

Motivating Innovation

GUSTAVO MANSO*

ABSTRACT

Motivating innovation is important in many incentive problems. This paper shows that the optimal innovation-motivating incentive scheme exhibits substantial tolerance (or even reward) for early failure and reward for long-term success. Moreover, commitment to a long-term compensation plan, job security, and timely feedback on performance are essential to motivate innovation. In the context of managerial compensation, the optimal innovation-motivating incentive scheme can be implemented via a combination of stock options with long vesting periods, option repricing, golden parachutes, and managerial entrenchment.

BERLE AND MEANS's (1932) seminal contribution brought to light the potential drawbacks associated with the separation of ownership and control. For example, in large corporations, shareholders delegate decision rights to a manager, who has the ability to manage resources to his own advantage. To alleviate possible conflicts of interest between shareholders and managers, incentive plans that align their interests are commonly used in practice.

A large literature, which includes Harris and Raviv (1978) and Holmstrom (1979), studies this issue using principal-agent models. In these models, the principal offers the agent an incentive plan that induces the agent to act in the principal's best interest. Most of the papers in this literature focus on the problems of inducing the agent to exert effort or deterring the agent from tunneling resources away from the corporation.

Here, I study a different problem, namely, how to structure incentives when the principal needs to motivate the agent to be more innovative. Such a problem arises naturally in different situations. For instance, shareholders may want to motivate a CEO to pursue more innovative business strategies; managers of

*Gustavo Manso is at MIT Sloan School of Management. I thank John Roberts, Anat Admati, Darrell Duffie, and Ed Lazear for valuable discussions and suggestions. I also thank Peter DeMarzo, Bob Gibbons, Thomas Hellman, Bengt Holmstrom, Jim March, Pedro Miranda, Stew Myers, Tomasz Sadzik, Maria Salgado, Yuliy Sannikov, Antoinette Schoar, Andy Skrzypacz, Ilya Strebulaev, Alexei Tchisty, Jeff Zwiebel, and seminar participants at the University of California at Berkeley, University of Chicago, Columbia University, Duke University, Harvard University, London Business School, London School of Economics, Massachusetts Institute of Technology, Northwestern University, New York University, Princeton University, Stanford University, University of California at Los Angeles, University of California at San Diego, University of Iowa, University of Oregon, and Wharton for helpful comments and Forrest Funnell for outstanding research assistance. Financial support from a Stanford Institute for Policy Research (SIEPR) Fellowship and the Lehman Brothers Fellowship for Research Excellence in Finance is gratefully acknowledged.

large corporations may want to induce their employees to be more innovative; or regulators may want to stimulate entrepreneurship, say, through the design of bankruptcy laws.

The key contribution of the paper is to show that incentive schemes that motivate innovation should be structured differently from standard pay-for-performance schemes used to induce effort or avoid tunneling. Innovation involves the exploration of new untested approaches that are likely to fail. Therefore, standard pay-for-performance schemes that punish failures with low rewards and termination may in fact have adverse effects on innovation. In contrast, the optimal incentive scheme that motivates innovation exhibits substantial tolerance (or even reward) for early failure and reward for long-term success. Under this incentive scheme, compensation depends not only on total performance, but also on the path of performance; an agent who performs well initially but poorly later earns less than an agent who performs poorly initially but well later or even an agent who performs poorly repeatedly. The paper also shows that commitment to a long-term compensation plan, job security, and timely feedback on performance are essential to motivate innovation.

In the context of executive compensation, the optimal contract that motivates innovation can be implemented via a combination of stock options with long vesting periods, option repricing, golden parachutes, and managerial entrenchment. Stock options with long vesting periods combined with option repricing and golden parachutes bring on tolerance for early failure and reward for long-term success, so that compensation depends not only on total performance but also on the path of performance as described above. Managerial entrenchment gives the manager job security, since an entrenched manager may keep his job even if it is ex-post efficient for the shareholders of the firm to fire him.

In the public debate on corporate governance, critics often hold that golden parachutes, option repricing, and managerial entrenchment are detrimental because they protect or even reward the manager after poor performance, potentially undermining the incentives for the manager to exert effort. Occasionally there are proposals to adopt regulations that restrict the use of some these practices.¹ However, as argued here, these practices may be part of an optimal incentive scheme that motivates innovation, in which case regulations that restrict their use may have an adverse effect on innovation. To assess the actual impact of such regulations, it remains to be studied empirically the actual contribution of these practices to innovation.

To model the innovation process, I use a class of Bayesian decision models known as bandit problems.² In bandit problems, the agent is uncertain about the true distribution of the available actions' payoffs. Innovation in this setting is the discovery, through experimentation and learning, of actions that are superior to previously known actions. I focus on the central concern

¹ See, for example, Bebchuk and Fried (2004) and "Rewards for Failure," British DTI consultation, June 2003.

² Berry and Fristedt (1985) provide an introduction to the statistical literature on bandit problems. Bergemann and Valimaki (2006) survey the applications of bandit problems to economics.

that arises in bandit problems: the tension between the *exploration* of new untested actions and the *exploitation* of well-known actions. Exploration of new untested actions reveals information about potentially superior actions, but is also likely to waste time with inferior actions. Exploitation of well-known actions ensures reasonable payoffs, but may prevent the discovery of superior actions.

To study the incentives for exploration and exploitation, I embed a bandit problem into a principal-agent framework. The principal-agent relationship could be between, for example, shareholders of a firm designing the compensation package offered to a manager, the manager of a firm designing an incentive plan for a worker, or a venture capitalist financing an entrepreneur.

The model has two periods and two possible outcomes in each period (success or failure). In each period, the agent can choose between shirking, exploiting a well-known approach, or exploring a novel approach that has an unknown probability of success. There are two important special cases. First, if exploration and exploitation are costless to the agent, there is no conflict of interest between the principal and the agent. The model then reduces to the two-armed bandit problem that captures the tension between exploration and exploitation. Second, if exploration is extremely costly to the agent, the agent chooses between exploitation or shirking. The model then reduces to a standard principal-agent model in which the principal must motivate the agent to exert effort. Therefore, the model developed here incorporates the tension between exploration and exploitation present in bandit problems, as well as the tension between working and shirking present in standard principal-agent models.

The optimal contracts that motivate exploitation and exploration are fundamentally different from each other. Since exploitation is just the repetition of well-known actions, the optimal contract that motivates exploitation is similar to standard pay-for-performance contracts used to motivate repeated effort. On the other hand, since with exploration the agent is likely to waste time with inferior actions, the optimal contract that motivates exploration exhibits substantial tolerance (or even reward) for early failures. Moreover, since exploration reveals information that is useful for future decisions, the optimal contract that motivates exploration rewards long-term success.

The threat of termination after poor performance also affects the incentives for exploration and exploitation. Since the threat of termination discourages the agent from shirking or exploring new actions, termination facilitates the provision of incentives for exploitation. Excessive termination may thus be optimal to motivate exploitation. In contrast, the effects of termination on the incentives for exploration are ambiguous. On the one hand, the threat of termination prevents the agent from shirking. On the other hand, the threat of termination encourages the agent to exploit conventional actions. Depending on which of these two effects is more important, termination may either facilitate or hinder the provision of incentives for exploration. Thus, either excessive termination or continuation may be optimal to motivate exploration.

The roles of feedback on performance and commitment to a long-term contract also differ depending on whether the principal wants to motivate exploration or exploitation. While not important to motivate exploitation, commitment to a long-term contract and timely feedback on performance are essential to motivate exploration.

The model produces testable empirical implications. For example, one can study whether incentive practices used in tasks for which motivating innovation is more (less) important resemble the incentive practices that, according to this paper, motivate exploration (exploitation). Moreover, one can study whether the adoption of incentive practices that motivate exploration (exploitation) leads to more (less) innovation. Later in the paper, I discuss how to interpret and test these predictions in the context of executive compensation, bankruptcy laws, and intrapreneurship in large corporations.

The paper is organized as follows. Section I discusses the related literature. Section II discusses the tension between exploration and exploitation in a single-agent decision problem. Section III introduces the tension between exploration and exploitation to a principal-agent model. Section IV studies incentives for exploration and exploitation. Section V studies implementation without commitment. Section VI studies the optimal termination policy of the principal. Section VII studies the provision of feedback. Section VIII discusses empirical implications and applications of the model to corporate governance and executive compensation, bankruptcy laws and entrepreneurship, and the incentives for intrapreneurship in large corporations. Section IX contains additional discussion and Section X concludes. All proofs are in the Appendix.

I. Related Literature

Prior papers study the incentives for innovation from an optimal contracting perspective. Holmstrom (1989) proposes an alternative explanation for why incentive schemes that motivate innovation must exhibit tolerance for failures. He argues that performance measures for innovative activities are noisier, and therefore to motivate innovation the principal should rely on compensation schemes that are less sensitive to performance. In the same vein, Aghion and Tirole (1994) argue that the outcomes of innovation activities are unpredictable and, hence, hard to contract upon *ex ante*. In an incomplete contract framework, they derive the optimal allocation of control rights that motivates innovation. These two papers focus on measurability and contractability aspects of innovation activity. In contrast, the present paper models the innovation process explicitly and focuses on the central trade-off that arises in innovation activities, the trade-off between exploration and exploitation.

The model of the innovation process adopted here follows a long tradition in the study of innovation. Schumpeter (1934) argues that innovation results from experimentation with “new combinations” of existing resources. Arrow (1969) associates innovation with the production of knowledge and proposes the use of Bayesian decision models to study innovation. Bandit problems

are Bayesian decision models that allow for knowledge acquisition through experimentation. Weitzman (1979) applies a simple bandit problem to study the innovation process. March (1991) uses the terms exploration and exploitation to describe the fundamental tension that arises in learning through experimentation. The literature in industrial organization, including Roberts and Weitzman (1981), Jensen (1982), Battacharya, Chatterjee, and Samuelson (1986), and Moscarini and Smith (2001), relies extensively on bandit problem and related models of learning through experimentation to study the innovation process. Also, recent papers on growth theory, such as Jovanovic and Rob (1990), Jovanovic and Nyarko (1996), and Aghion (2002), develop detailed models of innovation as the result of learning from the exploration of new technologies. In contrast to the above papers, which study individual decision problems, I embed bandit problems into a principal-agent framework to study incentives for exploration and exploitation.

Other papers study principal-agent models in which the agent's choice is not limited to the level of effort. Holmstrom and Milgrom (1991) develop a multitask principal-agent model in which the agent allocates effort across multiple tasks and the principal observes a performance measure for each of these tasks. They show that increasing compensation in one task will cause some reallocation of attention away from other tasks, and therefore pay-for-performance contracts may not be optimal. The current model resembles the multitask model in that the agent can choose to allocate effort to exploration or exploitation, and the intuition for why standard pay-for-performance is suboptimal when motivating innovation is related to the multitask intuition. However, even though the two models share some features, modeling the innovation process explicitly as a bandit problem and analyzing optimal incentives in that setting leads to richer predictions and insights about how to provide incentives for innovation. In particular, by examining a dynamic model, I am able to show that the compensation of the agent depends not only on total performance but also on the path of performance. Moreover, the results on lack of commitment, termination, and feedback are not present in multitask principal-agent models.

The paper is also related to the managerial short-termism literature, which argues that managers are biased toward short-term projects due to career concerns (Narayanan (1985)), takeover threats (Stein (1988)), concerns about near-term stock prices (Stein (1989)), the presence of noise traders (Shleifer and Vishny (1990)), and herding behavior (Zwiebel (1995)). In a survey of financial executives, Graham, Harvey, and Rajgopal (2005) find that the majority of managers would pass on a positive net present value (NPV) project to avoid missing the current quarter's consensus earnings forecast. Using firm-level data, Dechow and Sloan (1991) and Bushee (1998) find that short-termism is more prevalent for CEOs near retirement and in firms held by transient institutional investors.

To study the financing of innovation, Bergemann and Hege (2005) develop a principal-agent model in which there is learning about the quality of the project. The tension between exploration and exploitation does not arise in

their model though, as the agent can only choose one type of project. Moreover, their paper only considers implementation with a sequence of short-term contracts. Also related are Hellmann (2007) and Hellmann and Thiele (2009), who study incentives for innovation using a multitask principal-agent model.

