NBER WORKING PAPER SERIES

THE TENUOUS TRADEOFF BETWEEN RISK AND INCENTIVES

Canice Prendergast

Working Paper 7815
http://www.nber.org/papers/w7815

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2000

Thanks to Jim Baron, Wouter Dessein, Luis Garicano, Bob Gibbons, Tom Hubbard, Tano Santos, Michael Raith and participants at the University of Chicago, NBER and the American Compensation Association conference, for helpful comments. This work was supported by the NSF. Any errors are my own. The views expressed herein are those of the author and do not necessarily reflect those of the National Bureau of Economic Research.

The Tenuous Tradeoff Between Risk and Incentives
Canice Prendergast
NBER Working Paper No. 7815
July 2000
JEL No. J3

ABSTRACT

Empirical work testing for a negative tradeoff between risk and incentives, a cornerstone of agency theory, has not had much success. Indeed, the data seem to suggest a positive relationship between measures of uncertainty and incentives, rather than the posited negative tradeoff. I argue that the existing literature fails to account for an important effect of uncertainty on incentives through the allocation of responsibility to employees. When workers operate in certain settings, the activities that they should engage in are well known, and firms are content to assign tasks to workers and monitor their inputs. By contrast, when the situation is more uncertain, firms know less about how workers should be spending their time. As a result, the delegate responsibility to workers but, to constraint heir discretion, base compensation on observed output. Hence, uncertainty and output-based pay are positively related. I argue that parts of the existing empirical literature are better explained through this lens than with the standard model.

Canice Prendergast
Graduate School of Business
The University of Chicago
1101 E. 58th Street
Chicago, IL 60637
and NBER
canice.prendergast@gsb.uchicago.edu

1 Introduction

Much of the empirical and theoretical work on agency issues concerns the tradeoff of risk and incentives. From this perspective, the cost of offering a pay-for-performance contract to a (risk averse) employee is that it imposes risk on his compensation, which causes higher wage costs. Consequently, when choosing higher performance pay, firms trade off the benefits of more effort against higher wage costs. The risk imposed on workers is increasing in the uncertainty of the environment so that the standard test of the tradeoff is to show that incentive pay is lower in more uncertain environments. Unfortunately, empirical research has not shown a convincing relationship between pay-for-performance and observed measures of uncertainty. Indeed, as described below, for a range of occupations the data suggest that observed measures of uncertainty are *positively correlated* with incentive provision. The purpose of this paper is to understand why this is the case, and to offer an alternative theory of how uncertainty affects incentive provision.

In agency models, the uncertainty of the environment typically has one effect: it adds observation error to performance measures. This leads to the negative relationship between uncertainty and incentives. This paper argues that there is another effect of uncertainty on incentive provision which may be more important, namely, the delegation of responsibility to employees. The paper focuses on the distinction between instances where an employer tells his agent what to work on, and situations where the agent is given discretion over the activities that he spends time on. The results of the paper are based on two implications of this choice. First, delegation is more likely when there is greater uncertainty about what the agent should be doing. Second, output-based incentive pay is more likely to be observed in cases when discretion is limited: there is little need to base pay on output when inputs are monitored. So uncertain environments result in the delegation of responsibilities, which in turn generates incentive pay based on output. Thus, uncertainty and incentives are positively correlated.

The idea is best described by the following example. A firm is involved in large scale construction projects around the world, and uses project managers to run those projects. Compare projects being carried out in Canada to those in, say, Armenia. The company has considerable experience in Canada and "knows the ropes" for doing business there. By contrast, the company is very much in the dark when doing projects in Armenia, both because the economic environment is so different from what they are accustomed to and because they have little experience there. If one were to make predictions about compensation using the standard trade-off between risk and incentives, one would expect to see more pay-for-performance for the project manager in Canada than in Armenia, as the manager's performance can be measured more precisely in Canada. This is exactly the opposite of what happens: in fact, the Canada manager is paid a salary (with a small bonus), while the Armenia manager's pay is tied to the profitability of the project. The reasoning is simple. The company believes that as it has a good idea of how business should be done in Canada, it takes control of most decision-making and monitors the manager largely on the basis of his inputs. Headquarters feels that they can make effective decisions in Canada and therefore does not delegate much decision-making power to the manager. By contrast, headquarters has little ability to determine the profitability of the Armenian projects, as this depends on many pieces of information that it does not have. Because they believe themselves to have such a poor handle on the business, they choose an alternative strategy: they (largely) delegate decision-making to the project manager. But they also offer him an output-based contract, as this is the only way to monitor his performance. In effect, output-based pay is used because in uncertain environments, there are no other good measures by which to align incentives. Thus, incentive pay and uncertainty are positively correlated, unlike the standard model of trading off risk and incentives. The modeling sections of the paper are largely devoted to outlining this effect, discussing its implications, and analyzing the robustness of its conclusions.

I begin in Section 2 by describing the existing empirical evidence on the tradeoff of risk and incentives. Evidence on the tradeoff comes from three areas: (i) executives, (ii) sharecroppers, and (iii) franchisees. First, the evidence on executive compensation is about whether risk and incentives are substitutes: some studies find evidence in favor of the tradeoff, others do not. Executives compensation also provides weak evidence on another implication of the tradeoff, namely, relative performance evaluation, where common risk should be filtered out of compensation plans. Second, the evidence on agricultural sharecropping clearly points to a *positive* relationship between measures of risk and the fraction of output received by sharecroppers. Third, firms often decide whether to franchise their retail outlets or to retain them as company-owned. For franchisees, strong output-based contracts are the norm, while in company owned stores, pay-for-performance is more muted. Contrary to the tradeoff of risk and incentives, again there is a clear *positive* relationship between measures of risk and the decision to franchise: thus, pay-for-performance is more common in risky settings.[1]

Most work in agency theory assumes that the tasks carried out by employees are fixed, and

[1]Beyond these systematic studies, the theory also seems a little strained at a more anecdotal level in that the theory suggests that pay for performance would be more likely to be observed in stable industries and time periods, where there is little extraneous risk on workers. While I know of no recent systematic work on the cross-industry nature of pay-for-performance, it appears that much of the use of incentive pay is in volatile industries, such as the use of options in high tech industries and bonuses in the financial sector. Lazear (2000) also makes this point. If the tradeoff of risk and incentives is the primary force determining pay-for-performance, these are hardly the candidates that would be predicted to be most likely to use such risky instruments.

then considers the optimal output-based contracts given that restriction. At one extreme of this spectrum, workers are residual claimants on output, while at the other they are offered few incentives without output-based pay. But realistically, when output-based pay disappears, firms do not simply resign themselves to workers exerting little effort. Instead, they find other means of monitoring, namely by assigning the agent to carry out certain tasks and by direct observation of the agent's inputs. This in itself is not a problem for the existing theory unless the marginal cost of using this combination of directed actions and input monitoring depends on the uncertainty in the environment. But I argue that there is a natural relationship between the effectiveness of directed action and uncertainty. In particular, in stable scenarios, a principal has a good idea of what the agent should be doing, so that by observing efforts, he can be pretty sure that private and social benefits are aligned. However, in less certain environments, the principal may be able to monitor inputs (for example, whether the agent is keeping busy) but is likely to have less idea about what the agent should be spending his time on. In the absence of an effective mechanism for revealing this information, the principal is likely to respond by offering a pay-for-performance contract. In other words, input monitoring will be used in stable settings, but less so in more uncertain environments where workers will be offered more discretion but will have their actions constrained by tying pay to performance.

The reason for this result is simply that the expected marginal return to using output-based contracts to align socially and privately optimal actions is greater in uncertain situations than in stable ones, so that incentive pay is positively related to uncertainty. Obvious though this point may be, it remains unobserved in the literature, which has typically assumed that the marginal return to actions is independent of the underlying riskiness of the environment. I argue that this assumption is not plausible in many settings, and introducing a relationship between uncertainty and the marginal return to agent actions explains why we see more pay-for-performance in rapidly changing industries such as the high-tech sector,[2] while in more stable settings, input-based contracts will be the norm.

The basic insight of the model is described in Section 3, where the optimal tradeoff between input-monitoring and output-monitoring is shown to depend on the underlying uncertainty of the environment. Section 4 considers further applications and extensions of the model. First, I argue that unlike the implications of the existing literature, this paper suggests that pay-for-performance contracts are more likely to be found in complex than in straightforward jobs. The reason for this is that it is more difficult to monitor in complex positions, as the optimal action is hard to pinpoint, compared to jobs where there is little doubt over the right course of action. Hence complexity and

[2]One needs to be a little careful here to distinguish these firms from start-ups, which may offer options and stocks rather than high salaries for liquidity constraints reasons.

incentive pay go hand-in-hand. Second, I consider a series of modeling extensions, where I show that the insights of the model are robust to such extensions as allowing communication between the principal and agent, and to cases where partial delegation of responsibilities may be optimally used. Finally, I argue that the basic results are also likely to hold with other forms of uncertainty.

