What Trade-off of Risk and Incentives?

By CANICE PRENDERGAST*

The premise of agency theory is not simply that individuals respond to incentives, but that contracts reflect the costs and benefits of inducing appropriate behavior from agents. Foremost among these concerns is the trade-off of risk and incentives, which holds that one factor constraining pay-for-performance by firms is that it imposes risk on employees, which will be reflected in higher levels of compensation. As a result, incentives will be muted in risky environments. However, although the trade-off of risk and uncertainty has occupied center stage in the literature on compensation since early contributions such as Bengt Holmstrom (1979), empirical research has not shown a convincing relationship between pay-for-performance and observed measures of uncertainty. In a recent survey (Prendergast, 1999), I examined this relationship at some length. Many of the empirical studies on the trade-off look at executive compensation; some authors find evidence in favor of the trade-off, while others find none. Studies that consider compensation for non-executives find little evidence of such a trade-off. Beyond these systematic studies, the theory also seems a little strained at an anecdotal level, in 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. If the trade-off of risk and incentives is the primary force determining pay-for-performance, these are hardly the industries that would be predicted to rely heavily on such risky instruments.

I argue here that the theories may be 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. I show that there are a number of simple reasons why one might expect to find a positive relationship between risk and incentive provision which would counter the usual negative relationship that the theory would predict. I discuss each in turn in the context of risk-neutral agents; for more formal details, see Prendergast (2000).

* University of Chicago, Graduate School of Business, 1101 East 58th St., Chicago, IL 60637. I thank Robert Gibbons for very helpful comments.

I. Input Monitoring Is Less Effective in Uncertain Environments

Most work in agency theory restricts attention to output-based contracts, where one extreme is residual claimancy on output and the other is a salary, where pay is independent of (output) performance. But realistically, this is not what firms do when output-based pay disappears. Instead, firms find other means of resolving agency concerns, namely, directing the agent's actions and observing his inputs. For instance, a principal can simply tell the agent what to do and may intensify the monitoring of his inputs on those tasks in the absence of an output-based contract. The fact that a firm can substitute for output-based agency contracts by directing the agent's actions and monitoring is not in itself a problem for the existing theory unless the marginal cost of using such input monitoring depends on uncertainty in the environment. But there is a natural relationship between the effectiveness of input monitoring 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 and identifying what the agent should be working on, she can be pretty sure that private and social benefits are aligned. However, in less certain environments, the principal may be able to monitor inputs (e.g., whether the agent is keeping busy) but be 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.

For example, consider the following simple case. Assume that a firm hires an agent to exert effort on one of two possible tasks. The agent

chooses an effort level to exert, e_i, to activity 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 $C'(0) = 1$. Output y_i from exerting effort on task i depends on the effort level and on a random variable ρ_i in the following way: $y_i = \rho_i + e_i$, $i = 1, 2$. Assume that the random variables are both uniformly distributed, $\rho_i \sim \mathcal{U}[-x + \bar{\rho}_i, x + \bar{\rho}_i]$. All individuals are risk-neutral.

By renormalization, let the distributions be $\rho_1 \sim \mathcal{U}[-x, x]$ and $\rho_2 \sim \mathcal{U}[-x + \Delta, x + \Delta]$, where $\Delta = \bar{\rho}_2 - \bar{\rho}_1$. Therefore, the two activities are uniformly distributed with common variance $x^2/3$, but where activity 2 has a mean that is Δ higher than activity 1. Below, I show that increasing the variance leads to greater returns to output-based contracts, in contrast to the existing literature.

Agents often have important information that is not available to the principal. I assume that the agent knows the true values of ρ_i, while the principal knows only the distribution of the ρ_i's. To conclude the agency problem, I assume that the agent has personal preferences over the two actions. The agent knows his private benefits. As a concrete example, consider the case where the agent has personal preferences such that he is indifferent to one of the activities ($B_i = 0$) but gains a small benefit $B > 0$ from the other, where the principal has no idea of the identity of i. Specifically, 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.

The principal can potentially collect two pieces of information to determine how 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 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 risk costs or multitasking concerns.

