance on those tasks can be measured, cf. Holmström and Milgrom (1991). Alternatively, the agent may have private information about how measured performance depends on his actions, cf. Baker (1992).[27] In both cases, it may be optimal for the

[27] Although in Baker's (1992) model there is only a single task, I consider this paper, as well as Baker, Gibbons and Murphy (1994), as part of the "multitask" literature. Cf. also Gibbons (1998) and Prendergast (1999).

principal to offer low-powered incentives to counteract the agent's incentive to direct his effort in ways that maximize his compensation but not necessarily the principal's profit.

Feltham and Xie (1994), Datar, Kulp and Lambert (2001) and Baker (2000, 2002) extend these models to investigate a principal's choice between different (aggregate) performance measures. These papers argue that a principal typically faces a tradeoff between performance measures that involve little risk for the agent but are also distortionary, and measures that are better aligned with the principal's objectives but are risky. In particular, if the principal cannot measure an agent's performance on all relevant tasks, she may prefer to provide incentives based on broader and less distortionary, but noisier measures, rather than provide no incentives at all.

This paper differs from the multitask literature in several respects. The first concerns the role of an agent's private information. Multitask models with private information (Baker 1992, Baker, Gibbons and Murphy 1994) assume that the agent has information about how his effort is related to measured performance, and emphasize the resulting incentives to "game the system". Here, in contrast, the agent has private information that is directly relevant for his (first-best) optimal actions and only indirectly to measured performance. Most information about market conditions or the firm's production technology would appear to be of this kind. In this kind of setting, the principal wants to encourage the use of private information rather than suppress it.

Second, specific knowledge is a reason for using output-based incentives *in addition* to problems that may arise because of multitask problems. The models of Feltham and Xie (1994) and Baker (2002) predict a greater weight on output measures the more distortionary the available input measures are. If this were the only problem, firms could attempt to solve it by developing new ways to measure agents' inputs. Balanced scorecards, customer and employee surveys, and integrated information systems all represent efforts to obtain information about the minutiae of employees' activities. In practice, however, the abundance of information has proven to be of only limited use for the design of incentive plans. As I have argued in this paper, the root of this problem is that firms' principals simply do not know very well how their employees' actions contribute to firm performance, whereas employees themselves very often have an information advantage. Consequently, the information gathered through balanced scorecards etc. may help to guide managers in their decision making, but is bound to be of only limited use for performance evaluation.

Third, a tradeoff between distortion and noise is postulated in Feltham and Xie (1994) and Baker (2000, 2002), but arises endogenously here. In the models of Feltham and Xie and Baker, the distortion associated with any performance measure is parametric, i.e. exogenous. Baker (2002) 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. Here, in contrast, a tradeoff arises endogenously as a result of specific knowledge. Output-based compensation is risky, but aligns the agent's objectives with the principal's. Input-based compensation is safe, but distortionary in the sense that 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.

10 Concluding Remarks

This paper attempts to fill a gap in the theory of incentives that has been noted for a long time. Standard agency models assume that a principal cannot observe her agent's actions but knows what these actions should be. Yet in many of the most important occupations, the problem is the opposite one: "the reason shareholders entrust their money to selfinterested CEOs is based on shareholder beliefs that CEOs have superior skill or information in making investment decisions" (Murphy 1999, p.2521), whereas the observability of those decisions is a less salient problem.

Meanwhile, a second strand of literature has developed that recognizes the importance of specific knowledge but focuses almost entirely on the question of when a principal should delegate decision rights to an agent (see the references in footnote 6). In this literature, incentives are either absent, or appear in the background as costs of delegation. Although both the theory of incentives and the literature on delegation shed light on agency relationships within firms, their perspectives on them are very different and without overlap.

The theory of incentives has been unable to account for the role of specific knowledge because models that combine hidden actions with hidden information are notoriously difficult to handle. Simple extensions of workhorse principal-agent models quickly prove to be analytically intractable. To my knowledge, this paper is the first to study a model of incentive contracting in the presence of specific knowledge that leads to simple, closed-form solutions and intuitive comparative statics.