In Section 5 I show that one critical assumption generates the positive tradeoff between uncertainty and incentives: the availablility of good measures of output. It is well known that when output measures are unreliable, the desirability of output-based contracting falls due to multitasking concerns. This *in itself* is not a problem for the theory; what does matter is how multitasking concerns vary with the uncertainty of the environment. Specifically, if multitasking concerns are greater in uncertain environments, then the results described above can reverse in that the standard negative tradeoff can now be atttained. To see why, consider the example of the construction company in Armenia and Canada. An obvious multitasking problem which can arise with construction projects is that the accounts can be "doctored" to show that a project is doing well when in fact it is a bust. This doctoring is done by the project manager to avoid cancelling an inefficient project from which the manager gains private benefits. If the extent to which accounts can be manipulated does not vary between Canada and Armenia, the qualitative results of the basic model are unchanged: multitasking makes output-based contracting less desirable, but there are no qualitative implications for how uncertainty affects the tradeoff between input and output monitoring. But if it is easier to distort performance measures in more uncertain environments (for example, as accounting methods in Armenia are non-standard), then the results can reverse in that a negative tradeoff is now possible. The reason for this reversal is that there is now a countervailing effect whereby output-based contracts become increasingly distorted in uncertain settings. If the distortion increases sufficiently rapidly with uncertainty, it can swamp the beneficial returns to delegating. The upshot of this is that a positive tradeoff of risk and incentives should be predicted only in situations when good measures of performance are available, such as sharecropping and franchising, rather than for occupations where observed measures are a poor reflection of performance.

The paper is largely motivated by the absence of a negative tradeoff between incentives and observed measures of uncertainty in the empirical literature. I return to the empirical evidence in Section 6, where I consider the implications of this model for observed outcomes. I argue that the observed evidence is better explained by this model than the standard agency model. In standard econometric parlance, a difficulty with the existing empirical work is omitted variable bias. This model argues that uncertainty affects the responsibilities offered to workers, which in turn affects incentives. But responsibilities and discretion are rarely observed by the econometrician, so that omitted variable bias arises. I use existing evidence from the literature on franchising and sharecropping to address this, and to argue that the data appear consistent with the insights of the model.

In the most general incarnation of the standard agency model, such as Grossman and Hart (1983), it is difficult to generate simple agency contracts where one can talk about more or fewer incentives. The benchmark here is not such a general model, but the more commonly used model offered by Holmstrom and Milgrom (1991), where linear contracts are generated by assuming a dynamic model with exponential utility and normally distributed errors. For our purposes, this benchmark has two attractive features. First, optimal contracts are linear in measures of output, and second, the (sole) measure of incentives is decreasing in the (sole) measure of risk for a normal distribution, the variance. Anyone familiar with recent contibutions to agency theory will recognize the by now common optimal piece rate $\frac{1}{1+rC''(e)\sigma^2}$, where r is the degree of absolute risk aversion, $C(e)$ is the cost of effort to the agent and σ^2 the variance of the measurement error on performance. This yields the simple prediction of the piece rate falling with the variance. I construct a model which is in spirit similar to this, though with risk neutral agents, and show how the optimal piece rate increases with the measure of uncertainty, σ^2.

I begin in Section 2 by documenting empirical evidence on the relationship between measures of risk and the incentives provided to workers, and illustrate the paucity of evidence on a negative relationship between these variables. Following this, Section 3 considers the tradeoff between directing the actions of employees (and monitoring inputs) and delegating the choice to the worker (with endogenously optimal output based monitoring). I show that the cost of assigning tasks to the agent is increasing in the uncertainty of the environment so that firms prefer to use output-monitoring if the environment is risky enough. Section 4 considers some implications and extensions of the model, while Section 5 considers the importance of good output measures. I examine the empirical implications of the model in Section 6 and I conclude with a brief discussion.

2 The Evidence

This section provides a backdrop by surveying the existing evidence on how pay-for-performance varies with uncertainty. In each of the studies below, I use the following characterization: a *(-)* implies that there is statistically significant evidence (at the 5% level) of a negative relationship, as predicted by the theory, a *(+)* implies a significant positive relationship, and a *(0)* means that there is no statistically significant relationship. For reasons that will become more clear below, I consider four different classes of occupations where the tradeoff has been tested: (i) executives, (ii) sharecroppers, (iii) franchisees, and (iv) salesforce workers and others.

Table 1 gives the evidence on chief executive officers (and sometimes other executives), providing both the measure used and its results.[3] Here the evidence is inconclusive, with two studies finding a statistically significant negative effect, one a significant positive relationship, and five others find no relationship between risk and incentives.

Table 1: The Tradeoff of Risk and Incentives for Executives.

Authors	Measure of Risk	Result
Lambert and Larker (1987)	Volatility of Returns	-
Aggarwal and Samwick (1998a)	Volatility of Returns	-
Core and Guay (1999)	Idiosynchratic Risk	+
Bushman et al. (1996)	Volatility of Returns	0
Conyon and Murphy (1999)	Volatility of Returns	0
Garen (1994)	Volatility of Returns	0
Ittner et al. (1997) (Full sample)	Correlation of Financial and Accounting Returns	0
Yermack (1995) (Options Only)	Variance of Returns	0

An additional implication of the tradeoff of risk and incentives is relative performance evaluation, where firms filter out common risk from compensation contracts. In the context of executives, this takes the form of rewarding executives on how their firm does relative to some benchmark, such as competetitor or market performance. There have been a number of tests of such relative performance evaluation, as described in Table 2.

[3]In each of these cases, the test carried out addresses how the fraction of the firm held by the executive varies with some measure of uncertainty. The theory addresses how "piece rates" vary with uncertainty, so this is the closest available measure. Several studies present many sets of results, and I provide the most comprehensive model. Some elaboration is needed for two cases here. First, Ittner et al. (1997) offer some results with a significant negative effect and others with no significance. In a simple OLS regression there is a negative relationship, but when endogeneity is controlled for via a structural estimation technique for their full sample, no noise variable is significant at the 5% level. I cite this latter result here. Second, Core and Guay (1999) offer a series of results, mostly on total equity compensation rather than on the share owned by the executive. The results cited here are taken from Table 7 in their paper, which carries out the desired regression with % ownership as the dependent variable, and shows a positive relationship with idiosynchratic risk.

Table 2: Relative Performance Evaluation for Executives.

Authors	Evidence of Relative Performance Evaluation
Gibbons and Murphy (1990)	Yes
Aggarwal and Samwick (1998b)	No
Barro and Barro (1990)	No
Antle and Smith (1986)	No
Janakiraman et al. (1992)	No
Murphy (1999) [Contracts Observed]	27% use RPE

As can be seen, only one of the five studies that directly tests for relative performance evaluation finds evidence in its favor. Furthermore, the most recent evidence on executive contracting by Murphy (1999), who observes contracts, finds that only about a quarter of firms benchmark performance in this way. Overall, Tables 1 and 2 offer weak evidence about the extent to which incentives are traded off against risk. While there may be such a tradeoff, it hardly jumps out in the data.

Yet executives appear to be the occupation with the *strongest* evidence in favor of the tradeoff. To see this, I now consider two other occupations that have attracted some testing of this hypothesis, namely, sharecropping and franchising.

Table 3: The Tradeoff of Risk and Incentives for Sharecroppers.

Authors	Measure of Risk	Result
Rao (1971)	Variance of Profits	+
Allen and Lueck (1992)	Coefficient of Variation in Yield	+
Allen and Lueck (1995)	Coefficient of Variation in Yield [Within Crop]	+

Thus, the evidence on sharecropping is that the fraction of output sharecroppers keep is increasing in the noisiness of the financial returns, the opposite of the outcome suggested by the theory. Equally, consider the relationship between the decision to franchise (with high pay-for-performance) and the decision to keep a store company owned (with less pay-for-performance).[4] If the traditional tradeoff is correct, we should expect to see a negative relationship between franchising and riskiness. However, Table 4, taken from Lafontaine and Slade (1998), suggests the opposite.

[4]Obviously, this raises the question of why asset ownership affects incentives. See Holmstom and Milgrom (1991) on this.

Table 4: Risk and the Decision to Franchise.

Authors	Industry	Result
Martin (1988)	Panel Across Sectors	+
Norton (1988)	Restaurants and Hotels	+
Lafontaine (1992)	Many Sectors (Business Format Franchising)	+
Anderson and Schmittlein (1984)	Electronics Components	0
John and Weitz (1988)	Industrial Firms	0

As Lafontaine and Slade (1998) conclude, "these results suggest a robust pattern that is unsupportive of the standard agency model" (page 10). Tables 3 and 4 should, I believe, be seen in combination as strong evidence against the traditional negative tradeoff. The most generous interpretation of these data is that there is little evidence in its favor (particularly when one factors in the possibility of publication bias). Indeed, there is more evidence of a positive relationship between uncertainty and incentive provision, particularly outside the market for executives. It is this issue I explore in the theoretical sections of the paper.[5] Before doing so, for the sake of completeness I include other tests on the tradeoff for a range of occupations, which show little systematic pattern.