Some contemporaneous papers find empirical support for a few of the results derived here. For example, Acharya and Subramanian (2009) find that debtor-friendly bankruptcy laws lead to more innovation. Acharya, Baghai-Wadji, and Subramanian (2009) find that stringent labor laws that restrict the dismissal of employees encourage firm-level innovation. Francis, Hasan, and Sharma (2009) show that golden parachutes, as well as long-term incentives in the form of vested and unvested options, have a positive and significant effect on patents and citations to patents. Tian and Wang (2010) find that firms backed by venture capitalists that tolerate failure are significantly more innovative. Seru (2010) provides evidence that high-level managers in conglomerates are more reluctant to invest in novel projects because of the threat of reallocation of resources by headquarters. Azoulay, Graff Zivin, and Manso (2010) show that funding policies with tolerance for early failure and long horizons to evaluate results motivate creativity in scientific research. In a laboratory experiment, Ederer and Manso (2010) show that compensation schemes that tolerate early failure and reward long-term success encourage innovation. Several empirical predictions remain untested, however, and are discussed in more detail in Section VIII.

More generally, the paper is related to the literature that studies the effect of institutions on financial organization and economic growth. Previous work in this literature documents the effect of the institutional environment (e.g., corporate governance, investor protection, and political risk) on financial development (LaPorta et al. (1997)) and economic growth (King and Levine (1993), Rajan and Zingales (1998), Acemoglu, Johnson, and Robinson (2002), Bekaert, Harvey, and Lundblad (2005)). Since innovation is an essential ingredient for growth (Romer (1986), Aghion and Howitt (1992)), the findings of the current paper complement this literature and provide directions for future research on the topic.

II. The Tension between Exploration and Exploitation

In this section, I review the classical two-armed bandit problem with one known arm. This model illustrates the tension between exploration and exploitation in a single-agent decision problem.

The agent lives for two periods. In each period, the agent takes an action $i \in \mathcal{I}$, producing output S (success) with probability p_i or output F (failure) with probability $1 - p_i$. The probability p_i of success when the agent takes action $i \in \mathcal{I}$ may be unknown. To obtain information about p_i , the agent needs to engage in experimentation. I let $E[p_i]$ denote the unconditional expectation of p_i , $E[p_i|S, j]$ denote the conditional expectation of p_i given a success on action j , and $E[p_i|F, j]$ denote the conditional expectation of p_i given a failure on

action j . When the agent takes action $i \in \mathcal{I}$, he only learns about the probability p_i , so that

$$E[p_j] = E[p_j|S, i] = E[p_j|F, i] \quad \text{for } j \neq i.$$

The central concern that arises when the agent learns through experimentation is the tension between exploration of new actions and exploitation of well-known actions. To focus on the tension between exploration and exploitation, I assume that in each period the agent chooses between two actions. Action 1, the conventional work method, has a known probability p_1 of success such that

$$p_1 = E[p_1] = E[p_1|S, 1] = E[p_1|F, 1].$$

Action 2, the new work method, has an unknown probability p_2 of success such that

$$E[p_2|F, 2] < E[p_2] < E[p_2|S, 2].$$

I assume that the new work method is of exploratory nature. This means that when the agent experiments with the new work method, he is initially not as likely to succeed as when he conforms to the conventional work method. However, if the agent observes a success with the new work method, then the agent updates his beliefs about the probability p_2 of success with the new work method, so that the new work method becomes perceived as better than the conventional work method. This is captured as follows:

$$E[p_2] < p_1 < E[p_2|S, 2]. \tag{1}$$

The agent is risk-neutral and has a discount factor normalized to one. The agent thus chooses an action plan (i_k^j) to maximize his total expected payoff

$$\begin{aligned} R((i_k^j)) = & \{E[p_i|S] + (1 - E[p_i])F\} \\ & + E[p_i] \{E[p_j|S, i]S + (1 - E[p_j|S, i])F\} \\ & + (1 - E[p_i]) \{E[p_k|F, i]S + (1 - E[p_k|F, i])F\}, \end{aligned} \tag{2}$$

where $i \in \mathcal{I}$ is the first-period action, $j \in \mathcal{I}$ is the second-period action in the case of success in the first period, and $k \in \mathcal{I}$ is the second-period action in the case of failure in the first period.

Two action plans need to be considered. Action plan (1_1^1) , which I call exploitation, is just the repetition of the conventional work method. Action plan (2_1^2) , which I call exploration, is to initially try the new work method, stick to the new work method in the case of success in the first period, and revert to the conventional work method in the case of failure in the first period. The total payoff $R((2_1^2))$ from exploration is higher than the total payoff $R((1_1^1))$ from exploitation if and only if

$$E[p_2] \geq p_1 - \frac{p_1(E[p_2|S, 2] - p_1)}{1 + (E[p_2|S, 2] - p_1)}. \tag{3}$$

If the agent tries the new work method, he obtains information about p_2 . This information is useful for the agent's decision in the second period, since the agent can switch to the conventional work method if he learns that the new work method is not worth pursuing. The agent may thus be willing to try the new work method even though the initial expected probability $E[p_2]$ of success with the new work method is lower than the probability p_1 of success with the conventional work method. The second term on the right-hand side of equation (3) represents the premium in terms of first-period payoff that the agent is willing to pay to obtain information about p_2 .

The agent is willing to sacrifice more in the first period if he lives for multiple periods. With multiple periods, the benefits of experimenting with the new work method are higher, since the agent can use the information he learns from experimentation for a longer period of time. The same is true if the problem is to maximize the output of a team. In a team, the optimal action plan involves more sacrificing of first period output for at least one of the agents. If the agent discovers that the new work method is better than the conventional work method, the whole team benefits from his discovery.

III. The Principal-Agent Problem

In this section, I introduce incentive problems to the classical two-armed bandit problem with one known arm as reviewed in the previous section.

The principal hires an agent to perform the task described in the previous section. In each period, the agent incurs private costs $c_1 \geq 0$ if he takes action 1, the conventional work method, and private costs $c_2 \geq 0$ if he takes action 2, the new work method, but can avoid these private costs by taking action 0, shirking.

The costs c_1 and c_2 associated with the new and conventional work methods will be important in determining the form of the optimal contract. When c_2 is high relative to c_1 , it is more costly for the agent to employ the new work method than the conventional work method, perhaps because it takes more effort for the agent to search and implement new ideas. When c_1 is high relative to c_2 , it is more costly for the agent to employ the conventional work method than the new work method, perhaps because the agent dislikes monotonous work or extracts private benefits from learning new work methods.

Shirking has a lower probability of success than either of the two work methods, so that

$$p_0 < E[p_i] \quad \text{for } i = 1, 2. \quad (4)$$

I assume that the principal does not observe the actions taken by the agent.³ As such, before the agent starts working, the principal offers the agent a contract $\vec{w} = \{w_F, w_S, w_{SF}, w_{SS}, w_{FF}, w_{FS}\}$ that specifies the agent's wages

³This assumption is important because it implies that the principal will rely on output-based pay to provide incentives to the agent. In Section IX, I discuss alternatives to the nonobservability assumption and the consequences of relaxing it.

contingent on future performance. The agent has limited liability, meaning that his wages cannot be negative.

Both the principal and the agent are risk-neutral and have a discount factor of one. When the principal offers the agent a contract \bar{w} and the agent takes action plan $\langle i_k^j \rangle$, the total expected payments from the principal to the agent are given by

$$\begin{aligned} W(\bar{w}, \langle i_k^j \rangle) = & \{E[p_i]w_S + (1 - E[p_i])w_F\} \\ & + E[p_i] \{E[p_j|S, i]w_{SS} + (1 - E[p_j|S, i])w_{SF}\} \\ & + (1 - E[p_i]) \{E[p_k|F, i]w_{FS} + (1 - E[p_k|F, i])w_{FF}\}. \end{aligned}$$

When the agent takes action plan $\langle i_k^j \rangle$, the total expected costs incurred by the agent are given by

$$C(\langle i_k^j \rangle) = c_i + E[p_i]c_j + (1 - E[p_i])c_k.$$

I say that \bar{w} is an optimal contract that implements action plan $\langle i_k^j \rangle$ if it minimizes the total expected payments from the principal to the agent,

$$W(\bar{w}, \langle i_k^j \rangle),$$

subject to the incentive compatibility constraints,⁴

$$W(\bar{w}, \langle i_k^j \rangle) - C(\langle i_k^j \rangle) \geq W(\bar{w}, \langle l_n^m \rangle) - C(\langle l_n^m \rangle). \quad (\text{IC}\langle l_n^m \rangle)$$

This is a linear program with six unknowns and 27 constraints. When more than one contract solves this program, I restrict attention to the contract that pays the agent earlier.⁵

The principal's expected profit $\Pi(\langle i_k^j \rangle)$ from implementing action plan $\langle i_k^j \rangle$ is given by

$$\Pi(\langle i_k^j \rangle) = R(\langle i_k^j \rangle) - W(\bar{w}(\langle i_k^j \rangle), \langle i_k^j \rangle), \tag{5}$$

where $R(\langle i_k^j \rangle)$ is the principal's total expected revenue when the agent uses action plan $\langle i_k^j \rangle$, and $\bar{w}(\langle i_k^j \rangle)$ is the optimal contract that implements action plan $\langle i_k^j \rangle$. The principal thus chooses the action plan $\langle i_k^j \rangle$ that maximizes $\Pi(\langle i_k^j \rangle)$.

Both the classical two-armed bandit problem and the standard work-shirk principal-agent model are special cases of this model. On the one hand, when $c_1 = c_2 = 0$, there is no conflict of interest between the principal and the agent. Therefore, the principal does not need to provide incentives to the agent, and hence just solves the two-armed bandit problem described in Section II. On the

⁴ For simplicity, I assume that the agent has zero reservation utility. The participation constraint is thus not binding, since the agent has limited liability.

⁵ The other contracts that solve the above program are similar to the contract analyzed here except that the principal acts as a bank, keeping the wages of the agent to be paid later without obtaining any additional benefits from doing so. The contract analyzed here is also the contract that arises if the agent is slightly more impatient than the principal.

other hand, when $c_2 = \infty$, it is too costly for the agent to employ the new work method. The agent either shirks or employs the conventional work method. The principal's problem is thus just to prevent the agent from shirking, as in standard principal-agent models.

The assumptions in the principal-agent problem studied here are standard except that there is learning about the technology being employed. This gives rise to the tension between exploration and exploitation, since there is nothing to be learned about the conventional technology, but a lot to be learned about the new technology.

IV. Incentives for Exploration and Exploitation

In this section, I study the optimal contracts that implement exploration and exploitation. In the Internet Appendix,⁶ I study the principal's choice between exploration and exploitation and the distortions relative to the first-best produced by agency costs.

For clarity of exposition, I restrict attention to

$$c_2/c_1 \geq (E[p_2] - p_0)/(p_1 - p_0). \quad (6)$$

The right-hand side of equation (6) is lower than one. Restricting attention to (6) thus rules out situations in which the cost of employing the new work method is much lower than the cost of employing the conventional work method. Similar results hold without this restriction. However, the analysis is more complicated and does not add new insights.

A. Incentives for Exploitation

Proposition 1 derives the optimal contract that implements exploitation. Recall from Section II that exploitation is given by the action plan $\langle 1_1^1 \rangle$. The following definitions will be useful in stating Proposition 1:

$$\alpha_1 = \frac{c_1}{p_1 - p_0},$$

$$\beta_1 = \frac{(E[p_2] - p_0) + E[p_2](E[p_2|S, 2] - p_0)}{(p_1 - p_0) + E[p_2](p_1 - p_0)}.$$

PROPOSITION 1: *The optimal contract \bar{w}_1 that implements exploitation is such that*

$$w_F = w_{SF} = w_{FF} = 0,$$

$$w_{SS} = w_{FS} = \alpha_1,$$

⁶An Internet Appendix for this article is available online at <http://www.afajof.org/supplements.asp>.




Figure 1. The optimal contract that implements exploitation under the base-case parameters. The contract resembles a repetition of standard pay-for-performance contracts. Total pay of the agent depends only on total output, except if c_2/c_1 is small, in which case the agent gets an extra reward for early success to prevent exploration.

$$w_S = \alpha_1 + \frac{c_1(1 + E[p_2])}{(p_1 - E[p_2])} \left(\beta_1 - \frac{c_2}{c_1} \right)^+,$$

where $(x)^+ = \max(x, 0)$.