The model helps to explain why some predictions derived from traditional agency models have found only scant empirical support, and why incentive contracts in practice have a strong weight on "output" measures that is difficult to reconcile with standard theory. More importantly, the model generates many novel testable predictions, and in particular relates properties of incentive contracts to the gap between a principal and an agent in knowledge pertaining to the agent's job. At a more methodological level, the model illustrates that the real problem with standard agency models is not that the notion of "effort" they rely on is too narrow, but that specific knowledge is ignored. Once it is incorporated in the analysis, effort exertion and decision making in general become interchangeable terms, at least unless decision making also covers dimensions beyond expected profit such as risk taking.

The present paper makes a first step towards integrating specific knowledge, and thus some insights from the literature on delegation, into the theory of incentives. Much more remains to be done to bridge the gap between these strands of the literature. For instance, in contrast to what I have assumed here, the agent can often communicate at least some of his knowledge to the principal, which interacts with the optimal assignment of decision rights as well as with the optimal design of incentives. At any rate, bridging the gap in the literature is important because the assignment of decision rights cannot be isolated from the problem of how to provide incentives to agents with decision rights, and how to measure their performance. Conversely, incentives and performance measurement are difficult to understand unless one takes into account why principals delegate decision rights to agents in the first place, namely because of their knowledge.

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

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

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

$$\begin{aligned}
\pi &= Y(\boldsymbol{\theta}, \mathrm{s}) - w(\boldsymbol{\theta}, \mathrm{s}) \\
&= Y(\boldsymbol{\theta}, \mathrm{s}) - \beta(a_1(\mathrm{s}) + a_2(\mathrm{s})) - \gamma \left[\frac{e}{2} + (1 - e)Y(\boldsymbol{\theta}, \mathrm{s})\right] \\
&= [1 - \gamma(1 - e)]Y(\boldsymbol{\theta}, \mathrm{s}) - \beta(a_1(\mathrm{s}) + a_2(\mathrm{s})) - \frac{\gamma e}{2}. \qquad (12)
\end{aligned}$$

To obtain the expected value of (12) over $(\boldsymbol{\theta}, \mathrm{s})$, we need to evaluate the expected values of $a_1(\mathrm{s}) + a_2(\mathrm{s})$ and $Y(\boldsymbol{\theta}, \mathrm{s})$. First, notice from (7) that

$$a_1 + a_2 = \frac{1}{2d}[2\beta + (1 - e)(\widehat{\theta}_1(\mathrm{s}) + \widehat{\theta}_2(\mathrm{s}))\gamma,$$

where the expected value of $\widehat{\theta}_1(\mathrm{s}) + \widehat{\theta}_2(\mathrm{s})$ over s is $2\bar{\theta}$. Hence, the expected value of $a_1 + a_2$ in (12) over s is given by

$$\frac{\beta + \gamma(1 - e)\bar{\theta}}{d}. \qquad (13)$$

Next, we have

$$
\begin{aligned}
Y(\boldsymbol{\theta}, \mathbf{s}) &= a_1^*(\mathbf{s})\theta_1 + a_2^*(\mathbf{s})\theta_2 \\
&= \frac{\beta}{2d}(\theta_1 + \theta_2) + \frac{(1-e)\gamma}{d(1-\phi)}\left[\theta_1(\widehat{\theta}_1 - \phi\widehat{\theta}_2) + \theta_2(\widehat{\theta}_2 - \phi\widehat{\theta}_1)\right] \qquad (14)
\end{aligned}
$$