Table 5: Risk and Incentives in Other Industries.

Authors	Industry	Result
Kawasaki and McMillan (1987)	Japanese Subcontractors	-
Leffler and Rucker (1991)	Timber Tracts	+
Mulherin (1986)	Natural Gas Contracts	+
MacLeod and Parent (1999)	Many Sectors	0
Coughlin and Narasimhan (1992)	Salesforce Workers	0
John and Weitz (1989)	Salesforce Workers	0
Hallagan (1978)	California Gold Mining	0

[5]There are at least two reasons why we might find little relationship between observed measures of risk and incentive pay. First, it may be that the empirical measures of risk and pay-for-performance are a poor reflection of the true environment facing employees. This is the standard empirical measurement error problem and I have little to say on about it. This is less likely to be a problem for the literature on executives or franchisees, since we observe the contracts which they receive and can relatively easily compute the relevant measures of risk. For studies that address employees on implicit contracts, as in Brown (1990) or MacLeod and Parent (1999), this problem may be more severe. Second, it could be that our theories are missing something important about the relationship between the desire to induce individuals to exert effort and the riskiness of the environments in which they find themselves. It is this second point that I address here.

3 Delegation and the Choice between Monitoring Inputs and Outputs

Firms do not choose compensation plans independent of other strategic decisions. In this section, I consider how the decision to delegate decision-making power affects contracts offered to workers and how that decision depends on the uncertainty of the economic environment. Specifically, firms delegate decision-making power more in uncertain environments but offer output-based contracts in order to constrain the possibility that they use their discretion in harmful ways. By contrast, in more certain environments, firms assign tasks to workers and find it more profitable to monitor actions directly.

This section formalizes this in the following way. A principal hires an agent to exert effort on *one* of n possible tasks. For the action on which the agent is employed i, he chooses an effort level to exert, e_i, where i refers to activity i. Output from exerting effort on task i depends on the effort level and on a random variable ρ_i. Specifically, the performance of the firm is given by $y_i = \rho_i + e_i$. The cost of effort on activity i is $C(e_i)$ which has the following standard properties: $C'(e_i) > 0, C''(e_i) > 0$, and $C'(0) = 0$. All individuals are risk neutral and throughout the paper, the reservation utility of the agent is normalized to zero.[6] The distribution of the n random variables ρ_i is given by Φ_i. The distributions differ only in their means: they are distributed according to a common distribution with mean $\overline{\rho}_i$ and variance σ^2. In what follows, I vary σ^2 where an increase in σ^2 identifies a more uncertain environment. This is the measure of risk considered throughout the paper.

The premise of this section is that agents often have information not available to the principal.[7] I assume that the agent knows the true values of ρ_i for all i, while the principal only knows the distribution of the ρ_i. I also assume that the technology is such that the agent can carry out only a single activity, i.e., the fixed cost of engaging in two activities is too large to make it worthwhile to get involved in more than one.

The principal can potentially collect two pieces of information with which to reward the agent. First, she can observe the efforts exerted by the agent, e_i at a monitoring cost m_e. Second, she can collect information on output produced by the agent. This costs m_y to collect. Throughout

[6] I model the agents as risk neutral in order to ignore the standard tradeoff. It is important to note here that I am not claiming that we should necessarily see a positive relationship between observed measures of risk and incentive provision based on the insights provided here. Instead, the claim is that there are plausible influences that can cause a positive relationship, and that there is no necessary reason why the clear and simple logic characterized by the traditional negative tradeoff of risk and incentives should be reflected in the data.

[7] For instance, an agent may know the right customers to focus on, which subordinate to let go, or more generally the right strategies that should be followed.

this section, I assume that $m_y > m_e$. The monitoring costs of output are a metaphor to reflect any costs to introducing a pay-for-performance plan, such as multitasking concerns. (See Section 4 for details.)

If the agent is indifferent about which activity he carries out, the problem has the following trivial solution, assuming that it is at least profitable to monitor inputs. First, as $m_e < m_y$, the principal monitors inputs and offers a contract to the worker $w(e_i) = C(e_i)$ for all i. In other words, simply offer the worker a contract that pays him his costs of effort. As the agent knows the true value of ρ_i, he will choose the activity that yields the highest value of output among the n realized and will exert optimal effort on that activity. This optimal level of effort on that activity is given by e^* where $C'(e^*) = 1$ (the i subscript is dropped for simplicity). This yields the first best allocation of effort and activity selection and dominates any output-based contract, as $m_e < m_y$.

But I assume that the agents have personal preferences over which activity they enjoy most, where the activities have a personal benefit given by B_i. To keep matters simple I assume that these benefits are small enough so that they do not affect the calculation of the optimal activity to carry out.[8] The principal can, of course, extract these expected benefits through the salary paid to the agent.[9] As in Aghion and Tirole (1997) and Prendergast and Topel (1996), I also assume that the principal does not know the preferences of the agent, i.e, which activities have which benefits. The agent knows his private benefits of the various activities. As a concrete example, I consider the case where the agents have personal preferences such that they are indifferent to $n-1$ of the activities ($B_i = 0$) but gain a small benefit $B > 0$ from the last one, where the principal has no idea of the identity of i. The principal believes the distribution over the preferred activity to be uniform. I assume that there is no correlation between B_i and ρ_i, though see Section 4 for a relaxation of this assumption.

Assigned Actions and Input-Based Contracts The principal has two functions: (i) to assign an allowable set of tasks that the worker can carry out, and (ii) to decide how to reward the agent for the allowed set of tasks. First, consider a contract which rewards solely on effort (an input-based contract), and leaves the choice of tasks entirely to the agent. If the principal offers the contract $w(e_i) = C(e_i)$ for all i, the agent will simply carry out the activity that he enjoys most, as he gets rents of B from that activity. This will in expectation yield surplus (and hence benefits to the

[8] If the benefits were large, the principal could allow the agent to carry out an activity even when they are output-dominated by another, and charge the agent for carrying out that activity.

[9] I ignore the agent's individual rationality constraint here, by assuming that the worker signs a contract before observing ρ_i. This allows up-front transfers such that the usual mechanism design issues with an informed agent can be ignored: instead, the principal's objective is to maximize surplus.

principal) of

$$\frac{\sum_{i=1}^{n}\overline{\rho_i}}{n}+1-C(e^*)+B-m_e,$$

where $1 = C'(e^*)$. If $\rho_i \neq \rho_j$, for some i and j and B is small, this is dominated by a strategy where the principal restricts the allowed activities, and offers $w(e_i) = C(e_i)$ only for those activities. (If the agent carries out any other activities, he is simply penalized a sufficiently large amount of money that he will never do that activity.) Let activity k be the task with highest mean $\overline{\rho}_k$. If this is unique,[10] the profits for the firm from this strategy are

$$\overline{\rho}_k+1-C(e^*)+\frac{B}{n}-m_e.$$

As the private benefits are assumed small, this dominates delegating the activity choice to the agent if $\rho_i \neq \rho_j$, for some i and j. If the firm could monitor only inputs, this would be the optimal solution.[11] Note also that with assigned tasks, there is no value to offering an output-based contract, as $m_e < m_y$.

Delegated Actions and Output-Based Contracts Now consider another option, to delegate the choice of action to the agent. This can be optimal only if the agent is paid on output; otherwise, the agent chooses the action with highest private benefits. If output contracting is used, the optimal piece rate is 1 so that the agent chooses the optimal level of effort, e^*, *and chooses the correct activity* j, i.e., the one that maximizes $y_i - C(e_i) + B_i$. In other words, the purpose of offering pay-for-performance here is not simply to induce effort, as an input-based contract could do this, but to induce the agent to carry out the right kinds of efforts. (This, of course, is nothing more than a relabelling of what effort means but here the marginal return to inducing this kind of effort varies with the riskiness of the environment.)

In order to determine the return to offering output-based contracts, consider the distribution of the first order statistic of the realization of the ρ_i. The reason this is necessary is that if the

[10]In the case that m of the n observations are tied with the highest mean $\overline{\rho}_k$, the optimal input contract allows the agent to work on any of these m tasks, and the profits from this strategy are $\overline{\rho}_k + 1 - C(e^*) + \frac{mB}{n} - m_e$. In the case where all ρ_i have identical means, the agent is allowed to work on any activity.

[11]I have restricted attention to the case where $w(e_i) = C(e_i)$. With input monitoring, there is no better contract, though there are obviously others which can replicate the efficient outcome. In another possible mechanism, the principal auctions off the right to carry out various activities, where for example the principal could offer a price at which the agent could carry out activity i. A natural case would be to offer the agent the opportunity to carry out activity k at no price but a positive price of B to carry out any other activity. But this cannot do any better than a contract that simply mandates the agent to carry out a task, as the agent is ex ante indifferent to all ρ_i in an input monitoring plan.

agent is offered an output-based contract, he will choose the highest realization of ρ_i among the n. By contrast, if the agent is offered an input-based contract, he is assigned to activity k, which has mean $\overline{\rho}_k$. What matters then is difference between the distribution of the first order statistic ρ'_1 from the n realisations relative to $\overline{\rho}_k$. The idea here is that this difference is likely to be increasing in the variance of the environment. In other words, when the variance of the distribution of the ρ_i is large, the value of sampling the first order statistic is larger than when it is small. If the variance is very small, the principal knows that if he simply assigns a task to the agent, the expected marginal cost of being mistaken is likely to be small. As a result, there is little cost to input monitoring. On the other hand, when σ is large, the firm will likely use output based monitoring, as there is little certainty about the right kinds of activities to engage in.