The firm has two possibly optimal choices on how to pay the worker. First, it could simply tell the worker which activity he should be engaged in, monitor effort levels on that task, and reward him for that task if he exerts the optimal level of effort. This input-based contract has monitoring costs of m_e. Obviously, under an input-monitoring compensation plan, the agent will be assigned to work on activity 2, as this has a mean which is Δ higher than the other and has expected return $\Delta + 1 - C(e_2^*) - m_e + (B/2)$, where $C(e_2^*)$ is the optimal level of effort. Second, the firm could incur cost m_y and offer an output-based compensation plan with an optimal piece rate of 1. If offered this output-based compensation plan, the agent chooses the highest value of ρ_i (as the private benefits are small), and expected profits are $E[\rho_1'] + 1 - C(e_i^*) + (B/2) - m_y$, where ρ_1' is the first-order statistic of the two realizations.

The value of offering an output-based contract then depends on the importance of the agent choosing the right action (i.e., the difference between $E[\rho_1']$ and Δ. But for the uniform distribution, this is given by

$$(1)\quad E[\rho_1'] - \Delta = \begin{cases} \frac{1}{8}\left(1 - \frac{\Delta}{2x}\right)^2 (\Delta + 2x) & \text{if } \Delta < 2x \\ 0 & \text{otherwise.} \end{cases}$$

Then output-based monitoring is preferred if

$$\max\left\{0, \tfrac{1}{8}\left(1 - \frac{\Delta}{2x}\right)^2 (\Delta + 2x)\right\} \geq m_y - m_e.$$

But remember that the variance of the uniform distribution is $x^2/3$, so that by substitution, output monitoring is preferred if

$$\left(1 - \frac{\Delta}{\sqrt{12}\sigma}\right)^2 (\Delta + \sqrt{12}\sigma) \geq 8(m_y - m_e).$$

The left-hand side of this expression is increasing in σ^2. 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. In essence, the more uncertainty there is in the environment, the more important it is to induce the agent to choose the correct activity rather than assigning him one, which can only be done by basing pay on output.

II. Sorting Is Less Effective in Uncertain Environments

The human-resources management literature, typically carried out by organizational psychologists, contains an enormous amount of information about why firms carry out performance appraisals. I find it interesting that the desire to tie merit pay to the evaluations often ranks quite low on the list of reasons for doing them. As discussed in some of the principal textbooks on human resources management, such as Arthur Sherman et al. (1999) and George Milkovich and Jerry Newman (1999), two common reasons for doing evaluations are to provide feedback to employees on their perceived strengths and weaknesses and to identify talent within the firm.

In itself, this may not be a problem for agency if there are no important interactions between the purposes of these evaluations. However, the optimal agency contract may be affected by these other reasons for evaluation if supervisors do not always tell the truth in performance evaluations, as the data suggest. Supervisors often distort their reports on workers based on their personal preferences. For example, there is considerable evidence of favoritism within firms, where workers who are simply liked by the boss are more likely to get a good performance rating, independent of performance. Equally, there is evidence that supervisors are unwilling to impose low ratings on subordinates, a phenomenon known in the literature as "leniency bias." Most important for our purposes, these distortions are more common in cases where there is "money on the line": when pay is tied to performance, accuracy seems to fall (see Milkovich and Newman [1999] for the evidence).

Consider a firm that uses its evaluation procedures for two purposes: to provide incentives and to identify individual talents for allocating workers to different tasks. This would seem a reasonable assumption in many firms. Assume also that supervisors can distort evaluations to reward favored employees and harm disliked employees. Such distortions come at a personal cost to the supervisor, as the supervisor may have to cover his tracks or may be penalized if such distortions are identified. It follows that, when incentive pay is tied to supervisor evaluations, supervisors distort more, as their actions have more effect on the pay of their favored and unfavored employees. But remember that the firm is using performance appraisals to allocate workers to tasks based on their perceived talents. Then if supervisors lie more, more mistakes are made in the assignment of workers to tasks, as sometimes the "teacher's pet" gets a better job even though he is not talented enough for that position, and vice versa for a disliked employee. This naturally generates a trade-off between incentives and sorting, where high incentive pay causes agents to work hard, but at the cost of supervisors' reports being less informative about talents.

The optimal degree of incentive pay from this perspective then depends on the value of the supervisor's information for task assignment and the importance of the agent's efforts. The relationship between incentive pay and risk is then straightforward. Consider a risky environment, one where the supervisor obtains a noisy measure of the agent's talents. Even if the supervisor tells the truth in his evaluation, the information is of limited use because the environment is so noisy. Consequently, the marginal cost to the firm of the supervisor's distorting his evaluation is low. But if the value of the information for sorting purposes is low, the marginal cost of providing incentives to the worker is also low, because the cost of incentive provision is distorted performance evaluations. As a result, the firm provides the worker with considerable pay-for-performance in uncertain environments. By contrast, in less risky environments, supervisors have more valuable information, and so incentive pay is eschewed because the value of truthful reporting of that information is high. Once again, this insight generates a positive relationship between risk and incentives.