(arguments of $\widehat{\theta}_i$ omitted). The expected value of $\theta_1 + \theta_2$ in the first term is $2\bar{\theta}$. Obtaining the expected value of the term in []- brackets requires evaluating the term for each of the 16 permutations of $(\theta_1, \theta_2, s_1, s_2)$ and weighting it with the corresponding probability of the permutation, given by $(1 + ks_1t_1)(1 + ks_2t_2)(1 + \rho t_1t_2)/16$. The computation is complex, but the result is very simple, it is $2(1 - \phi + k^2t^2\eta)$, with η as defined in (2). Thus, the expected value of (14) is

$$
\frac{\bar{\theta}\beta}{d} + \frac{(1-e)\gamma\bar{\theta}^2}{d(1-\phi)}(1 - \phi + k^2t^2\eta) \qquad (15)
$$

Substitute (13) and (15) into (12) to obtain the principal's expected profit:

$$
[1 - \gamma(1-e)]\left[\frac{\bar{\theta}\beta}{d} + \frac{(1-e)\gamma\bar{\theta}^2}{d(1-\phi)}(1 - \phi + k^2t^2\eta)\right] - \beta\frac{\beta + \gamma(1-e)\bar{\theta}}{d} - \frac{\gamma e}{2}. \qquad (16)
$$

Differentiating (16) with respect to β and γ leads to the first-order conditions

$$
\frac{[1 - 2(1-e)\gamma]\bar{\theta} - 2\beta}{d} = 0 \quad \text{and} \qquad (17)
$$

$$
-\frac{e}{2} + \frac{(1-e)\bar{\theta}\left\{[1 - 2(1-e)\gamma](1 - \phi + k^2t^2\eta)\bar{\theta} - 2(1-\phi)\beta\right\}}{d(1-\phi)} = 0 \qquad (18)
$$

The solution of (17) and (18) for β and γ is stated in the proposition. Since

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

$$
\frac{\partial^2 \mathrm{E}(\pi)}{\partial \beta^2}\frac{\partial^2 \mathrm{E}(\pi)}{\partial \gamma^2} - \left(\frac{\partial^2 \mathrm{E}(\pi)}{\partial \beta \partial \gamma}\right)^2 = \frac{4(1-e)^2\eta k^2t^2\bar{\theta}^2}{d^2(1-\phi)},
$$

$\mathrm{E}(\pi)$ is strictly concave in β and γ and so the solution of (17) and (18) 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 for γ^* in (8) must be positive; otherwise the optimum is a corner solution with $\gamma = 0$. It is straightforward to verify that this condition is equivalent to assumption (A1).

2. The agent's effort must be positive for each realization of $\mathbf{s}$ for (7) to be valid. There are two candidates for the lowest value of $a_1(\mathbf{s})$, namely $a_1(-1,1)$ and $a_1(-1,-1)$. Using (7) and (3), the sign of the difference $a_1(-1,-1) - a_1(-1,1)$ can be computed as

$$\phi(1-k^2\rho^2) - \rho(1-k^2). \tag{19}$$

If (19) is negative, the smallest value of a_1 is given by

$$a_1(-1,-1) = \frac{(1+k^2\rho)\beta + (1-e)\gamma\bar{\theta}[(1-k)(1-k\rho) + k(1-t)(1+\rho)]}{2d(1+k^2\rho)},$$

which is always positive. If (19) is positive, the smallest value of a_1 is given by $a_1(-1,1)$. Using (7) and substituting for $\widehat{\boldsymbol{\theta}}$ from (3) and for β and γ from (8), one can verify that the condition for $a_1(-1,1)$ to be positive is given by (A2).

3. Finally, $\theta_1 a_1 + \theta_2 a_2$ must not exceed 1. The largest value of $\theta_1 a_1 + \theta_2 a_2$ is attained when $\boldsymbol{\tau} = \mathbf{s} = (1,1)$, and using (7) and (3) leads to

$$Y((1,1),(1,1)) = \frac{(1+t)\bar{\theta}}{d}\left\{\beta + (1-e)\frac{[1+k^2\rho+(1+\rho)kt]\bar{\theta}\gamma}{1+k^2\rho}\right\},$$

which is less than 1 as long as

$$d(1+k^2\rho) > (1+t)\bar{\theta}[(1+k^2\rho)\beta + (1-e)(1+k^2\rho+(1+\rho)kt)\bar{\theta}\gamma). \tag{20}$$