The formal proofs of all the propositions are in the Appendix. However, the main intuition behind Proposition 1 is as follows. To implement exploitation, the principal must prevent the agent from both shirking and exploring. If c_2 is high relative to c_1 , only shirking constraints are binding, and thus the optimal contract that implements exploitation is similar to the optimal contract used to induce the agent to exert effort in a standard word-shirk principal-agent model. If c_2 is low relative to c_1 , the exploration constraint is binding. To prevent exploration, the principal must pay the agent an extra premium in the case of success in the first period. This extra premium is decreasing in c_2/c_1 , since as c_2/c_1 increases the agent becomes less inclined to explore. Figure 1 shows the optimal contract \bar{w}_1 that implements exploitation for different values of c_2/c_1 under the base-case parameters.⁷

⁷ The base-case parameters used in all the figures are $p_0 = 0.25, E[p_2] = 0.3, p_1 = 0.5, E[p_2|S, 2] = 0.7$, and $c_1 = 1$. From Bayes's rule, $E[p_2|F, 2] = 0.129$. Each of the graphs in the figure corresponds to a wage paid to the agent in a particular contingency for different values of c_2/c_1 . When a node has no graphs, it is because the wage paid to the agent in that contingency is zero.

B. Incentives for Exploration

Proposition 2 derives the optimal contract that implements exploration. Recall from Section II that exploration is given by action plan (2_1^2) . The form of the optimal contract that implements exploration will depend on whether exploration is moderate or radical.

DEFINITION 1. *Exploration is radical if*

$$\frac{1 - E[p_2]}{1 - p_1} \geq \frac{E[p_2]E[p_2|S, 2]}{p_1^2},$$

and moderate otherwise.

Exploration is radical if the likelihood ratio between exploration and exploitation of a failure in the first period is greater than the likelihood ratio between exploration and exploitation of two consecutive successes. I call this exploration radical because it has a high expected probability of failure in the first period relative to the probability of failure of the conventional action.

The following definitions will also be useful in stating Proposition 2:

$$\alpha_2 = \max_{j \in \{0,1\}} \frac{(1 + E[p_2])c_2 - p_0c_j}{E[p_2]E[p_2|S, 2] - p_0E[p_j]} + \frac{(E[p_2] - p_0)p_0\alpha_1}{E[p_2]E[p_2|S, 2] - p_0E[p_j]},$$

$$\beta_2 = \frac{(E[p_2]E[p_2|S, 2] - p_0p_1) + E[p_2](p_1E[p_2|S, 2] - p_0p_1)}{(p_1^2 - p_0p_1) + E[p_2](p_1^2 - p_0p_1)}.$$

PROPOSITION 2: *The optimal contract \bar{v}_2 that implements exploration is such that*

$$w_{FS} = \alpha_1, \quad \text{and} \quad w_S = w_{SF} = w_{FF} = 0.$$

If exploration is moderate, then $w_F = 0$ and

$$w_{SS} = \alpha_2 + \frac{p_1 - p_0}{E[p_2]E[p_2|S, 2] - p_0p_1} \frac{p_1(1 + E[p_2])c_1}{E[p_2]E[p_2|S, 2] - p_1^2} \left(\frac{c_2}{c_1} - \beta_2 \right)^+.$$

If exploration is radical, then

$$w_F = \frac{p_1(1 + E[p_2])c_1}{E[p_2]E[p_2|S, 2] - p_1E[p_2]} \left(\frac{c_2}{c_1} - \beta_2 \right)^+$$

and

$$w_{SS} = \alpha_2 + \frac{E[p_2] - p_0}{E[p_2]E[p_2|S, 2] - p_0p_1} \frac{p_1(1 + E[p_2])c_1}{E[p_2]E[p_2|S, 2] - p_1E[p_2]} \left(\frac{c_2}{c_1} - \beta_2 \right)^+.$$

To implement exploration, the principal must prevent the agent from shirking or exploiting. The principal does not make payments to the agent after a

failure in the second period, since this only gives incentives for the agent to shirk. Moreover, the principal does not make payments to the agent after a success in the first period for two reasons. First, rewarding first-period success gives the agent incentives to employ the conventional work method in the first period, since the initial expected probability $E[p_2]$ of success with the new work method is lower than the probability p_1 of success with the conventional work method. Second, in the case of success in the first period, additional information about the first-period action is provided by the second-period performance, since the expected probability of success with the new work method in the second period depends on the action taken by the agent in the first period. Delaying compensation to obtain this additional information is thus optimal. Although there are 27 incentive compatibility constraints, it is easy to see that only a few may bind. The relevant incentive compatibility constraints are

$$(p_1 - p_0)w_{FS} \geq c_1, \tag{IC}_{(2_0^2)}$$

$$(E[p_2]E[p_2|S, 2] - p_0^2)w_{SS} - (E[p_2] - p_0)w_F - (E[p_2] - p_0)p_1w_{FS} \geq c_2 + E[p_2](c_2 - c_1) + p_0c_1, \tag{IC}_{(0_1^0)}$$

$$(E[p_2]E[p_2|S, 2] - p_0p_1)w_{SS} - (E[p_2] - p_0)w_F - (E[p_2] - p_0)p_1w_{FS} \geq c_2 + E[p_2](c_2 - c_1), \tag{IC}_{(0_1^1)}$$

and

$$(E[p_2]E[p_2|S, 2] - p_1^2)w_{SS} + (p_1 - E[p_2])w_F + (p_1 - E[p_2])p_1w_{FS} \geq (1 + E[p_2])(c_2 - c_1). \tag{IC}_{(1_1^1)}$$

The first three incentive compatibility constraints are associated with shirking. The last incentive compatibility constraint is associated with exploitation. One important thing to note is that w_F enters with a positive sign on the left-hand side of the incentive compatibility constraint associated with exploitation. Rewarding the agent for first-period failures may be useful to prevent the agent from exploiting, since the initial expected probability $(1 - E[p_2])$ of failure when the agent employs the new work method is higher than the probability $(1 - p_1)$ of failure when the agent employs the conventional work method.

The first incentive compatibility constraint is always binding. To prevent the agent from shirking in the second period after a failure in the first period, the principal pays the agent $w_{FS} = \alpha_1$ just as in standard principal-agent models. It remains to discuss how the principal uses w_{SS} and w_F to induce the agent to experiment with the new work method.

If $c_2/c_1 < \beta_2$, then exploitation is too costly for the agent. Only incentive compatibility constraints associated with shirking are binding. To prevent the agent from shirking in the first period and in the second period after a success in the first period, the principal pays the agent $w_{SS} = \alpha_2$.

If $c_2/c_1 \geq \beta_2$, then exploitation is not too costly for the agent. The incentive compatibility constraint associated with exploitation is binding. To prevent




Figure 2. The optimal contract that implements exploration under the base-case parameters. Under this contract, an agent who succeeds early and fails later has lower total compensation than an agent who fails early and succeeds later or even an agent who fails twice if c_2/c_1 is high.

exploitation, the principal can either reward the agent for failure in the first period or reward the agent for two consecutive successes. The principal’s choice between these two instruments depends on whether exploration is moderate or radical. With moderate exploration, it is cheaper for the principal to provide incentives through w_{SS} , since two consecutive successes are a stronger signal that the agent explored and not exploited than a failure in the first period. With radical exploration, it is cheaper for the principal to provide incentives through w_F , since a failure in the first period is a stronger signal that the agent explored and not exploited than two consecutive successes. Rewarding the agent for failure, however, induces the agent to shirk in the first period. To prevent shirking, delayed compensation w_{SS} must also be used.

Figure 2 shows the optimal contract that implements exploration for different values of c_2/c_1 under the base-case parameters. The optimal contract that implements exploration rewards long-term success but not short-term success. On the contrary, it may even reward short-term failure. This safety net is provided even though the agent is risk-neutral. The intuition is that if the agent is not protected against early failures, then the agent may prefer to exploit in order to avoid early failures.

An alternative way to interpret the optimal contract that implements exploration is to look at how it compensates different performance paths. The total compensation $w_F + w_{FS}$ when performance is FS is higher than the

total compensation $w_S + w_{SF}$ when performance is SF . An agent who recovers from failure has higher compensation than an agent who obtains short-lived success. Rewards are thus contingent on the performance path, and not only on the number of successes or failures obtained by the agent. If $w_F > 0$, then the total compensation $w_F + w_{FF}$ when performance is FF is higher than the total compensation $w_S + w_{SF}$ when performance is SF . Even an agent who fails twice may have higher compensation than an agent who obtains short-lived success. Because of the risky nature of exploration, failing twice may be a stronger signal for the principal that the agent explored and not exploited than obtaining a short-lived success.

Some principal-agent models assume that the agent can destroy output. In a static setting, this assumption implies that the optimal contract is nondecreasing. In a dynamic setting, however, this is not necessarily true. For example, in the model developed here, setting $p_0 = 0$ will have the same effect as allowing the agent to destroy output. Still, from Proposition 2, the optimal contract that implements exploration may have $w_S < w_F$. Under this contract, if the agent decides to destroy output at the end of the first period to obtain w_F , the agent forgoes the opportunity of earning w_{SS} in the second period.

V. Lack of Commitment

In contrast to the previous section, I now assume that the principal cannot commit to a long-term contract. In practice, commitment to a long-term contract may be achieved through explicit contracts, such as stock options or vesting stock, or through implicit contracts, based on reputation. Sometimes these options are not available, and the principal can only offer the agent a short-term contract specifying the agent's compensation contingent on the current period's performance. The problem becomes similar to the one proposed in Section III, except that there are additional constraints to guarantee that the principal is willing to keep the promised wages in the second period.

Fudenberg, Holmstrom, and Milgrom (1990) provide conditions under which a sequence of short-term contracts perform just as well as the optimal long-term contract. The model proposed here violates two of these conditions. First, there may not be common knowledge of technology. With learning through experimentation, the agent may be better informed than the principal about the technology in the second period, since first-period actions affect second-period expected probability of success. Second, the utility frontier may not be downward sloping, since the agent has limited liability.

PROPOSITION 3: *The optimal contract \bar{w}_1 that implements exploitation, derived in Proposition 1, can be realized through a sequence of short-term contracts.*

To implement exploitation, a sequence of short-term contracts performs just as well as the optimal long-term contract, because the optimal long-term contract that implements exploitation derived in Proposition 1 relies only on short-term incentives. Commitment is thus irrelevant to implement exploitation.

The following definition will be useful in stating Proposition 4:

$$\beta_4 = \frac{(E[p_2] - p_0)(1 + p_1)}{(p_1 - p_0) \left(1 + p_1 \frac{E[p_2] - p_0}{E[p_2|S, 2] - p_0} \right)}.$$

PROPOSITION 4: *The optimal contract \bar{w}_2 that implements exploration, derived in Proposition 2, cannot be replicated by a sequence of short-term contracts. Moreover, if*

- $c_2/c_1 \geq \beta_4$, then exploration is not implementable via short-term contracts
- $c_2/c_1 < \beta_4$, then the optimal sequence of short-term contracts \bar{w}_4 that implements exploration is such that

$$w_S = \frac{c_2}{E[p_2] - p_0} - p_0 w_{SS} + p_0 w_{FS},$$

$$w_F = w_{SF} = w_{FF} = 0,$$

$$w_{SS} = \frac{c_2}{E[p_2|S, 2] - p_0},$$

$$w_{FS} = \frac{c_1}{p_1 - p_0}.$$

Without commitment, the principal can only use short-term incentives to implement exploration. When $c_2/c_1 \geq \beta_4$, short-term incentives are not enough to induce exploration. If the principal rewards the agent for success in the first period, then the agent employs the conventional work method, which is relatively cheaper and yields a higher probability of success than the new work method. If, in contrast, the principal rewards the agent for failure in the first period, then the agent shirks, which is cheaper and yields a higher probability of failure than the new work method. Therefore, if $c_2/c_1 \geq \beta_4$, exploration cannot be implemented with a sequence of short-term contracts.⁸ When $c_2/c_1 < \beta_4$, short-term incentives may be enough to implement exploration. If the principal rewards the agent for success in the first period, it is too costly for the agent to employ the conventional work method. Exploration may thus be implementable with a sequence of short-term contracts. However, the cost $W(\bar{w}_4, \langle 2_1^2 \rangle)$ of implementing exploration with short-term contracts is higher than the cost $W(\bar{w}_2, \langle 2_1^2 \rangle)$ of implementing exploration with a long-term contract. When a long-term contract is used, the principal can wait until the second period to pay the agent, gathering more information about the agent's first-period action.