I illustrate this with two cases where I can get simple closed form solutions, the normal distribution and the uniform distribution. There is one important difference between these two examples, namely between the case where all actions look *ex ante* identical (Example 1), and the case where some actions look better than others from an *ex ante* perspective (Example 2).

Example 1: The Normal Distribution Assume that $n = 2$, and that both random variables $\rho_i \sim \mathcal{N}(0, \sigma^2)$. Therefore, all activities look identical to the principal. If input monitoring is used, the principal allows the agent to work on any activity and the expected surplus from input monitoring is $1 - C(e^*) - m_e + B$. With output monitoring, the agent chooses the activity that maximizes surplus, as a piece rate of 1 is optimal. As B is small, this implies that the agent chooses the activity with the highest ρ_i. The expected value of the first order statistic of the two random variables $E[\rho'_{1\{2\}}]$ is

$$E[\rho'_{1\{2\}}] = \frac{\sigma}{\sqrt{\pi}}.$$

Note that the first order statistic is increasing in σ^2 and that the profits from output monitoring are $\frac{\sigma}{\sqrt{\pi}} + 1 - C(e^*)) + \frac{B}{n} - m_y$, also increasing in σ^2. Output contracting is therefore preferred if and only if

$$\frac{\sigma}{\sqrt{\pi}} \geq m_y - m_e + \frac{B(n-1)}{n} \tag{1}$$

(subject to the right hand side being non-negative, as is necessarily the case when $m_y > m_e$). The agent is more likely to be offered an output based contract in riskier environments, as the critical level of variance, σ^{2*}, above which output monitoring is used, is given by

$$\sigma^{2*} = (m_y - m_e + \frac{B(n-1)}{n})^2 \pi. \tag{2}$$

This trivial example provides the intuition for the paper's main results. When the environment is more uncertain, the cost of assigning the agent to carry out a particular action is high, as there is likely to be another with significantly better return. As a result, the firm optimally chooses to delegate choice of action in sufficiently uncertain settings, but constrains the agent's choice in that setting by basing pay on output. Thus, incentive pay goes hand-in-hand with uncertainty.

A More General Result An important part of Example 1 is that the returns to all the actions are drawn from the same distribution. For this case, a more general result arises. Assume that all the random variable are drawn from the same distribution and that $y_i = \rho_i + e_i$. In that case, there is a single critical value of σ above which output contracts are optimal and below which assigned actions with input contracts are optimal. This arises from the fact that for a distribution ρ_i with common mean $\overline{\rho}$ and variance σ^2, the expected value of the first order statistic from n draws of a p.d.f. of the form $\sigma^{-1} g\{\frac{(\rho_i - \overline{\rho})}{\sigma}\}$ can be characterized as

$$E[\rho'_{1\{n\}}] = \overline{\rho} + \sigma H_n,$$

where H_n is independent of σ^2 and $\overline{\rho}$, but depends on n. See Cox and Hinkley (1992) for details. Thus, as in Example 1, the value of the first order statistic is linearly increasing in σ, while the return to input contracts is not. Therefore, one can easily generalize the insights from Example 1 to other distributions.

A difficulty in proving general results arises for the case where the random variables are drawn from distributions with different means. Problems arise in finding closed form solutions to carry out comparative statics. For this case, I consider the simplest distribution where closed form solutions are possible, namely, the uniform distribution, where we find results similar to those above.

Example 2: The Uniform Distribution Assume that $\rho_i \sim \mathcal{U}[-x + \overline{\rho}_i, x + \overline{\rho}_i]$ and that $n = 2$. By renormalization, let the distributions be $\rho_1 \sim \mathcal{U}[-x, x]$ and $\rho_2 \sim \mathcal{U}[-x + \Delta, x + \Delta]$, where $\Delta = \overline{\rho}_2 - \overline{\rho}_1$. Therefore, the two activities are uniformly distributed with common variance $\frac{x^2}{3}$, but where activity 2 has a mean which is Δ higher than activity 1. With input monitoring, the agent will be assigned to work on activity 2 and paid his costs of effort. This has expected return $\Delta + 1 - C(e^*) - m_e + \frac{B}{2}$. If offered an output-based compensation plan, the agent chooses the highest value of ρ_i, and expected profits are $E[\rho'_{1\{2\}}] + 1 - C(e^*) + \frac{B}{2} - m_y$. Then similar calculations to those in Example 1 yield

$$E[\rho'_{1\{2\}}] - \Delta = \begin{cases} \frac{1}{6}(1 - \frac{\Delta}{2x})^2(2x - \Delta) & \text{if } \Delta < 2x, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Then delegation is preferred if

$$max\{0, \frac{1}{6}(1 - \frac{\Delta}{2x})^2(2x - \Delta)\} \geq m_y - m_e. \tag{4}$$

But $max\{0, \frac{1}{6}(1 - \frac{\Delta}{2x})^2(2x - \Delta)\}$ is non-decreasing in x (and non-increasing in Δ) and the variance of the uniform distribution, $\frac{x^2}{3}$, is increasing in x. Therefore, as the variance increases, so also (weakly) does the return to using an output-based contract to induce the agent to choose the activity correctly. Thus, once again, this example points to the positive correlation between risk and incentive pay based on outputs.[12]

To summarize, consider a more concrete example, the franchise decision. There is more use of pay-for-performance for franchisees than managers of company-owned stores. Furthermore, we saw from Table 4 that franchises are more common in uncertain environments than in certain settings. The interpretation that I place on this is that in uncertain environments, headquarters has less idea of the kinds of products that should be offered, their prices, the number of employees to hire, and so on, than in more certain situations. As a result, they respond by offering output-based contracts, which are less necessary than when the headquarters knows with more certainty what should be done.

At a more general level, this section raises what I feel is an aspect of the agency literature that is often overlooked, namely, that uncertainty is likely to affect both the compensation and optimal distribution of actions of employees. Typically in agency theory we treat uncertainty in the economic environment as synonymous with measurement error: yet, as illustrated here, uncertainty has other effects on the employment relationship which confound the usual negative tradeoff between risk and incentives.

[12]It is important to note that I am not claiming that accuracy of monitoring is irrelevant to contracting problems. In fact, one can easily reinterpret these results to say that in more certain environments some measure of total incentives is at least as high as in less certain environments; all that differs is that monitoring is on inputs in the certain environments, and outputs in the less certain environments. But note that empirical researchers never see inputs, so that the objective of this paper has been to understand why empirical work which traces the relationship between outputs and uncertainty is unlikely to find a negative relationship.

4 Extensions

In this section, I consider some implications of the model, and also show that the basic insights are robust to other modeling assumptions.

4.1 Complexity

Recent contributions to agency theory focused on multitasking (Holmstrom and Milgrom, 1991) suggest that complex jobs are less likely to use incentive pay. Complex jobs have many dimensions of performance, some of which are poorly observed, so that rewarding on the observed dimensions typically has harmful effects on the unobserved dimensions. Thus, the marginal cost of output contracting in complex positions is high. However, it is also the case that the benefits of contracting on output are likely to be especially high in complex positions. Indeed, a simple parameterization of this model suggests a *positive* relationship between complexity and the likelihood of incentive pay. The reason for this is that it is more difficult to monitor complex positions than those for which it is easy to identify the optimal course of action; as a result, output-based contracts are more desirable in complex positions. A natural measure of complexity in the model is given by n, the number of activities that the individual can carry out. In this model, as the number of activities increases, so also does the desire to induce the agent to choose the right one to work on. As a result, output-based pay is more likely in complex settings.

To see this, consider Example 1 with the normal distribution, but where the agent can now choose among three activities, $n = 3$ rather than $n = 2$ in the basic setup. Let $E[\rho'_{1\{n\}}]$ be the expected value of the first order statistic fron n draws. Then it is the case that

$$E[\rho'_{1\{3\}}] = \frac{3\sigma}{2\sqrt{\pi}} > \frac{\sigma}{\sqrt{\pi}} = E[\rho'_{1\{2\}}]$$

and output monitoring is used when $\sigma^2 < \tilde{\sigma}^{2*}$, where

$$\tilde{\sigma}^{2*} = \frac{4\pi(m_y - m_e + \frac{B(n-1)}{n})^2}{9}. \tag{5}$$

Remember that the critical variance when $n = 2$ is given by

$$\sigma^{2*} = (m_y - m_e + \frac{B}{2})^2\pi > \tilde{\sigma}^{2*}$$

for B low, as is assumed here. Therefore, when the number of activities increases, the range of (variance) parameters where output-based contracts are used also increases. What this suggests is that the returns to offering incentive based pay are greatest for complex jobs, where there is an overall measure of performance.