Substituting for β and γ in (20) from (8), and using (2), then leads to condition (A3). ∎

Proof of Proposition 2: Part (a) follows from inspection of (8) (recall that η does not depend on $\bar{\theta}$). For part (b), we first need to derive an expression for the agent's expected wage. Since $w = Y - \pi$, the expected wage is the difference between (16) and (15), which equals

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

Substituting for β and γ from Proposition 1 and simplifying leads to

$$E(w) = \frac{4(1-e)^2k^2t^2\eta(1-\phi+k^2t^2\eta)\bar{\theta}^4 - (1-\phi)^2d^2e^2}{16(1-\phi)(1-e)^2dk^2t^2\eta\bar{\theta}^2} \tag{21}$$

It is then straightforward to verify that the derivative of (21) with respect to $\bar{\theta}$ is positive.

∎

Proof of Proposition 3: From inspection of (8), part (a) holds if and only if $k^2\eta$ is increasing in k. That is the case, since substituting from (2) one can compute

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

For part (b), it is straightforward to verify that (21) is increasing in $k^2\eta$. The previous argument then implies the result. ■

Proof of Proposition 4: Part (b) follows from inspection of (8). For part (a), it follows from inspection of (8) that β^* is increasing in e. The derivative of γ^* with respect to e is

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

which is negative if and only if

$$d(1+e)(1-\phi^2) - 4(1-e)k^2t^2\eta\bar{\theta}^2 > 0. \tag{22}$$

Next, use (1) and (15) and substitute for β and γ from (8) to obtain the (ex-ante) expected value of y:

$$E(y) = \frac{2(1-e)(1-\phi+k^2t^2\eta) - (1-\phi)de}{4d(1-\phi)} \tag{23}$$

The derivative of $E(y)$ with respect to e has the same sign as

$$d(1-\phi) - 2(1-\phi+k^2t^2\eta)\bar{\theta}^2 \tag{24}$$

The difference (22) minus (24) is

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

which means that γ^* is decreasing in e whenever $E(y)$ is increasing in e, proving part (b). For part (c), it is straightforward to verify that (21) is decreasing in e and increasing in t. ■

Proof of Proposition 5: Inspection of (8) 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 (9)), is positive.

The l.h.s. of (9) equals $\phi(2-k^2)>0$ for $\rho=0$, equals $-2(1-\phi)(1-k^2)^2<0$ for $\rho=1$, and its derivative of the l.h.s. of (9) 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 existence of critical value $\widetilde{\rho}\in(0,1)$ such that (9) holds for all $\rho\leq\widetilde{\rho}$.

Next, $a_1(1,-1)>a_1(1,1)$ holds if and only if (19) is positive. Subtracting $2(1-k^2)$ times (19) from (9) yields $k^2\phi(1-\rho^2)(1-k^4\rho^2)$, which is positive. This means that whenever $a_1(1,-1)>a_1(1,1)$, (9) also holds, implying that γ^* is decreasing in ρ.

For part (b), it is straightforward to show that (21) is increasing in η, and then the result follows from the arguments above. ■

Proof of Proposition 7: Part (a): Inspection of (8) 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

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

Part (b): The derivative of (21) with respect to ϕ is

$$\frac{[(1-\phi)^2d^2e^2+4(1-e)^2k^4t^4\eta^2\bar{\theta}^4][\eta+(1-\phi)\frac{\partial\eta}{\partial\phi}]}{16(1-e)^2(1-\phi)^2k^2t^2\eta^2d\bar{\theta}^2} \tag{25}$$

The denominator and the first term in the numerator of (25) are both positive; (25) therefore has the same sign as

$$\eta+(1-\phi)\frac{\partial\eta}{\partial\phi}. \tag{26}$$

Evaluating $\frac{\partial\eta}{\partial\phi}$ and using the definition of η from (2), (26) simplifies to $(1-\rho)^2/(1-k^2\rho)>0$. ■