⁸ Hermalin and Katz (1991) show that an action is implementable if there does not exist a randomization over actions that induces the same density over outcomes and costs less to the agent. When $c_2/c_1 \geq \beta_4$, a randomization over actions 0 and 1 induces the same density over first-period outcomes as action 2 and costs less to the agent.




Figure 3. Cost of implementing exploration when the principal can (solid line) and cannot (dashed line) commit to a long-term contract under the base-case parameters. Even if it is feasible to motivate exploration with a sequence of short-term contracts (low c_2/c_1), it is less costly to implement exploration if the principal can commit to a long-term contract.

Figure 3 compares the cost of implementing exploration when the principal can and cannot commit to a long-term contract. Exploration is implementable via a sequence of short-term contracts only if c_2/c_1 is low. Even if this is the case, the cost $W(\bar{w}_4, \langle 2_1^2 \rangle)$ of implementing exploration with short-term contracts is higher than the cost $W(\bar{w}_2, \langle 2_1^2 \rangle)$ of implementing exploration with a long-term contract.

This section contrasts the effect of lack of commitment on the implementation of exploration and exploitation. To implement exploitation, a sequence of short-term contracts performs as well as the optimal long-term contract. On the other hand, to implement exploration, the optimal long-term contract performs better than any sequence of short-term contracts. For some parameters, it is even impossible to implement exploration with a sequence of short-term contracts. These results show the importance of commitment when implementing exploration.

VI. Termination

In this section, I allow the principal to terminate the agent after a failure in the first period. Termination can be interpreted, for example, as the decision to fire a manager or worker in a corporation, or the decision to interrupt funding of a start-up company. The principal may use termination as a screening device, firing the agent if it is not worthwhile to keep him in the second period. In addition to that, the principal may use termination as a disciplinary device to induce the agent to take the appropriate action in the first period.

I derive the optimal contracts that implement exploitation with termination and exploration with termination. I then study when it is optimal for the principal to implement exploitation with termination instead of exploitation, and exploration with termination instead of exploration. Exploitation with termination is represented by action plan $\langle 1_t^1 \rangle$, and exploration with termination




Figure 4. The optimal contracts that implement exploitation (solid line) and exploitation with termination (dashed line) under the base-case parameters. With termination, the agent loses the rents he was able to extract in the second period after a failure in the first period. The threat of termination thus induces the agent to exert more effort in the first period, and therefore the principal needs to pay less to the agent for early success in order to implement exploitation.

is represented by action plan $\langle 2_t^2 \rangle$, where t means that the principal terminates the agent after a failure in the first period. For simplicity, the agent’s outside wages after termination are zero.⁹ The principal’s expected revenues when implementing exploration with termination and exploitation with termination are given by $R(\langle 1_t^1 \rangle)$ and $R(\langle 2_t^2 \rangle)$, which may incorporate, for example, the possibility of hiring a replacement agent after termination.

For brevity, the results on the implementation of exploitation with termination are derived in the Internet Appendix. Figure 4 compares the optimal contracts that implement exploitation and exploitation with termination for different values of c_2/c_1 under the base-case parameters. If the agent shirks or employs the new work method, he is more likely to fail in the first period than if he employs the conventional work method. To avoid failure, and consequently termination, the agent has more incentives to employ the conventional work method in the first period. Therefore, the principal needs to pay the agent lower first-period wages to implement exploitation with termination than to implement exploitation.

This result is similar to Stiglitz and Weiss (1983), where the principal uses termination to induce the agent to exert effort. Because termination serves this additional role in providing incentives to the agent, it is optimal for the principal to use excessive termination when implementing exploitation.

⁹ In this context, implementing action plan $\langle i_t^j \rangle$ is the same as implementing action plan $\langle i_0^j \rangle$ with $w_{FF} = w_{FS} = 0$.

Proposition 5 derives the optimal contract that implements exploration with termination. The following definitions will be useful in stating Proposition 5:

$$\alpha_5 = \max_{j \in \{0,1\}} \frac{(1 + E[p_2])c_2 - p_0c_j}{E[p_2]E[p_2|S, 2] - p_0E[p_j]}$$

$$\beta_5 = \frac{(E[p_2]E[p_2|S, 2] - p_0p_1) + E[p_2](p_1E[p_2|S, 2] - p_0E[p_2|S, 2])}{(p_1^2 - p_0) + E[p_2](p_1^2 - p_0)}$$

PROPOSITION 5: *The optimal contract \bar{w}_5 that implements exploration with termination is such that*

$$w_S = w_{SF} = 0.$$

If exploration is moderate, then $w_F = 0$ and

$$w_{SS} = \alpha_5 + \frac{p_1 - p_0}{E[p_2]E[p_2|S, 2] - p_0p_1} \frac{p_1(1 + E[p_2])c_1}{E[p_2]E[p_2|S, 2] - p_1^2} \left(\frac{c_2}{c_1} - \beta_5\right)^+$$

If exploration is radical, then

$$w_F = \frac{p_1(1 + E[p_2])c_1}{E[p_2]E[p_2|S, 2] - p_1E[p_2]} \left(\frac{c_2}{c_1} - \beta_5\right)^+$$

and

$$w_{SS} = \alpha_5 + \frac{E[p_2] - p_0}{E[p_2]E[p_2|S, 2] - p_0p_1} \frac{p_1(1 + E[p_2])c_1}{E[p_2]E[p_2|S, 2] - p_1E[p_2]} \left(\frac{c_2}{c_1} - \beta_5\right)^+.$$

The effects of termination on the incentives for the agent to employ the new work method in the first period will depend on the relative costs c_2/c_1 of the new and conventional work methods. If $c_2/c_1 \geq \beta_5$, then the incentive compatibility constraint associated with exploitation with termination is binding. Termination makes it harder to provide incentives for the agent to employ the new work method in the first period, since to avoid failure and termination, the agent has more incentives to employ the conventional work method in the first period. If $c_2/c_1 < \beta_5$, then the incentive compatibility constraint associated with shirking is binding. Termination makes it easier to provide incentives for the agent to employ the new work method in the first period, since, to avoid failure and termination, the agent has less incentives to shirk in the first period.

Figure 5 compares the optimal contracts that implement exploration and exploration with termination for different values of c_2/c_1 under the base-case parameters. If c_2/c_1 is high, the principal pays higher wages w_F to implement exploration with termination than to implement exploration. If c_2/c_1 is low, the principal pays lower wages w_{SS} to implement exploration with termination than to implement exploration.

I now compare the total expected profits of the principal when he implements exploration with the total expected profits of the principal when he implements




Figure 5. The optimal contract that implements exploration (solid line) and exploration with termination (dashed line) under the base-case parameters. With termination, the agent loses the rents he was able to extract in the second period after a failure in the first period. To prevent the threat of termination from drawing the agent away from exploration in the first period, the principal needs to pay an additional bonus to the agent after early failure if c_2/c_1 is high.

exploration with termination. It is optimal for the principal to implement exploration with termination instead of exploration if

$$R(\langle 2_1^2 \rangle) - R(\langle 2_t^2 \rangle) < W(\bar{w}_2, \langle 2_1^2 \rangle) - W(\bar{w}_5, \langle 2_t^2 \rangle).$$

To keep the agent working in the second period after a failure in the first period, the expected payments from the principal to the agent are equal to $(1 - p_1)p_1\alpha_1$. It is thus ex post efficient for the principal to terminate the agent after a failure in the first period if

$$R(\langle 2_1^2 \rangle) - R(\langle 2_t^2 \rangle) < (1 - E[p_2])p_1\alpha_1. \tag{7}$$

When (7) holds, the benefits from inducing the agent to work in the second period after a failure in the first period are lower than the expected payments that the principal must make to the agent after a failure in the first period to keep the agent working in the second period.

DEFINITION 2. *There is excessive termination with exploration if*

$$W(\bar{w}_2, \langle 2_1^2 \rangle) - W(\bar{w}_5, \langle 2_t^2 \rangle) > (1 - E[p_2])p_1\alpha_1,$$

and there is excessive continuation with exploration if

$$W(\bar{w}_2, \langle 2_1^2 \rangle) - W(\bar{w}_5, \langle 2_i^2 \rangle) < (1 - E[p_2])p_1\alpha_1.$$

There is excessive termination with exploration if the actual threshold for termination is higher than the ex post efficient threshold for termination. There is excessive continuation with exploration if the actual threshold for termination is lower than the ex post efficient threshold for termination. Excessive continuation or termination may arise because the termination policy affects the incentives for the agent's first-period action.

Corollary 1 investigates the conditions under which there is excessive termination with exploration or excessive continuation with exploration. The following definitions will be useful in stating Corollary 1:

$$\kappa_m \equiv \frac{(p_1 - E[p_2])(E[p_2]E[p_2|S, 2] - p_1p_0)}{(p_1 - E[p_2])(E[p_2]E[p_2|S, 2] - p_1p_0) + (E[p_2] - p_0)(E[p_2]E[p_2|S, 2] - p_1^2)},$$

$$\kappa_e \equiv \frac{(1 - E[p_2])(E[p_2]E[p_2|S, 2] - p_0p_1)}{(1 - E[p_2])(E[p_2]E[p_2|S, 2] - p_0p_1) + E[p_2]E[p_2|S, 2](E[p_2] - p_0)}.$$

COROLLARY 1: *If $c_2/c_1 < \max(\kappa_m, \kappa_e)\beta_2 + (1 - \max(\kappa_m, \kappa_e))\beta_5$, then there is excessive termination with exploration. If $c_2/c_1 > \max(\kappa_m, \kappa_e)\beta_2 + (1 - \max(\kappa_m, \kappa_e))\beta_5$, then there is excessive continuation with exploration.*

As shown in Proposition 5, the effects of termination on the incentives for the agent to employ the new work method in the first period depend on c_2/c_1 . For low values of c_2/c_1 , the agent is inclined to shirk. The threat of termination allows the principal to pay the agent lower wages to prevent shirking, offsetting the losses from excessive termination. For high values of c_2/c_1 , the agent is inclined to exploit. Excessive continuation allows the principal to pay the agent lower wages in the first period, offsetting the losses from excessive continuation.

As shown in Corollary 1, there is excessive continuation with exploration even if exploration is moderate. This is in contrast to the results in Proposition 2, which say that there is reward for failure only if exploration is radical. With moderate exploration, the principal does not reward the agent for failure because it is cheaper to use rewards for long-term success to induce exploration. However, the surplus the agent obtains in the second period after a failure in the first period still provides incentives for the agent to explore when c_2 is high relative to c_1 .

This section contrasts the optimal termination policies when implementing exploitation and exploration. Similar to models of repeated effort, there is excessive termination when implementing exploitation. On the other hand, depending on which constraints are binding, there may be excessive termination or excessive continuation when implementing exploration. There is excessive termination if only shirking constraints are binding, while there is excessive continuation if the exploitation constraint is binding. To sum up, termination

is useful to prevent shirking but it may be harmful when implementing exploration as it may induce exploitation from the agent.

VII. Feedback

In this section, I study what happens if the principal is better able to evaluate performance than the agent. This could be relevant, for example, in studying the relation between a venture capitalist and an entrepreneur, in which the venture capitalist knows more about the commercial value of the enterprise than the entrepreneur. Also, firms often have better information about the potential market performance of products being developed by their R&D employees. The focus of the section will be on whether the principal should provide feedback on performance to the agent. I show that the optimal provision of feedback will depend on whether the principal wants to implement exploration or exploitation.

I assume that the principal privately observes interim performance at the end of the first period, yet the performance path is publicly observable at the end of the second period. If the principal does not reveal interim performance realizations, then only incentive compatibility constraints $IC_{(i_k^j)}$, where $j = k$ need to be satisfied, since without feedback the agent cannot adjust his action according to the realization of first-period performance. However, for the same reason, only action plans $\langle i_k^j \rangle$ with $j = k$ can be implemented without feedback. Therefore, if the action plan to be implemented involves repetitive actions, then it is optimal for the principal not to provide feedback on performance. On the other hand, if the action plan to be implemented requires adjustments in action depending on the realized interim performance, then feedback on performance must be provided.

For brevity, the detailed analysis of the provision of feedback when the principal implements exploitation is presented in the Internet Appendix. Because exploitation involves repetitive actions, it is optimal for the principal not to provide feedback on performance to the agent. This result is similar to Lizzeri, Meyer, and Persico (2002) and Fuchs (2007), who find that in a setting where the principal's problem is to induce the agent to exert repeated effort, it is optimal for the principal not to reveal information about performance to the agent.