The logic of this section is simply that for those employees who carry out well-defined jobs, where the activities that keep the person (optimally) busy from one end of the day to the other are known, can easily be rewarded on input-based contracts. However, when the range of activities that the person engages in increases, it becomes harder to identify what the person should be doing and so output-based contracts become necessary.

4.2 Partial Delegation

So far, I have described two options available to the principal: he either assigns a task to the agent, or he allows the agent unrestricted choice over actions. However, in some settings, it may be optimal to allow the agent to choose actions from a limited set, but where the set is not a singleton. This is not an issue when there are only two actions (as in Examples 1 and 2). However, when there are more than two actions, the principal may partially delegate tasks in the following way: (i) for low levels of uncertainty, the principal assigns an action and monitors inputs (as above), (ii) for high levels of uncertainty, the principal allows the agent unrestricted choice over actions but with an output-based contract (again, as above), but (iii) for intermediate levels of uncertainty, he allows the agent to choose between a subset of actions, and uses input monitoring. In effect, the principal excludes some actions, those which are believed to be the least profitable. As a result, allocation of tasks can vary more continuously with our measure of uncertainty (variance) than in the basic model above. Yet it remains the case that output contracting is only used only in sufficiently uncertain settings. I illustrate this by another simple example where $n = 3$ and the distributions are uniform, but with different means.

Example 3: The Uniform Distribution and $n = 3$. Assume that there are three activities, $1, 2$, and 3, where action 1 has the lowest mean payoff and action 3 has the highest. Specifically, let $\rho_i \sim \mathcal{U}[-x + (i-2)\Delta, x + (i-2)\Delta]$, that $y_i = \rho_i + e_i$. Therefore, the three activities are uniformly distributed with common variance $\frac{x^2}{3}$, but where activity 2 has a mean which is Δ higher than activity 1, and activity 3 has a mean which is Δ higher than activity 2.

Three optimal outcomes are possible. First, the agent can be assigned action 3 and paid on inputs. Second, the agent can be offered an output-based compensation plan and given complete discretion over which action to choose. So far, there is nothing conceptually different from the previous sections. However, there is also an intermediate strategy which can be optimal, where the agent is offered an input-based contract, and is then given discretion over actions 2 and 3 only: action 1 is excluded. This partial delegation occurs for intermediate levels of uncertainty.

To see this, consider the value of a strategy where the agent is monitored on inputs but can choose from actions 2 and 3. Let the expected value of the first order statistic *between these two* be given by $E[\rho'_{1\{2\}}]$. This is identical to the problem in Example 2 and is given by

$$E[\rho'_{1\{2\}}] - \Delta = \begin{cases} \frac{1}{6}(1-\frac{\Delta}{2x})^2(2x-\Delta) & \text{if } \Delta < 2x, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

Now consider the value of allowing the agent to also choose the final action 1. Let $E[\rho'_{1\{3\}}]$ refer to the first order statistic chosen from the realization of all three. Straightforward calculations show that

$$E[\rho'_{1\{3\}}] - E[\rho'_{1\{2\}}] = \begin{cases} \frac{1}{6}(1-\frac{\Delta}{x})^2(1-\frac{\Delta}{2x})(x-\Delta)\left(3-\frac{\Delta}{2x}\right) & \text{if } \Delta < x, \\ 0 & \text{otherwise.} \end{cases} \tag{7}$$

The right hand side of (7) is the return to allowing the agent access to the first action over the restricted choice if he always chooses the action which is preferred by the principal. The problem, of course, is that he may not do so with an input contract.

The principal chooses between these three options: delegate all tasks with output contracts, assign to task 3 with an input contract, or allow choice between 2 and 3 with an input contract. Assigning task 3 to the agent yields profits of $\Delta + 1 - C(e_2^*) - m_e + \frac{B}{3}$. Allowing the agent complete discretion with an output-based contract yields profits of $E[\rho'_{1\{3\}}] + 1 - C(e_i^*) + \frac{B}{3} - m_y$. Finally, if the agent is offered an input-based contract and allowed to choose between actions 2 and 3, expected profits are given by[13]

$$\frac{1}{3}\Delta + \frac{1}{3}E[\rho'_{1\{2\}}] + 1 - C(e_i^*) + \frac{2B}{3} - m_e.$$

It is simple to show that partial delegation will then arise if there exists some value of x (the parameter measuring variance) such that

$$\frac{1}{3}\Delta + \frac{1}{3}E[\rho'_{1\{2\}}] + \frac{B}{3} \geq max\{\Delta, E[\rho'_{1\{3\}}] - m_y - m_e\}. \tag{8}$$

[13]This arises as follows. With probablility $\frac{2}{3}$, the agent's preferred action is 2 or 3, and he carries out his preferred action. This has a conditional expected value of ρ_i of $\frac{\Delta}{2}$. With residual probability $\frac{1}{3}$, action 1 is his preferred action, but this is excluded and so he chooses the principal's preferred action, yielding an expected equilibrium value of ρ_i of $E[\rho'_{1\{2\}}]$.

Note that for $\Delta < x$, $E[\rho'_{1\{3\}}] - E[\rho'_{1\{2\}}]$ is increasing in x, as is $E[\rho'_{1\{2\}}]$. This implies that the value of each strategy can be ordered with respect to the variance of the distribution. As a result, (i) for low variance, the agent is assigned action 3, (ii) for intermediate levels, he may be excluded only from action 1 and offered an input-based contract, (iii) for high variance, the agent can choose any action, but is offered an output-based contract.[14] Thus, for some cases delegation is assigned in a more continuous way with respect to variance than in the stark outcomes above, yet it remains the case that output contracting is used in sufficiently uncertain settings.

4.3 Communication

Note that I have restricted attention to two types of contracts: where the worker is delegated the task, and where the principal simply chooses the set of tasks. Yet there is another option, where the agent communicates something to the principal, and the allocation is based on the message sent by the agent. There is one form of communication which can improve the allocation here, in that a Pareto-improving message can be sent: where the agent tells the principal the set of actions which are neither (i) the agent's preferred project nor (ii) the principal's preferred project. In the equilibrium where the principal chooses the action of the agent, excluding these projects benefits both parties.[15]

Yet the basic insight of the model remains with this extension. Consider the case where $n = 3$, and the random variables $\rho_i \sim \mathcal{N}(0, \sigma^2)$. If the agent is monitored on outputs, there is no role for communication and the outcome is as above, with expected return of $\frac{3\sigma}{2\sqrt{\pi}} + 1 - C(e^*) + \frac{B}{3} - m_y$. By contrast, if input contracts are used, there is now a role for communication between the agent and principal. The optimal communication mechanism is where the agent puts forward two projects, one of which is his own preferred project and the other the principal's preferred project; he excludes the project that neither wants.[16] (In the case where the agent's preferred option and the principal's perferred option are the same, he proposes that project plus a dummy project.) The principal then randomly chooses between the two projects that the agent puts forward, and expected surplus

[14] As a trivial example, consider the case where $m_y = \infty$, so that output based contracts are never profitable. Then the optimal actions are (i) for low x, assign action 3 to the agent, and (ii) for higher values of x, allow the agent to choose between actions 2 and 3. In both cases, input monitoring is used.

[15] It is the indifference of the agent over many actions which allows a role for communication here. If the agent is not indifferent between any activities (in terms of private benefits), there can be no role for communication in the absence of output being observed. But if output is observed, then there is no role for communication, as there are no further distortions beyond the monitoring cost.

[16] This mechanism uses the fact that the agent is indifferent between all projects which are not his preferred choice. As a result, given the equilibrium strategy that the principal randomizes between the two put forward by the agent, the agent truly is indifferent between proposing the principal's preferred choice and any other. In order to give this communication game the best chance possible, I assume that he does report the principal's preferred project as the other element of his recommended set.

is given by $\frac{3\sigma}{4\sqrt{\pi}} + 1 - C(e^*) + \frac{B}{2} - m_e$, and so the agent is delegated the task (and offered an output-based contract) if $\sigma^2 < \tilde{\sigma}^{2**}$, where

$$\tilde{\sigma}^{2**} = \frac{16\pi(m_y - m_e + \frac{B}{6})^2}{9}. \tag{9}$$

Without communication, when $n = 3$, output contracts are preferred when the variance parameter exceeds $\tilde{\sigma}^{2*}$ in (5), which is lower than $\tilde{\sigma}^{2**}$ for small B. Thus communication reduces the parameter values for which output-based contracts occur, but retains the qualitative relationship between contracts and uncertainty.