Proof of Proposition 8: The agent's individually optimal effort vector is given by (7) with β set to 0. Consequently, the principal's profit is also given by (16) with β set to 0:

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

Differentiation with respect to γ leads to the stated expression. The proof that (A1')-(A3') are the necessary and sufficient conditions for an interior solution proceeds in the same fashion as the proof of Proposition 1. ■

Proof of Proposition 9: By inspection, γ^* as stated in (10) is increasing in $\bar{\theta}$ and t and decreasing in k. It is increasing in k if and only if ηk^2 is, which was already shown in Proposition 3. γ^* is decreasing in e if

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

cf. (22). Proceeding in the same way as in Proposition 4, it can be shown that γ^* is decreasing in e whenever $\partial E(y)/\partial e \geq 0$. γ^* is decreasing in ρ if η is decreasing in ρ, the condition for which was derived in Proposition 5. Finally, γ^* is increasing in ϕ if and only if $(1-\phi)/(1-\phi+\eta k^2t^2)$ is increasing in ϕ, which is easily shown to be true. ■

Proof of Proposition 10: Proceeding as in Section 4.1, the agent's optimal effort can be determined as

$$a_i(\widehat{\theta}_i) = \frac{\beta + (1-e)\gamma\widehat{\theta}_i}{2d}, \tag{27}$$

where, by setting $\rho = 0$, (3) implies $\widehat{\theta}_i = (1+kts_i)\bar{\theta}$. The principal's expected profit as a function of β and γ is then given by

$$[1-\gamma(1-e)]\left[\beta + (1-e)\gamma\bar{\theta}(1+k^2t^2)\right]\frac{\bar{\theta}}{2d} - \beta[\beta+(1-e)\gamma]\frac{\bar{\theta}}{2d} - \frac{\gamma e}{2}, \tag{28}$$

which corresponds to (16) when $\phi = \rho = 0$ (such that $\eta = 1$) and when $E(a_1 + a_2)$ and $E(Y)$ are each half of the values given in (13) and (15) because the effort levels are only half as large (cf. (7) and (27)). Differentiation of (28) with respect to β and γ then leads to the expressions in (11).

It remains to be shown that Assumptions (A1") and (A3") are necessary and sufficient conditions for an interior solution. First, it is straightforward to show that (A1") is necessary and sufficient for γ^* to be positive. Next, a_1 attains its lowest value when $s_1 = -1$, in which case $a_1 = [2\beta + (1-e)(1-kt)\gamma\bar{\theta}]/(4d)$. Substituting for β and γ from (11), this expression equals

$$\frac{(1+2kt)de + (1-e)(1+3kt-2k^2t^2)kt\bar{\theta}^2}{2d(1-e)(8k^2t^2-1)\bar{\theta}} > 0,$$

implying that $a_i \geq 0$ always holds, i.e. that no counterpart of (A2) is needed in the two-agent case. Finally, the largest expected value of Y is given by

$$Y((1,1),(1,1)) = \frac{(1+t)\bar{\theta}}{2d}[2\beta + (1-e)(1+kt)\bar{\theta}\gamma],$$

which is less than 1 as long as

$$2d > (1+t)\bar{\theta}[2\beta + (1-e)(1+kt)\gamma\bar{\theta}). \tag{29}$$

Substituting for β and γ in (29) from (11) then leads to condition (A3"). ■

Proof of Proposition 11: With two exceptions, all results are obtained in straightforward fashion by differentiating β^* and γ^* in (11) with respect to the exogenous parameters. The two exceptions are the following. First, the derivative of β^* with respect to $\bar{\theta}$ has the same sign as $(1-e)(2-k^2t^2)\bar{\theta}^2 - 3de$. Subtracting $\frac{3}{2}$ times the l.h.s. of (A1") from this expression leads to $\frac{1}{2}(1-e)(5-4k^2t^2)\bar{\theta}^2 > 0$, which implies that when (A1") holds, β^* must be increasing in $\bar{\theta}$. Second, the derivative of γ^* with respect to e is negative if