In contrast, as shown by the next proposition, feedback is essential to motivate exploration, as it permits efficient experimentation.

PROPOSITION 6: *To implement exploration, the principal must provide feedback to the agent. The optimal contract \bar{w}_6 that implements exploration when the principal is better able to evaluate performance than the agent is the same as the optimal contract \bar{w}_2 that implements exploration derived in Proposition 2.*

The function of feedback here is to provide information that improves the agent's future performance. No punishment is associated with feedback. On the contrary, for some parameters, the agent is even rewarded in the case of

failure. If the agent is not protected against failures, then the agent is inclined to employ the conventional work method to avoid failures.¹⁰

This section contrasts the feedback policy when implementing exploitation and exploration. It shows that, similar to repeated-effort models, the principal should never provide feedback on performance to the agent when implementing exploitation, but should always provide timely feedback on performance to the agent when implementing exploration.

VIII. Empirical Implications

The main contribution of the paper is to contrast incentive schemes that motivate exploration and exploitation. Motivating exploitation requires standard pay-for-performance schemes, excessive termination, short-term contracts, and no feedback on performance. In contrast, motivating exploration involves tolerance (or even reward) for early failure and reward for long-term success, so that not only total performance but also the path of performance matters for compensation. Moreover, excessive continuation, commitment to a long-term incentive plan, and timely feedback on performance are also important to motivate exploration.

Two empirical tests arise naturally. First, one can study whether incentive practices used in tasks for which innovation is more (less) important resemble the incentive practices that, according to this paper, motivate exploration (exploitation). Second, one can study whether the adoption of incentive practices that motivate exploration (exploitation) leads to more (less) innovation. In the remainder of this section, I discuss the application of the results of the paper to different settings, providing more details about how they can be tested empirically.

A. Executive Compensation and Corporate Governance

Executive compensation has increasingly been criticized as excessive and not related to performance. This public outcry has created pressure for regulations that limit the use of stock options, golden parachutes, entrenchment, and option repricing.¹¹ We argue here that stock options, golden parachutes, entrenchment, and option repricing may be part of an optimal contract that motivates innovation. Therefore, passing regulations that restrict their use may in some cases have an adverse effect on innovation. To assess the actual impact of such regulations, it remains to be studied empirically the actual contribution of these practices to innovation.

From Propositions 2 and 5, it is easy to see that the optimal compensation that motivates exploration with and without termination can be implemented

¹⁰ Other research provides alternative rationale for the provision of feedback. Ederer (2010a) shows that feedback may be useful when the agent is uncertain about his ability. Ray (2007) develops a model in which interim performance evaluation serves the purpose of screening bad projects.

¹¹ See, for example, Bebchuk and Fried (2004) and "Rewards for Failure," British DTI consultation, June 2003.

via a combination of stock options with long vesting periods, option repricing, and golden parachutes. For simplicity, I assume that the firm does not pay dividends until the second period. By granting the manager stock options that vest and mature at the end of the second period and have a strike price of S , it is possible to commit to a payment of w_{SS} to the manager after two consecutive successes. If the manager fails in the first period, there are two situations to consider. When parameters of the model are such that the manager needs to be fired, he may leave the firm with a payment w_F , as shown in Proposition 5. This payment can be implemented by promising the manager a severance payment (e.g., in the form of a golden parachute) to be paid if the manager gets fired. When the manager is to stay in the firm, then the optimal contract can be implemented by repricing the original stock option. The option repricing needs to be done in a way that guarantees that the manager puts effort in the second period ($w_{FS} > c_1/(p_1 - p_0)$), and also that the manager gets an extra surplus (w_F) for early failures as shown in Proposition 2.¹²

Corollary 1 shows that, under the optimal contract that motivates exploration, shareholders may need to commit not to fire the manager even if it is ex-post efficient to do so. Managerial entrenchment can produce this desired excessive continuation, as it makes it harder for shareholders to fire the manager. The appointment of a board of directors that is friendly to the manager may lead to managerial entrenchment. Alternatively, dispersed ownership, typical in large public corporations, reduces the incentives for shareholders to intervene, effectively entrenching the manager.¹³

This naturally raises the question of why a lot of the innovation in the economy comes from firms financed by venture capital, which has concentrated ownership. Some of the practices adopted by venture capitalists may help motivate innovation.¹⁴ For example, excessive continuation may be achieved by delegating the decision to stop a project to the general partner. The bulk of the compensation of the general partner is in the form of carried interest, which is effectively a call option on the projects being financed. This provides incentives for the general partner to keep projects alive beyond the point under which it would be efficient to terminate them. Moreover, the entrepreneur in a start-up company financed by a venture capitalist typically does not earn rewards for early successes, and most of his compensation comes in the long term, when the company goes through an IPO or is sold to another firm. Venture capitalists are also known for providing detailed feedback on performance to entrepreneurs.

The theory developed here suggests that stock options with long vesting periods, entrenchment, and golden parachutes should be more often used in

¹² This extra surplus w_F can be paid in the first period, as in the optimal contract derived in Proposition 2. Alternatively, the extra surplus w_F can be paid only in the second period. Any such contract performs as well as the optimal long-term contract derived in Proposition 2 as long as after two periods an extra expected surplus w_F is paid to the manager if he fails in the first period.

¹³ Burkart, Gromb, and Panunzi (1997) and Myers (2000) develop models in which dispersed ownership serves as a way to reduce intervention by shareholders.

¹⁴ Gompers and Lerner (2004) describe in detail the practices used in the venture capital industry.

situations in which exploration and innovation are important. One potentially interesting direction of research is to study if these instruments are indeed more common in firms and industries involved in innovation. Francis et al. (2009) show that golden parachutes, as well as long-term incentives in the form of vested and unvested options, have a positive and significant effect on patents and citations to patents. Atanassov (2008) and Sapra, Subramanian, and Subramanian (2009) study the effects of corporate governance on innovation. Tian and Wang (2010) find that start-up firms financed by venture capitalists that are more tolerant of early failures are more innovative.

Previous studies such as Lambert (1986) and Feltham and Wu (2001) have developed static models in which the optimal compensation that encourages risk-taking is convex, resembling a stock option. Other studies derive optimal contracts that, for different reasons from that proposed here, involve golden parachutes, entrenchment, or option repricing. In a setting in which the manager observes a private signal about the future prospects of the firm, Inderst and Mueller (2010) and Diaz, Rayo, and Sapra (2009) show that stock options and golden parachutes may be optimal to induce the manager to reveal information to the board after bad outcomes. In an incomplete contracting framework, Almazan and Suarez (2003) show that a contract consisting of bonus and severance pay may be optimal to induce the incumbent manager to invest in firm-specific human capital when there is a threat that a better rival manager becomes available. In a setting where the only instruments available to the principal are at-the-money call options, Acharya, John, and Sundaram (2000) show that option repricing may be optimal because it motivates the agent to exert effort after poor performance.

B. Bankruptcy Laws and Entrepreneurship

Bankruptcy laws in Europe and in the United States have recently been under debate. On the one hand, to encourage entrepreneurial activity, in June of 2000 the European Council issued the “European Charter for Small Enterprises,” which states that

... failure is concomitant with responsible initiative and risk-taking and must be mainly envisaged as a learning opportunity.

The Charter declares that bankruptcy law reforms should become a clear priority for Member States and that new bankruptcy laws should allow failed entrepreneurs a fresh start. On the other hand, after 8 years of discussion, in April of 2005, the U.S. Congress passed a new creditor-friendly bankruptcy law, the “Bankruptcy Abuse and Consumer Protection Act,” which makes it more difficult for insolvent debtors to obtain exemptions and discharge of obligations.

The model developed in this paper sheds light on the incentive effects of different bankruptcy laws. If the entrepreneur borrows money to undertake some project and the project fails, then the entrepreneur will not have the funds to pay his debts and will be insolvent. From Propositions 1 and 2, the optimal contracts that motivate exploration and exploitation are quite different in the

way they treat insolvent debtors. The optimal contract that motivates exploration rewards the agent after failure. One can interpret this as a bankruptcy law based on the principle of a fresh start, as it provides the entrepreneur with generous exemptions and an immediate full discharge of debt, so that the entrepreneur keeps some surplus after failure (in a violation of the absolute priority of claims). By protecting the entrepreneur against early failure, these bankruptcy laws make the entrepreneur more inclined to explore. On the other hand, the optimal contract that implements exploitation does not reward the agent after failure. One can interpret this as a bankruptcy law based on the principle of absolute priority. The creditor seizes the goods owned by the entrepreneur and discharge takes several years.

A natural question to ask is why governments impose a single mandatory bankruptcy law instead of offering a menu of bankruptcy laws that contains the optimal law for different situations. By considering the incentives for exploration, this paper provides a potential explanation to this question. Due to knowledge spillovers and imperfect protection of intellectual property rights, individuals involved in exploratory activities cannot fully appropriate the economic value generated by the knowledge they produce. As Nelson (1959) argues, this leads to underexploration when compared to the socially efficient level of exploration. One way to alleviate the underexploration problem is by imposing a debtor-friendly bankruptcy law.

There is a large literature on the design of bankruptcy laws. Based on standard models of incentives, Jensen (1991) and Aghion, Hart, and Moore (1992) are strong proponents of bankruptcy laws that respect the absolute priority of claims. Other papers find beneficial effects of deviations from absolute priority. For example, Bebchuk and Picker (1993) and Berkovitch, Israel, and Zender (1997) show that deviations of absolute priority may encourage investments in firm-specific versus general human capital. Baird (1991) and Povel (1999) show that deviations of absolute priority induce the entrepreneur to reveal private information to creditors when bad outcomes occur.¹⁵ Ayotte (2007) shows that a mandatory debtor-friendly bankruptcy law may increase social welfare because it prevents the monopolist bank from extracting too much surplus from the entrepreneur. Acharya and Subramanian (2009) analyze the effect of bankruptcy laws on entrepreneurship using cross-sectional and time-series data of several countries, and find that debtor-friendly bankruptcy laws lead to more innovation.

C. Intrapreneurship in Large Corporations

Managers of large corporations often claim that it is hard to motivate their employees to be more creative.¹⁶ One potential explanation is the difficulty

¹⁵ Landier (2002) develops a model with multiple equilibria in which the stigma of failure may prevent entrepreneurs from abandoning bad projects.

¹⁶ "CEO Challenge 2004: Perspectives and Analysis," The Conference Board, Report 1353 (<http://www.conference-board.org/> and <http://www.ceochallenge.org/>).

large corporations face in credibly promising to reward employees for their discoveries and to tolerate early failures. In the case of executive compensation, companies can overcome this problem by using stock options with long vesting periods, option repricing, golden parachutes, and entrenchment. For lower-level employees, however, these types of contracts may not be available, since, for example, there may be no verifiable measures of the long-term performance of the employee.

To overcome these difficulties, business consultants have argued that nurturing a corporate culture that allows freedom to experiment and tolerates failures is essential to motivate innovation among employees of large corporations. Farson and Keyes (2002) and Sutton (2002) provide several examples of innovative corporations, such as IBM and 3M, that adopt such a culture. As shown in Proposition 4, the ability to commit to a long-term contract is essential to encourage exploration. Promises made in the form of a corporate culture can be enforced through reputation. From Proposition 2, a corporate culture that tolerates early failure and rewards long-term success is optimal to motivate exploration.

Innovative organizations may also rely on explicit long-term contracts to overcome the commitment problem and induce exploration. For example, research departments in business or academic organizations often grant tenure to their researchers. Knowing that they will not lose their jobs, researchers are willing to explore new research directions that are likely to fail but may lead to breakthroughs. Even before obtaining tenure, researchers in academic organizations are usually given a period of time under which they cannot be terminated. From Corollary 1, by committing not to terminate researchers, research departments are able to motivate exploration. Researchers are also often given explicit rewards for long-term success. Lerner and Wulf (2007) find that more long-term incentives to the heads of research and development departments are associated with more heavily cited patents, while short-term incentives are unrelated to measures of innovation.