4.4 Symmetric Information and Recontracting

I have assumed that the agent has information on the payoff to various actions, which the principal cannot ascertain. Yet there is another case where this idea becomes relevant, namely in situations of symmetric uncertainty but where there are recontracting costs. For example, consider a new firm, where there is enormous uncertainty about what the firm will be doing in the near future. In these circumstances, it is difficult to write an input-based contract for the simple reason that no one knows what the agent should be doing; to provide such a contract would involve recontracting extremely frequently, as actions which were optimal yesterday not be tomorrow. In such circumstances, firms may simply offer output-based contracts because they do not need to be frequently recontracted; the agent simply orients his actions to whatever increases the bottom line. Again, this idea predicts that when there is uncertainty about what the agent should be doing, output-based contracts should be used, though here the reason is that input-based contracts would have to be more frequently renegotiated, which is likely to be costly.

4.5 Another Form of Uncertainty

I have also considered the case where the private benefits of the agent are sufficiently small that they do not change the efficient allocation. However, consider an alternative where the only uncertainty arises from not knowing the (non-trivial) private benefits of carrying out various actions, rather than the profitability of the actions. Many employers give their employees tasks to carry out, and simply ask that they be completed in some specified period of time, rather than telling them what to do at every moment. Thus, they delegate the choice of when to do the task. The reason for this is that employers realize that workers have preferences which vary from hour to hour, and they allow their employees discretion to perform tasks when they most feel like it. At some points in the day, they may be tired and will do routine things, and save their most arduous tasks for times that they feel most attentive. As with the basic model, this points to the importance of uncertainty, in this case about the agent's costs because if an employer knew the agent's preferences at all points in time, he would simply tell him what to do. It is because of uncertainty about the agent's within-day or within-week costs that the agent is allowed such discretion. But when such discretion is offered, another means of monitoring is necessary, which is typically based on getting the task done (i.e., output). Thus, I feel that the insight about delegation and uncertainty extends to other forms of uncertainty.

4.6 Correlated Preferences

Thus far I have assumed that the preferences of the agent are uncorrelated with those of the principal. This assumption does not affect the qualitative results of the paper, but does change the critical value of uncertainty above which delegation occurs. Intuitively, it results in a greater region for which input monitoring occurs for the simple reason that the agent is now more likely to pick the right activity without the constraint of an output-based contract. To see this, consider a case where the agent's preferred action ($B_i > 0$) is correlated with probability p with the preferred action of the principal (i.e., the first order statistic of the two normally distributed variables). This increases the return to the strategy where you allow the agent to choose the action that he wants, but where inputs are monitored. This yields expected returns of

$$p\frac{\sigma}{\sqrt{\pi}} + 1 - C(e^*) + B - m_e, \tag{10}$$

so that output based contracts are preferred over this strategy only if $\sigma^2 \geq \tilde{\sigma}^{2*}$, where

$$\tilde{\sigma}^{2*} = \frac{(m_y - m_e + \frac{B(n-1)}{n})^2\pi}{(1-p)^2} > \sigma^{2*}. \tag{11}$$

As a result, it remains the case that there exists a critical variance above which output-based contracts are optimal, but that level is higher than in the uncorrelated case.

4.7 The Fixed Costs of Monitoring Output

I have assumed that the costs of monitoring output are fixed, given by m_y. This is meant as a metaphor for any costs that firms incur from monitoring outputs and rewarding agents based on those measures. Indeed, if these were simply monitoring costs, one could construct a random monitoring mechanism along the lines of Becker (1968) to overcome this cost. As a result, I see the fixed costs as representing any deadweight loss from the optimal output-based contract. It is for this reason that I have retained the standard agency issue of inducing effort decisions in the model, as the monitoring costs are meant to reflect any distortions in effort decisions made from contracting.[17]

The standard costs of output-based contracts in the literature are either risk costs imposed on the worker or costs that arise due to multitasking. These are not fixed costs, but costs which increase with the chosen piece rate. The more general insight from this section is that when considering a compensation plan, the costs of forcing an action on the agent must be balanced against the benefits of reducing risk on workers,[18] or reducing multitasking concerns, rather than simply a fixed monitoring expense.

To see this and also the important role for effort in the model, consider the following rent-seeking example provided in Prendergast (1999).[19] Assume the same technology as in the Normal distribution example above ($n = 2, \rho_i \sim \mathcal{N}(0, \sigma^2)$) and that the agent has a cost of effort given by $C(e_i) = \frac{ce_i^2}{2}$. If the principal uses input-based monitoring, he allows the agent to choose the activity (as both appear identical ex ante), observes e_i and rewards via the contract described above. This yields surplus of $\frac{1}{2c} + B - m_e$.

If the principal uses an output-based contract, there is no cost to observing output per se, but the agent can carry out some dysfunctional activity, b_i which increases observed output, $\tilde{y}_i$, but which has no effect on true surplus. For example, the agent could spend resources distorting accounting numbers which have little relation to true profitability. I call this a rent-seeking activity. Specifically, the agent can produce observed output $\tilde{y}_i = y_i + b_i$, where the (deadweight) cost of the rent-seeking activity is given by $K(b_i) = \frac{\kappa b_i^2}{2}$.

This model differs from that above in that there is no monitoring cost to observe output, but there is a cost to using output-based contracts in that they induce rent-seeking. This results in optimal piece rates of less than unity, even with risk neutral agents. Straightforward calculations show that the optimal output piece rate is given by $\beta_1^* = \frac{1}{1+\frac{c}{\kappa}}$ and the resulting surplus is

$$S^* = E[\rho'_{1\{2\}}] + \left(\frac{1}{1+\frac{c}{\kappa}}\right)\left(\frac{1}{c} - \frac{1}{1+\frac{c}{\kappa}}\left(\frac{1}{2c} + \frac{1}{2k}\right)\right) + \frac{B}{2}.$$

[17]Note that in both the input and output monitoring cases above, effort is at the first best level. For this reason, without further elaboration, it is not clear why effort is in the model at all.

[18]It is difficult to obtain a closed form solution for the costs of risk here, even in the commonly considered case of exponential utility functions initially used in Holmstrom and Milgrom (1991), which requires normally distributed errors. The reason for this is that the first order statistic of the normal distribution is not itself normally distributed and so the usual certainty equivalence calculation does not apply.

[19]This is a simple model of influence, introduced by Milgrom and Roberts (1988).

In this case, the analog to m_y in Section 3 is given by $\frac{1}{2c} - \frac{1}{1+\frac{c}{\kappa}}\left(\frac{1}{c} - \frac{1}{1+\frac{c}{\kappa}}\left(\frac{1}{2c} + \frac{1}{2k}\right)\right) > 0$ for $\kappa < \infty$. It follows by substitution that there exists a critical value of σ^2 above which output-based contracts are optimal, so that the logic generalizes to scenarios where the costs of using output-based contracts are not fixed.

5 On What Do These Results Depend?

One critical assumption generates the positive correlation between uncertainty and the use of output-based incentives, namely, that the measure of output be reliable and that the extent of reliability be not correlated with the uncertainty of the environment. Much of the recent work on agency theory starts from the premise that one difficulty with basing contracts on output is that observed measures typically do not reflect true output.[20] It is well known that such concerns reduce the likelihood of output-based pay. However, an additional implication in the context of this model is that if such dysfunctional behavior is more difficult to detect in uncertain environments, the positive correlation between uncertainty and output-based incentives can reverse.

To see this, I make one change to my simplest case, Example 1, where $n = 2$, and both random variables $\rho_i \sim \mathcal{N}(0, \sigma^2)$. Remember that in this case, if input monitoring is used, expected output is $1 - C(e^*) - m_e + B$. Now, however, I follow Baker (1992) and assume that the measure on which the agent can be rewarded is not true output y_i, but a corrupted measure of output $\hat{y}_i$, which is given by

$$\hat{y}_i = \rho_i + \mu e_i, \tag{12}$$

where $\mu \sim \mathcal{N}(0, \sigma_\mu^2)$. (I assume that μ is independent of i and ρ for simplicity.) Also assume that the cost of effort is given by $C(e_i) = \frac{e_i^2}{2}$ and that only the agent knows the true value of μ. The difficulty that this gives rise to is that the agent is rewarded on $\hat{y}_i$, which depends on μ, while surplus does not.

If the principal uses solely the monitoring of inputs, this extension has no effect on our earlier results. However, difficulties arise if only output contracting is used, as follows. Following standard techniques, if the firm chooses to reward the agent on $\hat{y}_i$, the optimal piece rate on output[21] is

[20] For instance, a recent newspaper article (Chicago Tribune, 1998) described the incentive contract offered to a professional basketball player, Tim Hardaway, who was offered an \$850,000 bonus for a 3-1 assist-to-turnover ratio in 1998. The difficulty with such a contract is that the player sometimes had an incentive to carry out actions to affect his assist-to-turnover ratio even in cases when this is not in the team's interest. The player admitted that he occasionally passed up an (efficient) shot in order to get an assist but described this as "a terrible way to play".