$$2d(1+e) - (1-e)(2k^2t^2-1)\bar{\theta}^2 > 0. \tag{30}$$

(30) is decreasing in k^2t^2, so a sufficient condition for (30) to hold is that it holds when $k = t = 1$, i.e. when

$$4d(1+e) - 2(1-e)\bar{\theta}^2 > 0. \tag{31}$$

Now, when $k = t = 1$, (A3") equals

$$d(7-5e) - 8(1-e)\bar{\theta}^2 \tag{32}$$

The difference $4*(31) - (32)$ equals $3d(3+7e) > 0$, which means that under (A3"), (31) must also hold, completing the proof that γ^* is decreasing in e for all values of e. ■

References

Aggarwal, Rajesh K., and Andrew Samwick (1999): The Other Side of the Trade-off: The Impact of Risk on Executive Compensation Source. *Journal of Political Economy* **107** (1), 65-105

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

Albuquerque, Ana (2003): Who Are Your Peers? A Study of Relative Performance Evaluation. Working paper, University of Rochester

Athey, Susan; and John Roberts (2001): Organizational Design: Decision Rights and Incentive Contracts. *American Economic Review Papers and Proceedings* **91** (2), 200-205

Baiman, Stanley; David F. Larcker, and Madhav V. Rajan (1995): Organizational Design for Business Units Source. *Journal of Accounting Research* **33** (2), 205-29

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

Baker, George P. (2000): The Use of Performance Measures in Incentive Contracting. *American Economic Review* **90**

Baker, George P. (2002): Distortion and Risk in Optimal Incentive Contracts. *Journal of Human Resources* **37** (4), 728-51

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

Baker, George P., and Brian J. Hall (2002): CEO Incentives and Firm Size. Harvard NOM Working Paper No. 02-37. *Journal of Labor Economics*, forthcoming

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

Bertrand, Marianne; and Sendhil Mullainathan (2001): Are CEOs Rewarded for Luck? The Ones without Principles Are. *Quarterly Journal of Economics* **116** (3), 901-32

Bushman, Robert M., Raffi J. Indjejikian, and Abbie Smith (1995): Aggregate Performance Measures in Business Unit Manager Compensation: The Role of Intrafirm Interdependencies. *Journal of Accounting Research* **33** (Suppl.), 101-28

Bushman, Robert M., Raffi J. Indjejikian, and Abbie Smith (1995): CEO Compensation: The Role of Individual Performance Evaluation. *Journal of Accounting and Economics* **21** (2), 161-93

Bushman, Robert M., and Abbie Smith (2001): Financial Accounting Information and Corporate Governance. *Journal of Accounting and Economics* **32** (1-3), 237-333

Datar, Srikant; Susan Kulp, and Richard A. Lambert (2001): Balancing Performance Measures. *Journal of Accounting Research* **39** (1), 76-92

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* **69** (4), 811-38

Drago, Robert; Garvey, Gerald T. (1998): Incentives for Helping on the Job: Theory and Evidence. *Journal of Labor Economics* **16** (1)

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

Garvey, Gerald T., and Todd T. Milbourn (2003): Asymmetric Benchmarking in Compensation: Executives are Rewarded for Good Luck but not Penalized for Bad, Working paper, Washington University

Gibbs, Michael, Kenneth A. Merchant, Wim A. Van der Stede, and Mark E. Vargus (2003a): Determinants and Effects of Subjectivity in Incentives. *Accounting Review*, forthcoming

Gibbs, Michael, Kenneth A. Merchant, Wim A. Van der Stede, and Mark E. Vargus (2003b): Performance Measure Properties & Incentives. Working paper, University of Chicago

Gibbons, Robert (1998): Incentives in Organizations. *Journal of Economic Perspectives* **12** (4), 115-32

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*

Hermalin, Benjamin (2003): Trends in Corporate Governance. Working paper, UC Berkeley

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

Holmström, Bengt (1982): Moral Hazard in Teams. *Bell Journal of Economics* **13** (2), 324-340

Holmström, Bengt (1992): Contracts and the Market for Executives: Comment. In: Lars Wein and Hans Wijkander (eds.), *Contract Economics*, New York: Blackwell Publishers.