The way a corporation organizes its internal capital markets may also have an impact on innovation. Seru (2010) finds evidence that high-level managers in conglomerates are more reluctant to invest in novel projects because of the threat of relocation of resources by headquarters in the case of failure. Sometimes a corporation is able to overcome this problem by creating a central fund that provides resources for experiments in different units, so that an experiment failure does not affect the budget situation of the division. Thomke (2002) discusses the case of Bank of America, which created a central fund to finance experiments and assigned 25 of its branches to be used as real-life laboratories where new products and service concepts were tested. Consistent with the predictions of the model developed here, the incentive scheme of the exploration team responsible for these branches is very different from the incentive schemes of the rest of the workers in Bank of America. Initially, the exploration team was assigned a target failure rate of 30%. In the first year of operation, the team had only 10% failure. As the leader of the exploration team explained, "We are trying to sell ourselves to the bank. If we have too many failures, we

just won't be accepted" (p. 9). To make it clearer that failures are welcome, the top executives of Bank of America were contemplating an increase in the target failure rate from 30% to 40%.

IX. Additional Discussion

In the preceding sections of the paper, I assume that the agent is risk-neutral and has limited liability. Results similar to those above hold if the agent is risk-averse. The critical elements influencing the optimal contracts are the likelihood ratios between the different action plans, not the agent's preferences. If the agent is risk-neutral, then the problem of finding the optimal contract that implements a given action plan simplifies to a linear programming problem. This allows me to focus on incentive issues rather than on risk-sharing issues.

For tractability, I restrict the analysis to a model with two periods and two possible outcomes in each period. Having more periods can strengthen the results obtained here. As I discuss in Section II, if the agent lives for multiple periods, the agent is willing to sacrifice even more output in the first period by employing the new work method, since the information learned in the first period can be used for a longer period of time. On the other hand, having multiple possible outcomes in each period may change some of the results. For example, if the new work method can produce a big success in the first period, then it is possible that the optimal contract that implements exploration rewards the big success in the first period. Two considerations justify the restriction to two possible outcomes. First, most of the studies on innovation point to a high rate of failure in innovative projects as the fundamental difference between innovative and traditional projects. If this is indeed the case, then even with multiple possible outcomes in each period, the principal will still rely on reward for early failures as it is the cheapest way to distinguish exploration from exploitation. Second, if the agent is risk-averse, then both rewards for failure and rewards for big successes will be used. Since the probability of a big success in the first period when the agent employs the new work method is usually very low, in practice we will see rewards for failure more often than rewards for big successes.

Lazear (1986) makes the distinction between input-based pay, where the principal compensates the agent based on the action taken by the agent, and output-based pay, where the principal compensates the agent based on the output produced by the agent. I assume in this paper that the principal observes only the output produced by the agent, and consequently can only use output-based pay. Prendergast (2002) argues that, in uncertain environments, such as the innovation environment studied here, it is difficult for the principal to evaluate the different projects available to the agent and therefore the principal delegates responsibilities to the agent, which in turn leads to output-based pay. Even in such uncertain environments, however, a noisy signal about the action taken by the agent may be observable by the principal. One can show that if such a signal is available for contracting, then the principal uses both

input-based pay and output-based pay to compensate the agent. As the signal about the actions taken by the agent becomes more precise, the principal relies more on input-based pay, but still relies on output-based pay in the form studied in this paper. It is only in the extreme case in which the principal perfectly observes the actions taken by the agent that the principal does not rely on output-based pay.

In the model, the agent chooses between a conventional technology, which has known probability of success, and a new technology, which has unknown probability of success. In a strict interpretation, the model seems to apply better to mature firms (with an existing business to exploit). However, if one interprets the agent's choice as a choice between more innovative and less innovative strategies, then it is easy to see that the model can also be applied to start-up companies, since entrepreneurs in those companies often face this type of decision.

The paper analyzes the problem of motivating innovation as an individual incentive problem. It is common that an individual, such as a manager in an organization, has to choose between more innovative or less innovative projects. Moreover, there are often performance measures associated with the outcome of this choice that can be used to compensate this individual. Therefore, modeling the problem as an individual incentive problem seems reasonable and produces a rich set of predictions. Since innovation is often produced by teams of individuals working together on a problem, it is interesting to consider the team incentive problem. Ederer (2010b) extends the model in this paper to allow for multiple agents and finds some new results that arise from the strategic interaction between members of the innovation team.

X. Conclusion

This paper proposes a framework to study the incentives for innovation. In this framework, innovation is the result of learning through the exploration of untested approaches that are likely to fail. As a result, the optimal incentive scheme that motivates exploration is fundamentally different from standard pay-for-performance schemes used to motivate effort. Tolerance (or even reward) for early failure, reward for long-term success, excessive continuation, commitment to a long-term incentive plan, and timely feedback on performance are all important to motivate exploration.

Practices such as golden parachutes, managerial entrenchment, and debtor-friendly bankruptcy laws protect or even reward the agent when failure occurs. These practices are often criticized, because, by protecting or rewarding the agent after poor performance, they undermine the incentive for the agent to exert effort. This paper shows that these practices may arise as part of an optimal incentive scheme that motivates exploration. Therefore, regulations that limit their use may, in some cases, have an adverse effect on innovation. To assess the actual impact of such regulations, it remains to be studied empirically the actual contribution of these practices to innovation.

There are several potentially interesting extensions of the theoretical model proposed here. For example, if the agent has superior information about his own type, then contracts may be used to sort agents. This raises the question of how to design contracts that attract creative workers while avoiding conventional workers and shirkers. Answers to this problem could be relevant, for example, to firms trying to hire a turnaround manager or trying to attract a more creative workforce. Another interesting question is the effect of public versus private ownership on innovation. Earnings in public companies are transparent to the market, which may put pressure on the manager to meet short-term earnings expectations, and potentially reduce incentives for innovation. I leave these questions for future research.

Empirical research mentioned in the paper provides support to some of the results derived here. Some of the predictions of the model remain untested though, and additional empirical work seems warranted. For example, it would be interesting to investigate if the combination of stock options with long vesting periods, option repricing, golden parachutes, and managerial entrenchment is more prevalent in firms for which motivating innovation is important. It would also be interesting to study whether more feedback is provided when the goal is to motivate exploration. The venture capital industry may be a natural place to test this hypothesis as venture capitalists are known to use their expertise to provide feedback to entrepreneurs.

Appendix: Proofs

The following definitions will be useful in stating the incentive compatibility constraints:

$$V_S(\bar{w}, \langle i_k^j \rangle) = w_S + E[p_j|S, i]w_{SS} + (1 - E[p_j|S, i])w_{SF},$$

$$V_F(\bar{w}, \langle i_k^j \rangle) = w_F + E[p_k|F, i]w_{FS} + (1 - E[p_k|F, i])w_{FF}.$$

Proof of Proposition 1: The optimal contract \bar{w} that implements action plan $\langle 1_1^1 \rangle$ satisfies the following incentive compatibility constraints:

$$\begin{aligned} & (p_1 - E[p_i])(V_S(\bar{w}, \langle 1_1^1 \rangle) - V_F(\bar{w}, \langle 1_1^1 \rangle)) \\ & + E[p_i](p_1 - E[p_j|S, i])(w_{SS} - w_{SF}) \\ & + (1 - E[p_i])(p_1 - E[p_k|F, i])(w_{FS} - w_{FF}) \geq \\ & (c_1 + p_1c_1 + (1 - p_1)c_1) - (c_i + E[p_i]c_j + (1 - E[p_i])c_k). \quad (\text{IC}_{\langle i_k^j \rangle}) \end{aligned}$$

First, I show that $w_F = w_{FF} = w_{SF} = 0$. Suppose $w_F > 0$ or $w_{FF} > 0$. A contract \bar{w}' that is the same as \bar{w} but has $w_{F'} = 0$ and $w_{FF'} = 0$ satisfies all $\text{IC}_{\langle i_k^j \rangle}$ and has $W(\bar{w}', \langle 1_1^1 \rangle) < W(\bar{w}, \langle 1_1^1 \rangle)$. Suppose now that $w_{SF} > 0$. Let the contract \bar{w}' be the same as \bar{w} except that $w_{SF'} = 0$, $w_{SS'} = w_{SS} - w_{SF}$ and $w_{S'} = w_S + w_{SF}$. The contract \bar{w}' satisfies all $\text{IC}_{\langle i_k^j \rangle}$, $W(\bar{w}', \langle 1_1^1 \rangle) = W(\bar{w}, \langle 1_1^1 \rangle)$, but \bar{w}' pays the agent earlier than \bar{w} .

I now argue that some incentive compatibility constraints are redundant. If $(i, j) \neq (1, 1)$, then it follows from $IC_{(1^1_0)}$ and $IC_{(i^j)}$ that $IC_{(i^j_0)}$ are redundant. If $(i, k) \neq (1, 1)$, then it follows from $IC_{(1^0)}$ and $IC_{(i^j_k)}$ that $IC_{(i^j_k)}$ are redundant. If $(i^j_k) \neq (2^2_1)$ and either $i = 2, j = 2$, or $k = 2$, then it follows from $c_2/c_1 \geq (E[p_2] - p_0)/(p_1 - p_0)$ that $IC_{(i^j_k)}$ is redundant. Rewriting the incentive compatibility constraints that are not redundant, we have:

$$(p_1 - p_0)w_{SS} \geq c_1, \quad (IC_{(1^0)})$$

$$(p_1 - p_0)w_{FS} \geq c_1, \quad (IC_{(1^1_0)})$$

$$(p_1 - p_0)w_S + (p_1^2 - p_0p_1)w_{SS} - (p_1^2 - p_0p_1)w_{FS} \geq c_1, \quad (IC_{(0^1_1)})$$

and

$$(p_1 - E[p_2])w_S + (p_1^2 - E[p_2]E[p_2|S, 2])w_{SS} - (p_1^2 - E[p_2]p_1)w_{FS} \geq c_1 - c_2 + E[p_2](c_1 - c_2). \quad (IC_{(2^2_1)})$$

I now show that $IC_{(1^0)}$ and $IC_{(1^1_0)}$ are binding. If that is not the case, then either

$$\Delta_1 \equiv w_{SS} - \frac{c_1}{p_1 - p_0} > 0$$

or

$$\Delta_2 \equiv w_{SF} - \frac{c_1}{p_1 - p_0} > 0.$$

Let \bar{w}' be the same as \bar{w} except that $w'_{SS} = w_{SS} - \Delta_1$, $w_{S'} = w_S + p_1\Delta_1$, $w'_{FS} = w_{FS} - \Delta_2$, and $w_{F'} = w_F + (1 - p_1)\Delta_2$. The contract \bar{w}' satisfies the above constraints, $W(\bar{w}', \langle 1^1_1 \rangle) = W(\bar{w}, \langle 1^1_1 \rangle)$, and \bar{w}' pays the agent earlier than \bar{w} . The incentive compatibility constraints $IC_{(2^2_1)}$ and $IC_{(0^1_1)}$ then become

$$(p_1 - p_0)w_S \geq c_1 \quad (IC_{(0^1_1)})$$

$$(p_1 - E[p_2])w_S + E[p_2](E[p_2|S, 2] - p_1)\frac{c_1}{p_1 - p_0} \geq c_1 - c_2 + E[p_2](c_1 - c_2). \quad (IC_{(2^2_1)})$$

If $c_2/c_1 \geq \beta_1$, then $IC_{(0^1_1)}$ is binding. Otherwise, $IC_{(2^2_1)}$ is binding.