[21] I assume, as in Baker (1992), that more complex mechanisms are not possible to extract information on μ.

given by $\beta^* = \frac{1}{1+\sigma_\mu^2}$ and the ensuing surplus from the contract is given by

$$S = \frac{\sigma}{\sqrt{\pi}} + \frac{1}{2(1+\sigma_\mu^2)} + \frac{B}{n} - m_y.$$

If σ_μ^2 and σ^2 are uncorrelated, this extension does not pose any problem to the qualitative results of the paper: the dysfunctional responses imply only that the critical value of uncertainty above which output contracting occurs is higher than without such responses. This is because the multitasking concerns act like a fixed cost (in σ space) of contracting on output. Specifically, the critical value is now given by

$$\sigma^{2***} = (m_y - m_e + \frac{\sigma_\mu^2}{2(1+\sigma_\mu^2)} + \frac{B(n-1)}{n})^2\pi > \sigma^{2*}. \tag{13}$$

This result does not hold, however, if it is easier to engage in dysfunctional actions in more uncertain settings. It is not hard to imagine why such a correlation would arise. For example, consider a surgeon who is rewarded on the basis of mortality rates. Such a surgeon may avoid operating on particularly risky cases. Yet the opportunities for the surgeon to engage in such activities would likely be greater in the case of experimental surgery than with a routine procedure such as an appendectomy is being carried out. Because appendectomies are characterized by little uncertainty, peer review is likely to be alert to surgeon's attempts to avoid a tricky case. This will not be so with experimental procedures, which are are less well known and less subject to standardized guidelines. Thus, dysfunctional responses may be easier to get away with in more uncertain settings.

This correlation is important as it can cause the results above to reverse in such a way that incentives are no longer positively correlated with uncertainty. To see this, consider the following simple form of correlation between the two sources of uncertainty: $\sigma_\mu^2 = \tilde{\sigma}_\mu^2 + k\sigma^2$. Here, k picks up the degree of correlation in how multitasking opportunities vary with the underlying uncertainty of the environment. Then designing the optimal output-based contract still implies a piece rate of $\beta^* = \frac{1}{1+\sigma_\mu^2}$ and the ensuing surplus from the contract (if used) is given by $\frac{\sigma}{\sqrt{\pi}} + \frac{1}{2(1+\sigma_\mu^2)} + \frac{B}{n} - m_y$, which equals

$$\tilde{S} = \frac{\sigma}{\sqrt{\pi}} + \frac{1}{2(1+\tilde{\sigma}_\mu^2 + k\sigma^2)} + \frac{B}{n} - m_y. \tag{14}$$

Unlike in the previous sections, this term (which measures the value of an output-based contract) is no longer necessarily increasing in σ^2. In other words, it is no longer necessarily the case that the value of output contracting is increasing in uncertainty, as the opportunities for dysfunctional behavior also increase in the uncertainty of the environment. Only if k is small (i.e., when dysfunctional behavior is not much more a problem in uncertain than in certain settings) will the result necessarily hold. By contrast, for k sufficiently large, $\tilde{S}$ must decrease in σ^2, because the effect on the dysfunctional responses will be larger than the benefit from a better choice of activity.[22] Indeed, for the case where k is large enough and $m_e > m_y$, the results of the model above are reversed: for low uncertainty, output-based contracts are used, and for higher rates of uncertainty, input-based contracts are used.[23]

This implies that an important assumption to obtain a positive correlation between uncertainty and incentives is that good measures of output are available. The point of this is that a positive relationship between uncertainty and incentives is most likely in tests of franchising, executive pay, or sharecropping, where such measures are available, rather than, for example, expecting output-based incentives in exploratory surgery, which though characterized by considerable uncertainty, is subject to the type of multi-tasking concerns described in Leventis (1998).

6 Empirical Implications

The results of the paper rely on two key assumptions. First, when monitoring of inputs is costly, firms will respond by offering output-based contracts with delegation of tasks. Second, determining the optimal course of action is more difficult in uncertain circumstances, so that delegation of responsibility is more likely in risky settings. There are tests of the first assumption in the context of the franchising decision. These data are, again, taken from a recent survey by Lafontaine and Slade (1998). Specifically, they survey existing evidence to see how (more prosaic) monitoring costs affect the decision to franchise. The available research conclusively shows that franchising is more likely (a (+) in Table 6) when direct monitoring costs are high.

[22] There is one other strategy, ignored here, where the principal chooses to monitor *both* inputs and outputs. Although costly, it may be the optimal way to resolve the agency problems with both effort and activity choice. This becomes optimal when σ^2 is large enough. I have ignored it here simply because I wish to show how the results above can reverse, which will be the case if the monitoring costs are large enough that monitoring of both does not happen.

[23] Little can be said about intermediate levels, as $\tilde{S}$ need not be monotonic in σ^2.

Table 6: Direct Monitoring Costs and the Decision to Franchise.

Authors	Measure Used	Result
Brickley and Dark (1987)	Distance from HQ	+
Bercovitz (1998)	Distance from HQ	+
Minkler (1990)	Distance from HQ	+
Norton (1988)	Rural	+
Brickley et al. (1991)	Low Population Density	+
Carney and Gedajlovic (1991)	Low Population Density	+
Lafontaine (1992)	Absence of Other Franchises in the Zip Code	+
Kehoe (1996)	Absence of Other Franchises in the Zip Code	+

These papers use a measure of direct monitoring costs (how far it is to travel to the franchise, whether there are other franchises nearby, whether the franchise is situated in a rural area, and so on) and show that when monitoring costs are high, franchising is more likely. In terms of the model above, it is more effective simply to delegate decision-making to the agent in such settings, but where they are offered output-based contracts. This supports the first assumption of the model.

The second assumption is that monitoring is more costly in uncertain environments, and that delegation is more likely to occur in such uncertain environments. The empirical difficulty here is that worker discretion is typically unobserved which could bias econometric estimates. A natural implication of this is that without controlling for some measure of responsibility we are likely to find a positive relationship between uncertainty and incentives, but if we can control for task assignment, we would expect to see no such relationship.[24]

The most likely place to observe data on the correlates of agents' responsibilities is in the franchising literature; franchisees are offered more responsibilities than the managers of company-owned stores. This theory suggests that the decision to franchise (and hence delegate responsibility to the agent) will be positively correlated with uncertainty, but *conditional on franchising*, there is no reason to expect such a relationship: perhaps we would observe the traditional negative relationship if franchisees are risk averse. The only evidence I am aware of that tests this proposition is Lafontaine (1992), who considers how uncertainty affects both (i) the decision to franchise, and (ii) the royalty rate offered to the franchisee. She finds (see her Table 5) that the decision to franchise is statistically significant and *positively* related to her measure of uncertainty (the likelihood of bankruptcy), but royalty rates are *negatively* related, though with a t-statistic of only 1.4. This

[24]This model also suggests that workers who report that they have considerable discretion over their jobs should be more likely to be offered pay for performance. The only work that I am aware of which tests this is MacLeod and Parent (1999), who find such a relationship.

clearly suggests that risk plays a different role in the decision to allocate tasks to the agent than in the case where it is simply providing insurance and incentives. As such, I take this result as supportive of the model, in that it suggests that delegation of decision-making occurs in uncertain environments, but conditional on the delegation decision, there is little reason to expect such a relationship.

Rao's (1971) analysis of sharecropping also seems apposite. He examines the frequency of renting versus sharecropping for rice and tobacco farmers in Andhra Pradesh. Although rice yields are characterized by much less uncertainty than tobacco yields, rice farmers are less likely to have fixed rental contracts than tobacco farmers. This is obviously counter to the standard theory, which suggests fixed rents are more common for crops with little uncertainty.

To resolve this puzzle, he notes (as in this paper) that the standard agency model "omits any consideration of the scope for decision making in the face of uncertainty" (p. 582). Such decision making rights are particularly important in situations of uncertainty, as "relative economic certainty in the sense of limited decision making seems to be necessary for the prevalence of sharecropping", whereas "situations of high uncertainty may necessitate fixed-cash rents" (p. 582). Rao specifically attributes this difference to the greater role for "entrepreneurship" in responding to economic conditions for tobacco than for rice. One important feature is how farmers respond to changing relative prices of crops. Changing acerage under cultivation in response to changing relative prices is an important part of the effort decisions that farmers make: higher prices should result in the farming of more marginal land. Rao argues (in standard agency terminology) that the reason for the difference in contracts is that the marginal return to effort in response to such uncertainty is lower in rice than in tobacco. He notes that "the area under rice cultivation reponded little to changes in its relative price" (p. 583), because soil suitable for rice is typically unsuitable for other crops.[25] In effect, there is little response that farmers can make to area grown when the relative price of rice changes. As a result, there is less need to offer agents incentives for rice, whereas for tobacco there are more opportunities to respond, and so fixed rent contracts are more common. As with the model above, this approach stresses how uncertainty affects the returns to costly actions.