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

Ittner, Christopher D., and David F. Larcker (1998): Innovations in Performance Measurement: Trends and Research Implications, *Journal of Management Accounting Research* **10**, 205-38

Ittner, Christopher D., David F. Larcker, and Madhav V. Rajan (1997): Performance Measures in Annual Bonus Contracts. *Accounting Review* **72** (2), 231-55

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* **26**, 1-14

Keating, Scott A. (1997): Determinants of Divisional Performance Evaluation Practices. *Journal of Accounting and Economics* **24**, 243-273

Kole, Stacey R., and Kenneth M. Lehn (1999): Deregulation and the Adaptation of Governance Structure: The Case of the U.S. Airline Industry. *Journal of Financial Economics* **52** (1), 79-117

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

Lambert, Richard A. (2001): Contracting Theory and Accounting. *Journal of Accounting and Economics* **32** (1-3), 3-87

Lambert, Richard A., David F. Larcker (1987): An Analysis of the Use of Accounting and Market Measures of Performance in Executive Compensation Contracts. *Journal of Accounting Research* **25** (Suppl.), 85-125

Lazear, Edward (1995): *Personnel economics*, MIT Press

Marino, Anthony M., and John Matsusaka (2003): Decision Process, Agency Problems and Information: An Economic Analysis of Budget Procedures. *Review of Financial Studies*, forthcoming

Murphy, Kevin J. (1999): Executive Compensation. In: Ashenfelter and Card (eds.), *Handbook of Labor Economics, Vol. 3B*

Murphy, Kevin J., and Paul Oyer (2003): Discretion in Executive Incentive Contracts. Working Paper, USC

Murphy, Kevin J., and Jan Zabojnik (2003): Managerial Capital and the Market for CEOs. Working Paper, USC

Ortega, Jaime (2003): Employee Discretion at Work and the Scope of Performance Pay Contracts: Testing Multitask Principal-Agent Theory. Working paper

Prendergast, Canice (1999): The Provision of Incentives in Firms. *Journal of Economic Literature* **37** (1), 7-63

Prendergast, Canice (2002): The Tenuous Tradeoff Between Risk and Incentives, *Journal of Political Economy* **110** (5), 1071-1102

Raith, Michael (2003): Competition, Risk and Managerial Incentives, *American Economic Review* **93** (4), 1425-1436

Sasser, W. Earl, Jr.; and Samuel H. Pettway (1974): Case of Big Mac's Pay Plans, *Harvard Business Review*, July

Shi, Lan (2003): Respondable Risk and Incentives: Evidence from Executive Compensation. Working paper, University of Chicago

Sloan, Richard (1993): Accounting Earnings and Top Executive Compensation. *Journal of Accounting and Economics* **16**, 55-100

Smith, Clifford W. Jr., and Ross L. Watts (1992): The Investment Opportunity Set and Corporate Financing, Dividend, and Compensation Policies. *Journal of Financial Economics* **32** (3), 263-92

Stein, Jeremy (2002): Information Production and Capital Allocation: Decentralized versus Hierarchical Firms, *Journal of Finance* v57, n5, 1891-1921

Stern Stewart (1999): ABC, the Balanced Scorecard, and EVA: Distinguishing the Means from the End. *EVAluation* **1** (2) (Stern Stewart Europe Ltd.)

Van Ophem, Hans, Joop Hartog, and Wim P. Vijverberg (1993): Job Complexity and Wages. *International Economic Review* **34** (4), 853-72

Watts, Ross L. (2003): Conservatism in Accounting. Part I: Explanations and Implications. Working paper, University of Rochester

Zabojnik, Jan (1996): Pay-Performance Sensitivity and Production Uncertainty. *Economics Letters* **53** (3), 291-96