Proof of Proposition 2: The optimal contract \bar{w} that implements action plan $\langle 2^2_1 \rangle$ satisfies the following incentive compatibility constraints:

$$\begin{aligned} & (E[p_2] - E[p_i])(V_S(\bar{w}, \langle 2^2_1 \rangle) - V_F(\bar{w}, \langle 2^2_1 \rangle)) \\ & + E[p_i](E[p_2|S, 2] - E[p_j|S, i])(w_{SS} - w_{SF}) \\ & + (1 - E[p_i])(p_1 - E[p_k|F, i])(w_{FS} - w_{FF}) \geq \\ & (c_2 + E[p_2]c_2 + (1 - E[p_2])c_1) - (c_i + E[p_i]c_j + (1 - E[p_i])c_k). \quad (IC_{(i^j_k)}) \end{aligned}$$

First, I show that $w_S = w_{SF} = w_{FF} = 0$. Suppose $w_S > 0$. Let \vec{w}' be the same as \vec{w} except that $w'_S = 0$, $w'_{SS} = w_{SS} + \frac{w_S}{E[p_2|S,2]} - \epsilon$. There exists an $\epsilon > 0$ such that the contract \vec{w}' satisfies all $IC_{(i'_k)}$ and $W(\vec{w}', \langle 2_1^2 \rangle) < W(\vec{w}, \langle 2_1^2 \rangle)$. Now suppose $w_{SF} > 0$. Let the contract \vec{w}' be the same as \vec{w} except that $w'_{SF} = 0$ and $w'_{SS} = w_{SS} + \frac{1-E[p_2|S,2]}{E[p_2|S,2]}w_{SF} - \epsilon$. There exists an $\epsilon > 0$ such that the contract \vec{w}' satisfies all $IC_{(i'_k)}$ and $W(\vec{w}', \langle 2_1^2 \rangle) < W(\vec{w}, \langle 2_1^2 \rangle)$. Finally, suppose $w_{FF} > 0$. If the contract \vec{w}' is the same as \vec{w} , except that $w'_{FF} = 0$, and $w'_F = w_F + (1 - p_1)w_{FF}$, then all $IC_{(i'_k)}$ are still satisfied, $W(\vec{w}', \langle 2_1^2 \rangle) = W(\vec{w}, \langle 2_1^2 \rangle)$, and the contract \vec{w}' pays the agent earlier than \vec{w} .

It follows from $IC_{(2_0^2)}$ and $IC_{(i'_1)}$ that $IC_{(i'_0)}$ and $IC_{(i'_2)}$ are redundant. From $IC_{(2_0^2)}$, we have that $w_{FS} \geq \frac{c_1}{p_1-p_0}$ and $IC_{(i'_1)}$ implies $IC_{(i'_0)}$. Since $c_2 \geq \frac{E[p_2]-p_0}{p_1-p_0}c_1$, $IC_{(i'_1)}$ implies $IC_{(i'_2)}$.

Rewriting the incentive compatibility constraints that are not redundant, we have:

$$(p_1 - p_0)(w_{FS} - w_{FF}) \geq c_1, \tag{IC_{(2_0^2)}}$$

$$\begin{aligned} & (E[p_2]E[p_2|S, 2] - p_1E[p_j])w_{SS} + (p_1 - E[p_2])w_F \\ & + ((1 - E[p_2])p_1 - (1 - p_1)p_0)w_{FS} \\ & \geq (c_2 + E[p_2]c_2 + (1 - E[p_2])c_1) - (c_1 + p_1c_j), \end{aligned} \tag{IC_{(1_1^1)}}$$

$$\begin{aligned} & (E[p_2]E[p_2|S, 2] - p_0E[p_j])w_{SS} - (E[p_2] - p_0)w_F \\ & + ((1 - E[p_2])p_1 - (1 - p_0)p_0)w_{FS} \\ & \geq (c_2 + E[p_2]c_2 + (1 - E[p_2])c_1) - p_0c_j, \end{aligned} \tag{IC_{(0_1^1)}}$$

and

$$(E[p_2|S, 2] - E[p_j])w_{SS} \geq c_2 - c_j. \tag{IC_{(2_1^1)}}$$

The incentive compatibility constraint $IC_{(2_0^2)}$ is binding and $w_{FS} = \frac{c_1}{p_1-p_0}$. Suppose $w_{FS} > \frac{c_1}{p_1-p_0}$. If the contract \vec{w}' is the same as \vec{w} , except that $w'_{FS} = \frac{c_1}{p_1-p_0}$, and $w'_F = w_F + (1 - p_1)(w_{FS} - w'_{FS})$, then all $IC_{(i'_j)}$ are still satisfied, $W(\vec{w}', \langle 2_1^2 \rangle) = W(\vec{w}, \langle 2_1^2 \rangle)$, and the contract \vec{w}' pays the agent earlier than \vec{w} . On the other hand, the incentive compatibility constraints $IC_{(1_1^2)}$, $IC_{(1_1^0)}$, $IC_{(2_1^1)}$, and $IC_{(2_0^1)}$ are redundant. If $c_2 \geq c_1$, $IC_{(1_1^1)}$ implies $IC_{(1_1^2)}$, and if $c_2 < c_1$, $IC_{(1_1^2)}$ is trivially satisfied. Also, $IC_{(0_1^1)}$ and $IC_{(0_2^1)}$ imply $IC_{(1_1^0)}$. Moreover, $IC_{(0_2^1)}$ implies $IC_{(2_0^1)}$. Finally, $IC_{(1_1^1)} + \frac{p_1-E[p_2]}{E[p_2]-p_0}IC_{(0_1^1)}$ implies $IC_{(2_1^1)}$.

If $c_2/c_1 \geq \beta_2$, then one can show that $w_{SS} \geq \frac{c_1}{p_1-p_0} \geq \frac{c_1-c_2}{p_1-E[p_2]}$. Therefore, $IC_{(0_1^1)}$ implies $IC_{(0_1^0)}$ and $IC_{(0_2^0)}$. Either $w_F > 0$ and $IC_{(1_1^1)}$ and $IC_{(0_1^1)}$ are binding, or $w_F = 0$ and $IC_{(1_1^1)}$ is binding. When $IC_{(1_1^1)}$ and $IC_{(0_1^1)}$ are binding, the contract

is always feasible. Comparing the promised wages in each of the two possible contracts one can show that when

$$\frac{1 - E[p_2]}{1 - p_1} \geq \frac{E[p_2]E[p_2|S, 2]}{p_1^2},$$

the former contract is less costly for the principal than the latter contract. Otherwise, the latter contract is less costly.

If $c_2/c_1 < \beta_2$, then the candidate for the optimal contract is such that $IC_{(0_1^j)}$ and $IC_{(2_0^j)}$ are binding, $w_{SS} = w_{SS}^{0j1}$, and $w_F = 0$, where

$$j \in \arg \max_{j \in \{0,1\}} w_{SS}^{0j1} \equiv \frac{(1 + E[p_2])c_2 - p_0c_j}{(E[p_2]E[p_2|S, 2] - p_0E[p_j])} + \frac{(E[p_2] - p_0)p_0 \frac{c_1}{p_1 - p_0}}{(E[p_2]E[p_2|S, 2] - p_0E[p_j])}.$$

I first prove that the candidate contract is feasible. For that it suffices to show that $IC_{(1_1^j)}$ is satisfied. If $E[p_2]E[p_2|S, 2] \geq p_1^2$, then $IC_{(0_1^j)}$ implies $IC_{(1_1^j)}$. If $E[p_2]E[p_2|S, 2] < p_1^2$,

$$\begin{aligned} w_{SS}^{0j1} &< \frac{(1 + E[p_2])\beta_2c_1 - p_0c_1}{(E[p_2]E[p_2|S, 2] - p_0p_1)} + \frac{(E[p_2] - p_0)p_0 \frac{c_1}{p_1 - p_0}}{(E[p_2]E[p_2|S, 2] - p_0p_1)} \\ &= \frac{(1 + E[p_2])\beta_2c_1 - (1 + p_1)c_1}{(E[p_2]E[p_2|S, 2] - p_1^2)} - \frac{(p_1 - E[p_2])p_0 \frac{c_1}{p_1 - p_0}}{(E[p_2]E[p_2|S, 2] - p_0p_1)} \\ &< \frac{(1 + E[p_2])c_2 - (1 + p_1)c_1}{(E[p_2]E[p_2|S, 2] - p_1^2)} - \frac{(p_1 - E[p_2])p_0 \frac{c_1}{p_1 - p_0}}{(E[p_2]E[p_2|S, 2] - p_0p_1)}. \end{aligned}$$

In addition, $IC_{(0_1^j)}$ is not satisfied for any $w_{SS} < w_{SS}^{0j1}$. Therefore, it is impossible to improve on the candidate contract. ■

Proof of Proposition 3: The proof of the proposition follows from the fact that $IC_{(1_1^0)}$ and $IC_{(1_1^1)}$ are binding under the optimal long-term contract. ■

Proof of Proposition 4: To implement (2_1^2) , the following incentive compatibility constraints must be satisfied:

$$\begin{aligned} &(E[p_2] - E[p_i])(V_S(\bar{w}, \langle 2_1^2 \rangle) - V_F(\bar{w}, \langle 2_1^2 \rangle)) \\ &\quad + E[p_i](E[p_2|S, 2] - E[p_j|S, i])(w_{SS} - w_{SF}) \\ &\quad + (1 - E[p_i])(p_1 - E[p_k|F, i])(w_{FS} - w_{FF}) \\ &\geq (c_2 + E[p_2]c_2 + (1 - E[p_2])c_1) - (c_i + E[p_i]c_j + (1 - E[p_i])c_k). \quad (IC_{(i_k^j)}) \end{aligned}$$

Moreover, for the contract to be renegotiation-proof, we must have $j, k \in \mathcal{I}$ such that $IC_{(2_1^j)}$ and $IC_{(2_k^2)}$ bind.

If $c_2 \geq \frac{E[p_2|S, 2] - p_0}{p_1 - p_0} c_1$, from $IC_{(2_1^1)}$ we have that

$$w_{SS} = \frac{c_2 - c_1}{E[p_2|S, 2] - p_1},$$

$$w_{SF} = 0.$$

This contradicts $IC_{(1_1^1)} + \frac{p_1 - E[p_2]}{E[p_2] - p_0} IC_{(0_1^1)}$. Therefore, (2_1^2) is not implementable with a sequence of short-term contracts if $c_2 \geq \frac{E[p_2|S, 2] - p_0}{p_1 - p_0} c_1$. If $c_2 < \frac{E[p_2|S, 2] - p_0}{p_1 - p_0} c_1$, from $IC_{(2_1^0)}$ we have that

$$w_{SS} = \frac{c_2}{E[p_2|S, 2] - p_0},$$

$$w_{SF} = 0.$$

From $IC_{(2_k^2)}$, $k \in \{0, 2\}$ we have that

$$w_{FS} = \frac{c_1}{p_1 - p_0},$$

$$w_{FF} = 0.$$

Using the above equations, we can rewrite the following incentive compatibility constraints:

$$w_S \geq \frac{c_2(1 - p_0)}{E[p_2] - p_0} + p_0 \frac{c_1}{p_1 - p_0}, \tag{IC_{(0_1^2)}}$$

$$\begin{aligned} w_S &\geq \frac{c_2(E[p_2|S, 2] - p_0(1 - (p_1 - E[p_2])))}{(E[p_2|S, 2] - p_0)(E[p_2] - p_0)} \\ &\quad - \frac{c_1 p_0 (E[p_2|S, 2] - p_0)}{(E[p_2|S, 2] - p_0)(E[p_2] - p_0)} + p_0 \frac{c_1}{p_1 - p_0}, \end{aligned} \tag{IC_{(0_1^1)}}$$

and

$$w_S \geq \frac{c_2(E[p_2|S, 2] - p_0(1 + E[p_2]) + p_0^2)}{(E[p_2|S, 2] - p_0)(E[p_2] - p_0)} + p_0 \frac{c_1}{p_1 - p_0}. \tag{IC_{(0_k^0)}}$$

It is easy to show that given $c_2 < \frac{E[p_2|S, 2] - p_0}{p_1 - p_0} c_1$, $IC_{(0_1^0)}$ implies $IC_{(0_1^1)}$ and $IC_{(0_1^2)}$. Therefore, from $IC_{(0_1^0)}$, our candidate for w_S is

$$w_S = \frac{c_2}{E[p_2] - p_0} - \frac{p_0 c_2}{E[p_2|S, 2] - p_0} + p_0 \frac{c_1}{p_1 - p_0}.$$

It can be shown that the candidate contract satisfies all other incentive compatibility constraints if and only if

$$c_2 < \frac{(E[p_2|S, 2] - p_0)(1 + p_1)}{(p_1 - p_0) \left(\frac{E[p_2|S, 2] - p_0}{E[p_2] - p_0} + p_1 \right)} c_1.$$

In this case, the sequence of short-term contracts derived above is the optimal sequence of short-term contracts. ■

Proof of Proposition 5: The proof is similar to the proof of Proposition 2. ■

Proof of Corollary 1: The proof follows from comparing the costs of implementing exploration and exploration with termination from the contracts derived in Propositions 2 and 5. If $c_2/c_1 > \max(\kappa_m, \kappa_e)\beta_2 + (1 - \max(\kappa_m, \kappa_e))\beta_5$, then $W(\bar{w}_2, \langle 2_1^2 \rangle) - W(\bar{w}_5, \langle 2_t^2 \rangle) > (1 - E[p_2])p_1\alpha_2$ and there is inefficient continuation with exploration. If $c_2/c_1 < \max(\kappa_m, \kappa_e)\beta_2 + (1 - \max(\kappa_m, \kappa_e))\beta_5$, then $W(\bar{w}_2, \langle 2_1^2 \rangle) - W(\bar{w}_5, \langle 2_t^2 \rangle) < (1 - E[p_2])p_1\alpha_2$ and there is inefficient termination with exploration. ■