Finally, return to the literature on executives. Unlike the evidence on franchising and sharecropping, the literature on the tradeoff between risk and incentives is inconclusive. One possible reason for the absence of a positive correlation between risk and incentives is simply that there is little variation in the delegation of tasks across executives. For instance, if chief executive officers have control over their actions, and are little constrained by their boards of directors, there is

[25] As a result, low elasticities of areas cultivated to relative price are found for rice, in the range of 0.03 to 0.08 (see Rao for the appropriate references).

no reason for expect the omitted variable bias that generates the positive tradeoff. Only in cases where principals respond to uncertainty by delegating more responsibility will a positive tradeoff be expected.

7 Conclusion

The tradeoff between risk and incentives has become a mantra among economists working on agency issues, despite the rather lukewarm evidence in its favor. Indeed, in some occupations there is more convincing evidence of a positive relationship. The objective of this paper has been to understand why we might not expect to find such a relationship in the data.[26] I argue that the marginal returns to delegation are likely to be higher in more uncertain environments, as a principal may have little idea what the right kinds of efforts are in such cases. In more stable environments, they can simply tell the agent what to do, while in riskier settings, they may have little choice but to offer agency contracts to induce appropriate behavior, as there are no other good measures of performance on which to base pay. Thus, incentives and uncertainty are positively related.

[26] See Prendergast, 2000a, b, for some other reasons.

References

Aggarwal, Raj and Andrew Samwick (1998a), "The Other Side of the Tradeoff: The Impact of Risk on Executive Compensation", *Journal of Political Economy*, 107, 65-105.

Aggarwal, Raj and Andrew Samwick (1998b), "Executive Compensation, Strategic Competition and Relative Performance Evaluation: Theory and Evidence", forthcoming, *Journal of Finance*.

Aghion, Philippe and Jean Tirole (1997), "Formal and Real Authority in Organizations", *Journal of Political Economy*, 105, 1-29.

Allen, Douglas and Dean Lueck (1992), "Contract Choice in Modern Agriculture: Crop-Share versus Cash Rent", *Journal of Law and Economics*, 35(2), 397-426.

Allen, Douglas and Dean Lueck (1992), "Risk, Uncertainty and Contracts" Working Paper No.173, Center for Law, Economics and Public Policy, Yale Law School, 1995.

Anderson, E. and D. Schmittlein (1984), "Integration of the Sales Force: An Empirical Examination", *RAND Journal of Economics*, 15, 385-95.

Antle, Richard and Abbie Smith (1986), "An Emprical Investigation of the Relative Rerformance Evaluation of Corporate Executives", *Journal of Accounting Research*, 24. 1-39.

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

Barro, Jason and Robert Barro (1990), "Pay, Performance and the Turnover of Bank CEOs", *Journal of Labor Economics*, 8, 448-81.

Becker, G. (1968), "Crime and Punishment: An Economic Approach." *Journal of Political Economy*, 76, 169-217.

Bercovitz, J. (1998), "Franchising vs. Company Ownership", mimeo, Duke University.

Brickley, James and F. Dark (1987), "The Choice of Organizational Form: The Case of Franchising", *Journal of Financial Economics*, 18, 401-420.

Brickley, James, F. Dark (1987) and Michael Weisbach (1991), "An Agency Perspective on Franchising", *Financial Management*, 20, 27-35.

Brown, Charlie (1990), "Firms' Choice of Method of Pay", *Industrial and Labor Relations Review*, 43, 165-183.

Bushman, Robert, Raffi Indejikian, and Abbie Smith (1996), "CEO Compensation: The Role of Individual Performance Evaluation", *Journal of Accounting and Economics*, 21, 161-93.

Carney, M. and E. Gedajlovic (1991), "Vertical Integration in Franchise Systems: Agency Theory and Resource Explanations", *Strategic Management Journal*, 18, 401-420.

Coughlin, Anne and Chakravarti Narasimhan (1992), "An Empirical Analysis of Sales-Force Compensation Plans", *Journal of Business*, 65(1), 93-122.

Conyon, Simon and Kevin J. Murphy (1999), "The Prince and the Pauper? CEO Pay in the US and UK", mimeo, USC.

Core, John and Wayne Guay (1999), "The Use of Equity Grants to Manage Optimal Equity Incentive Levels", *Journal of Accounting and Economics*, 28, 151-84.

David Cox and D.V. Hinkley (1992), *Theoretical Statistics*, Chapman and Hall, London.

Garen, John (1994), "Executive Compensation and Principal-Agent Theory", *Journal of Political Economy*, 102, 1175-99.

Gibbons, Robert and Kevin J. Murphy (1990), "Relative Performance Evaluation for Chief Executive Officers", *Industrial and Labor Relations Review*, 43, 30-52.

Grossman, Sanford and Oliver Hart (1983), "An Analysis of the Principal-Agent Problem", *Econometrica*, 51, 7-45.

Hallagan, William (1978), "Share Contracting for California Gold", *Explorations in Economic History*, 15, 196-210.

"Hardaway's Chance" (1998), Chicago Tribune.

Holmstrom, Bengt (1979), "Moral Hazard and Observability", *Bell Journal of Economics*, 10, 74-91.

Holmstrom, Bengt and Paul Milgrom (1991), "Multi-Task Principal Agent Analyses: Linear Contracts, Asset Ownership and Job Design", *Journal of Law, Economics, and Organization*, 7, 24-52.

Ittner, Christopher, David Larker and Madhav Rajan (1997), "The Choice of Performance Measures in Annual Bonus Contracts", *The Accounting Review*, 72, 231-55.

Janakiraman, Surya, Richard Lambert and David Larker (1992), "An Empirical Investigation of the Relative Performance Evaluation Hypothesis", *Journal of Accounting Research*, 30, 53-69.

John, G. and Barton Weitz (1988), "Forward Integration into Distribution: An Empirical Test of Transaction Cost Analysis", *Journal Law, Economics and Organization*, 4, 337-355.

John, G. 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.

Kawasaki, Seiichi and John McMillan (1987), "The Design of Contracts: Evidence from Japanese Subcontracting", *Journal of the Japanese and International Economies*, 1, 327-49.

Kehoe, M. (1996), "Franchising, Agency Problems and the Cost of Capital", *Applied Economics*, 28, 1485-93.

Lafontaine, Francine (1992), "Agency Theory and Franchising: Some Empirical Results", *RAND Journal of Economics*, 23, 263-83.

Lafontaine, Francine and Margaret Slade (1998), "Incentive Contracting and the Franchise Design", in *Advances in Business Applications of Game Theory*, K. Chaterjee and W. Samuelson (eds.), Kluyer Academic Press.

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

Lambert, Richard, David Larker and Keith Weigelt (1993), "The Structure of Organizational Incentives", *Administrative Science Quarterly*, 38, 438-61.

Lazear, Edward (2000), "The Power of Incentives", forthcoming, *American Economic Review Papers and Proceedings*, May 2000.

Lazear, Edward and Sherwin Rosen (1981), "Rank Order Tournaments as Optimal Labor Contracts", *Journal of Political Economy*, 89, 841-64.

Leffler, Keith and Randal Rucker (1991), "Transactions Costs and the Organization of Production: Evidence from Timber Sales Contracts", *Journal of Political Economy*, 23, 263-83.

Leventis, Andrew (1997), "Cardiac Surgeons Under the Knife", mimeo, Princeton University.

MacLeod, Bentley and Daniel Parent (1999), "Job Characteristics and the Form of Compensation", *Research in Labor Economics*, Vol. 18, JAI Press, 177-242.

Martin, R.E., (1988), "Franchising and Risk Management", *American Economic Review*, 78, 954-68.

Minkler, A. (1990), "An Empirical Analysis of a Firm's Decision to Franchise", *Economics Letters*, 34, 77-82.

Milgrom, Paul and John Roberts (1988), "An Economic Approach to Influence in Organizations", *American Journal of Sociology*, 94, S154-79.

Mulherin, J. Harold (1986), "Complexity in Long-Term Contracts: An Analysis of NAtural Gas Contractual Provisions", *Journal of Law, Economics and Organization*, 2(1), 105-18.

Murphy, Kevin J. (1999), "Executive Compensation", forthcoming, *Handbook of Labor Economics, Vol. 3*, eds., Orley Ashenfelter and David Card, North Holland.

Norton, S.W. (1988), "An Empirical Look at Franchising as an Organizational Form", *Journal of Business*, 61, 197-217.

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

Prendergast, Canice (2000a), "What Trade-off of Risk and Incentives?", forthcoming, *American Economic Review Papers and Proceedings*, May.

Prendergast, Canice (2000b), "Incentives and Uncertainty", mimeo, University of Chicago.

Prendergast, Canice and Robert Topel (1996), "Favoritism in Organizations", *Journal of Political Economy*, 104:5, 958-78.

Rao, C.H. Hanumata (1971), "Uncertainty, Entrepreneurship, and Sharecropping in India", *Journal of Political Economy*, 85(3), 578-95.

Scott, F. (1995), "Franchising versus Company Ownership as a Decision Variable of the Firm", *Review of Industrial Organization*, 10, 69-81.

Yermack, D (1995), "Do Corporations Award CEO Stock Options Effectively?", *Journal of Financial Economics*, 39, 237-70.