III. Investigations Are Less Effective in Uncertain Environments

The third reason to expect a positive relationship between incentives and uncertainty concerns endogenous monitoring and investigations. One of the assumptions typically used in the agency literature is that the principal gets costless signals on the agent's efforts and always monitors the performance of the agent. But this is not realistic: in many situations, monitoring is sporadic and is

based on some endogenous observed phenomenon. For example, consider a bureaucrat who has the opportunity to engage in petty corruption. During the typical day, there is little chance of oversight, and verifiable information on the worker's performance arises only by a formal investigation. Such investigations are infrequent and typically triggered by an endogenous event such as a customer complaint, or when a suggestion of wrongdoing reaches the principal. Moreover, investigations occur only when bad news is expected. The role of noisy environments then arises in the relationship between suspicion of malfeasance and truth. In noisy environments, the link is less clear than in simple settings, where initial impressions are rarely overturned. Consequently, in noisy environments, agents realize that they will sometimes "get away with it" even if they are monitored: to overcome this danger, greater incentives are necessary in riskier environments.

More formally, consider a standard agency setting, but where sometimes the principal does not observe anything about the worker's performance. In other cases, he receives some initial *nonverifiable* impression of the worker's performance. For example, a student could claim that someone else has cheated on an exam; or a bureaucrat might be suspected of spending beyond his income. The principal uses this information to decide whether to investigate the performance of the agent. Two results arise from this setting, as described in Prendergast (2000). First, the principal has an incentive to monitor only when the performance of the worker is poor. Consequently, risk of the environment will play an important role. Specifically, it is shown that when the initial impression that induces the investigation is riskier (so there is more noise in the initial information), not surprisingly, the agent suffers less from an investigation. As a result, for a given contract, incentives are lower in riskier settings. In order to induce appropriate incentives for monitoring and effort, the firm then chooses higher pay-for-performance in cases where the environment is riskier.

IV. Career Concerns Are Less Effective in Uncertain Environments

The existing literature describes at least one other reason why one would expect to see incentives positively correlated with risk. This is the career-concerns model of Holmstrom (1982), where agents are sometimes willing to exert effort in order to generate a reputation. For example, consider an employee who operates in an environment with little uncertainty. He realizes that good performance today is informative of his talent, which increases his future pay. Thus, in a certain environment, the worker has an incentive to exert effort in the absence of a contract explicitly tying current pay to performance. Suppose instead that the environment is characterized by substantial uncertainty, where output is relatively uninformative of the worker's talent. In that case, career concerns do not function well as a means of inducing effort exertion. As a result, career concerns may be a weak incentive to perform better. To compensate for this, incentive pay must be higher in riskier situations, as this is the only plausible way to induce effort exertion. Once again, incentives and risk are positively correlated. See Robert Gibbons and Kevin J. Murphy (1992) and Prendergast (1999) for more details on this.

V. Conclusion

The trade-off between risk and incentives has become a mantra among economists working on agency issues, despite the lukewarm evidence in its favor. The objective of this paper has been to provide a series of theoretical reasons why one might not expect to find such a relationship in the data. It is important to note here that I do not claim that one 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 trade-off of risk and incentives should be reflected in the data.

REFERENCES

Gibbons, Robert and Murphy, Kevin J. "Optimal Incentive Contracts in the Presence of Career Concerns." *Journal of Political Economy*, June 1992, *100*(3), pp. 468–506.

Holmstrom, Bengt. "Moral Hazard and Observ-

ability." *Bell Journal of Economics*, Spring 1979, *10*(1), pp. 74–91.

_____. "Managerial Incentive Problems: A Dynamic Perspective." Working paper, Kellogg School of Management, Northwestern University, 1982.

Milkovich, George and Newman, Jerry (with Carolyn Milkovich). *Compensation*, 6th Ed. Boston, MA: Irwin, McGraw-Hill, 1999.

Prendergast, Canice. "The Provision of Incentives in Firms." *Journal of Economic Literature*, March 1999, *37*(1), pp. 7–63.

_____. "The Tenuous Trade-off of Risk and Incentives." Working paper, University of Chicago, 2000.

Sherman, Arthur; Bohlander, George and Snell, Scott. *Managing human resources*, 11th Ed. Cincinnati, OH: SouthWestern, 1999.