Proof of Proposition 6: Action plan $\langle 2_1^2 \rangle$ can only be implemented if the principal provides feedback on performance to the agent. ■

REFERENCES

- Acemoglu, Daron, Simon Johnson, and James Robinson, 2002, Reversal of fortune: Geography and institutions in the making of the modern world income distribution, *Quarterly Journal of Economics* 117, 1231–1294.
- Acharya, Viral, Ramin Baghai-Wadji, and Krishnamurthy Subramanian, 2009, Labor laws and innovation, Working paper, London Business School.
- Acharya, Viral, Kose John, and Rangarajan Sundaram, 2000, On the optimality of resetting executive stock options *Journal of Financial Economics* 57, 65–101.
- Acharya, Viral, and Krishnamurthy Subramanian, 2009, Bankruptcy codes and innovation, *Review of Financial Studies* 22, 4949–4988.
- Aghion, Philippe, 2002, Schumpeterian growth theory and the dynamics of income inequality, *Econometrica* 70, 855–882.
- Aghion, Philippe, Oliver Hart, and John Moore, 1992, The economics of bankruptcy reform, *Journal of Law, Economics, and Organizations* 8, 523–546.
- Aghion, Philippe, and Peter Howitt, 1992, A model of growth through creative destruction, *Econometrica* 60, 323–352.
- Aghion, Philippe, and Jean Tirole, 1994, The management of innovation, *Quarterly Journal of Economics* 109, 1185–1209.
- Almazan, Andres, and Javier Suarez, 2003, Entrenchment and severance pay in optimal governance structures, *Journal of Finance* 58, 519–548.
- Arrow, Kenneth, 1969, Classificatory notes on the production and diffusion of knowledge, *American Economic Review* 59, 29–35.
- Atanassov, Julian, 2008, Quiet life or managerial myopia: Is the threat of hostile takeovers beneficial for technological innovation? Working paper, University of Oregon.
- Ayotte, Kenneth, 2007, Bankruptcy and entrepreneurship: The value of a fresh start, *Journal of Law, Economics and Organization* 23, 161–185.
- Azoulay, Pierre, Joshua Graff Zivin, and Gustavo Manso, 2010, Incentives and creativity: Evidence from the academic life sciences, Working paper, MIT Sloan School of Management.

- Baird, Douglas, 1991, The initiation problem in bankruptcy, *International Review of Law and Economics* 11, 223–232.
- Battacharya, Sudipto, Kalyan Chatterjee, and Larry Samuelson, 1986, Sequential research and the adoption of innovations, *Oxford Economic Papers* 38, 219–243.
- Bebchuk, Lucian, and Jesse Fried, 2004, *Pay Without Performance* (Harvard University Press, Cambridge, MA).
- Bebchuk, Lucian, and Randal Picker, 1993, Bankruptcy rules, managerial entrenchment, and firm-specific human capital, Working paper, Law & Economics Working Paper No. 16, The University of Chicago Law School.
- Bekaert, Geert, Campbell Harvey, and Christian Lundblad, 2005, Does financial liberalization spur growth? *Journal of Financial Economics* 77, 3–55.
- Bergemann, Dirk, and Ulrich Hege, 2005, The financing of innovation: Learning and stopping, *Rand Journal of Economics* 36, 719–752.
- Bergemann, Dirk, and Juuso Valimaki, 2006, Bandit problems, in S. Durlauf, and L. Blume, eds. *The New Palgrave Dictionary of Economics* (Macmillan Press, Basingstoke, UK).
- Berkovitch, Elazar, Ronen Israel, and Jaime Zender, 1997, Optimal bankruptcy law and firm-specific investments, *European Economic Review* 41, 487–497.
- Berle, Adolph, and Gardiner Means, 1932, *The Modern Corporation and Private Property* (Macmillan, New York, NY).
- Berry, Donald, and Bert Fristedt, 1985, *Bandit Problems: Sequential Allocation of Experiments* (Chapman and Hall, New York, NY).
- Burkart, Mike, Denis Gromb, and Fausto Panunzi, 1997, Large shareholders, monitoring, and the value of the firm, *Quarterly Journal of Economics* 112, 693–728.
- Bushee, Brian, 1998, The influence of institutional investors on myopic R&D investment behavior, *Accounting Review* 73, 305–333.
- Dechow, Patricia, and Richard Sloan, 1991, Executive incentives and the horizon problem, *Journal of Accounting and Economics* 14, 51–89.
- Diaz, Miguel, Luis Rayo, and Haresh Sapra, 2009, Market pressure, control rights, and innovation, Working paper, University of Chicago Graduate School of Business.
- Ederer, Florian, 2010a, Feedback and motivation in dynamic tournaments, *Journal of Economics & Management Strategy* 19, 733–769.
- Ederer, Florian, 2010b, Launching a thousand ships: Incentives for parallel innovation, Working paper, UCLA.
- Ederer, Florian, and Gustavo Manso, 2010, Is pay-for-performance detrimental to innovation? Working paper, MIT.
- Farson, Richard, and Ralph Keyes, 2002, *Whoever Makes the Most Mistakes Wins: The Paradox of Innovation* (The Free Press, New York, NY).
- Feltham, Gerald, and Martin Wu, 2001, Incentive efficiency of stock versus options, *Review of Accounting Studies* 6, 7–28.
- Francis, Bill, Iftekhar Hasan, and Zenu Sharma, 2009, Do incentives create innovation? Evidence from CEO compensation contracts, Working paper, Rensselaer Polytechnic Institute.
- Fuchs, William, 2007, Contracting with repeated moral hazard and private evaluations, *American Economic Review* 73, 1432–1448.
- Fudenberg, Drew, Bengt Holmstrom, and Paul Milgrom, 1990, Short-term contracts and long-term agency relationships, *Journal of Economic Theory* 51, 1–31.
- Gompers, Paul, and Josh Lerner, 2004, *The Venture Capital Cycle* (MIT Press, Cambridge, MA).
- Graham, John, Campbell Harvey, and Shiva Rajgopal, 2005, The economic implications of corporate financial reporting, *Journal of Accounting and Economics* 40, 3–73.
- Harris, Milton, and Artur Raviv, 1978, Some results on incentive contracts with applications to education, insurance, and law enforcement, *American Economic Review* 68, 20–30.
- Hellmann, Thomas, 2007, When do employees become entrepreneurs? *Management Science* 53, 919–933.
- Hellmann, Thomas, and Veikko Thiele, 2008, Incentives and innovation inside firms: A multi-tasking approach, Working paper, University of British Columbia.

- Hermalin, Benjamin, and Michael Katz, 1991, Moral hazard and verifiability: The effects of renegotiation in agency, *Econometrica* 59, 1735–1753.
- Holmstrom, Bengt, 1979, Moral hazard and observability, *Bell Journal of Economics* 10, 74–91.
- Holmstrom, Bengt, 1989, Agency costs and innovation, *Journal of Economic Behavior and Organization* 12, 305–327.
- Holmstrom, Bengt, and Paul Milgrom, 1991, Multi-task principal-agent analysis, *Journal of Law, Economics, and Organizations* 7, 24–52.
- Inderst, Roman, and Holger Mueller, 2010, CEO replacement under private information, *Review of Financial Studies* 23, 2935–2969.
- Jensen, Michael, 1991, Corporate control and the politics of finance, *Journal of Applied Corporate Finance* 4, 13–33.
- Jensen, Richard, 1982, Adoption and diffusion of an innovation of uncertain probability, *Journal of Economic Theory* 27, 182–193.
- Jovanovic, Boyan, and Yaw Nyarko, 1996, Learning by doing and the choice of technology, *Econometrica* 64, 1299–1310.
- Jovanovic, Boyan, and Rafael Rob, 1990, Long waves and short waves: Growth through intensive and extensive search, *Econometrica* 58, 1391–1409.
- King, Robert, and Ross Levine, 1993, Finance and growth: Schumpeter might be right, *Quarterly Journal of Economics* 108, 717–737.
- La Porta, Rafael, Florencio Lopez-de Silanes, Andrei Shleifer, and Robert Vishny, 1997, Legal determinants of external finance, *Journal of Finance* 52, 1131–1150.
- Lambert, Richard, 1986, Executive effort and selection of risky projects, *The Rand Journal of Economics* 17, 77–88.
- Landier, Augustin, 2002, Entrepreneurship and the stigma of failure, Working paper, University of Chicago Graduate School of Business.
- Lazear, Edward, 1986, Salaries and piece rates, *The Journal of Business* 59, 405–431.
- Lerner, Josh, and Julie Wulf, 2007, Innovation and incentives: Evidence from corporate R&D, *Review of Economics and Statistics* 89, 634–644.
- Lizzeri, Alessandro, Margaret Meyer, and Nicola Persico, 2002, The incentive effects of interim performance evaluations, Working paper, CARESS.
- March, James, 1991, Exploration and exploitation in organizational learning, *Organization Science* 2, 71–87.
- Moscarini, Giuseppe, and Lones Smith, 2001, The optimal level of experimentation, *Econometrica* 69, 1629–1644.
- Myers, Stewart, 2000, Outside equity, *Journal of Finance* 55, 1005–1037.
- Narayanan, M. P., 1985, Managerial incentives for short-term results, *Journal of Finance* 15, 1469–1484.
- Nelson, Richard, 1959, The simple economics of basic scientific research, *Journal of Political Economy* 67, 297–306.
- Povel, Paul, 1999, Optimal “soft” or “tough” bankruptcy procedures, *Journal of Law, Economics and Organization* 15, 659–684.
- Prendergast, Canice, 2002, The tenuous trade off between risk and incentives, *Journal of Political Economy* 110, 1071–1102.
- Rajan, Raghuram, and Luigi Zingales, 1998, Financial dependence and growth, *American Economic Review* 88, 559–586.
- Ray, Korok, 2007, Performance evaluation and efficient sorting, *Journal of Accounting Research* 45, 839–882.
- Roberts, Kevin, and Martin Weitzman, 1981, Funding criteria for research, development, and exploration projects, *Econometrica* 49, 1261–1288.
- Romer, Paul, 1986, Increasing returns and long-run growth, *Journal of Political Economy* 94, 1002–1037.
- Sapra, Haresh, Ajay Subramanian, and Krishnamurthy Subramanian, 2009, Corporate governance and innovation: Theory and evidence, Working paper, University of Chicago Graduate School of Business.

- Schumpeter, Joseph, 1934, *The Theory of Economic Development* (Harvard University Press, Cambridge, MA).
- Seru, Amit, 2010, Do conglomerates stifle innovation? Firm boundaries matter: Evidence from conglomerate and R&D activity, *Journal of Financial Economics*, forthcoming.
- Shleifer, Andrei, and Robert Vishny, 1990, Equilibrium short horizons of investors and firms, *American Economic Review* 80, 148–153.
- Stein, Jeremy, 1988, Takeover threats and managerial myopia, *Journal of Political Economy* 96, 61–80.
- Stein, Jeremy, 1989, Efficient capital markets, inefficient firms: A model of corporate behavior, *Quarterly Journal of Economics* 104, 655–669.
- Stiglitz, Joseph, and Andrew Weiss, 1983, Incentive effects of termination: Applications to the credit and labor markets, *American Economic Review* 73, 912–927.
- Sutton, Robert, 2002, *Weird Ideas that Work: 11 1/2 Practices for Promoting, Managing, and Sustaining Innovation* (The Free Press, New York, NY).
- Thomke, Stefan, 2002, Bank of America, Harvard Business School case studies 9–603–022 and 9–603–023 (Harvard Business School Publishing, Boston, MA).
- Tian, Xuan, and Tracy Wang, 2010, Tolerance for failure and corporate innovation, Working paper, Indiana University.
- Weitzman, Martin, 1979, Optimal search for the best alternative, *Econometrica* 47, 641–654.
- Zwiebel, Jeffrey, 1995, Corporate conservatism and relative compensation, *Journal of Political Economy* 103, 1–25.