Motivating Innovation

Gustavo Manso*

September 29, 2006

Abstract

Motivating innovation is an important concern in many incentive problems. For example, managers often claim that it is difficult to motivate their employees to devise innovative ways of doing things. The difficulty arises because innovation is the result of the exploration of untested approaches that are likely to fail, and failure is usually associated with low wages and termination. This paper shows that incentive schemes that motivate exploration are fundamentally different from standard pay-for-performance incentive schemes used to motivate effort. The optimal compensation scheme that motivates exploration exhibits substantial tolerance (or even reward) for early failure and reward for long-term success. Moreover, even though the principal can terminate the agent, inefficient continuation may be optimal to induce exploration since the threat of termination may prevent the agent from exploring new untested approaches. Finally, commitment to a long-term compensation plan and timely feedback on performance are essential ingredients to induce exploration. The institution of tenure, debtor-friendly bankruptcy laws, and golden parachutes are examples of schemes that protect the agent when failure occurs and thereby encourage exploration.

*MIT Sloan School of Management, 50 Memorial Way E52-403B, Cambridge, MA 02139 (e-mail: manso@mit.edu, url: www.mit.edu/~manso/). I thank John Roberts, Anat Admati, Darrell Duffie, and Ed Lazear for valuable discussions and suggestions. I also thank David Ahn, Snehal Banerjee, Peter DeMarzo, Jerker Denrell, Willie Fuchs, Bob Gibbons, Thomas Hellman, Ayca Kaya, Ilan Kremer, Jim March, Pedro Miranda, Stefan Reichelstein, Paul Pfleiderer, Leonardo Rezende, Tomasz Sadzik, Maria Salgado, Yuliy Sannikov, Ilya Segal, Kathryn Shaw, Andy Skrzypacz, Ilya Strebulaev, Alexei Tchistyi, Bob Wilson, Jeff Zwiebel and seminar participants at Berkeley, Chicago, Columbia, Duke, FGV, MIT, Northwestern, NYU, Princeton, PUC-Rio, Stanford, UCLA, UCSD, and UPenn for helpful comments. Financial support from a Stanford Institute for Policy Research (SIEPR) Fellowship and the Lehman Brothers Fellowship for Research Excellence in Finance is gratefully acknowledged.

Results? Why, man, I have gotten lots of results! I know several thousands of things that won't work. *Thomas Edison*

1 Introduction

Innovation is vital for the long-term growth and performance of organizations. Top executives in business organizations are aware of that. In a recent survey,[1] approximately 78 percent of the 540 CEOs interviewed said that "stimulating innovation, creativity, and enabling entrepreneurship" is a top priority of their organizations. Motivating innovation remains, however, a challenge for most organizations. Innovation results from the exploration of new untested approaches that are likely to fail. Therefore, standard pay-for-performance schemes that punish failures with low wages and termination may have adverse effects on innovation.

The key contribution of this paper is to show that the incentive schemes that motivate innovation are fundamentally different from standard pay-for-performance schemes. The optimal compensation scheme that motivates innovation exhibits substantial tolerance (or even reward) for early failure and reward for long-term success. Moreover, even though the principal can terminate the agent, inefficient continuation may be optimal to motivate innovation since the threat of termination may prevent the agent from exploring new untested approaches. Finally, commitment to a long-term compensation plan and timely feedback on performance are also essential ingredients to motivate innovation.

Principal-agent models in which the principal must motivate the agent to exert effort, such as Harris and Raviv (1978) and Holmstrom (1979), form the basis of the intuition most economists have about incentives. These models predict that pay-for-performance motivates the agent to exert more effort, improving performance. Empirical studies that support the predictions of the standard principal-agent models usually deal with simple routine tasks, in which effort is the main input of the worker. For example, Lazear (2000) finds that the productivity of windshield installers in Safelite Glass Corporation increases when their compensation method changes from hourly wages to piece-rate pay. In contrast, a substantial body of experimental and field research in psychology provides evidence that, in tasks that require exploration and creativity, pay-for-performance may actually undermine performance. McGraw (1978), McCullers (1978), Kohn (1993) and Amabile (1996) survey this line of research and conclude that pay-for-performance encourages the repetition of what has worked in the past, but not the exploration of new untested approaches.

Some commonly used incentive schemes protect or even reward the agent when failure occurs. For example, innovative corporations such as 3M and IBM are known for having corporate cultures that tolerate failure. Research departments in business and academic organizations grant researchers tenure, securing them a job even if their subsequent productivity is low. Executive compensation packages include golden parachutes and option repricing, rewarding executives

[1] "CEO Challenge 2004: Perspectives and Analysis," The Conference Board, Report 1353.

after poor performance. An entrenched manager may keep his job even if it is ex post efficient for the firm to fire him. Debtor-friendly bankruptcy laws allow entrepreneurs a fresh start after failure. These incentive schemes are often criticized because, by protecting or rewarding the agent after poor performance, they undermine the incentives for the agent to exert effort. This paper shows that these incentive schemes may arise as part of an optimal contract that motivates exploration. Restricting their use may thus have adverse effects on innovation.

To model the process of innovation, I use a class of Bayesian decision models known as bandit problems.[2] In bandit problems, the agent does not know the true distribution of payoffs of the available actions. 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 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 model has two important special cases. On one hand, if exploration and exploitation are costless to the agent, then there is no conflict of interest between the principal and the agent. The model reduces to the two-armed bandit problem that captures the tension between exploration and exploitation. On the other hand, if exploration is extremely costly to the agent, then the agent chooses between exploitation or shirking. The model reduces to a standard principal-agent model where 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. 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 exploration 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 relies on long-term incentives.

The ability of the principal to commit to a long-term contract has different effects on the incentives for exploitation and exploration. Since the optimal contract that motivates exploitation relies on short-term incentives, the ability to commit to a long-term contract is irrelevant to motivate exploitation. In contrast, the optimal contract that motivates exploration relies on long-term

[2]Berry and Fristedt (1985) provides an introduction to the statistical literature on bandit problems. Bergemann and Valimaki (2006) surveys the applications of bandit problems to economics.

incentives. Not only are short-term contracts strictly dominated by long-term contracts, but exploration may not even be implementable with a sequence of short-term contracts. The ability to commit to a long-term contract is thus essential to motivate exploration.

I study the effects of termination after poor performance on the incentives for exploration and exploitation. Since the threat of termination helps to prevent the agent from shirking or exploring new actions, termination facilitates the provision of incentives for exploitation. Inefficient termination may thus be optimal to motivate exploitation. In contrast, the effects of termination on the incentives for exploration are ambiguous. On one hand, the threat of termination prevents the agent from shirking. On the other hand, the threat of termination encourages the agent to exploit well known actions. Depending on which of these two effects is predominant, termination may either facilitate or hinder the provision of incentives for exploration. Either inefficient termination or continuation may thus be optimal to motivate exploration.

Finally, I study the role of feedback on performance when the principal is better able than the agent to evaluate performance. If the principal does not provide feedback on performance to the agent, then the agent cannot adjust his action according to performance. There are thus fewer deviations that the agent can attempt. Therefore, it is cheaper to motivate exploitation if the principal does not provide feedback on performance to the agent. In contrast, since exploration requires the agent to adjust his action after poor performance, feedback on performance is essential to motivate exploration.

The model of the innovation process adopted here follows a long tradition in the study of innovation. Schumpeter (1934) argues that innovation results from the 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. March (1991) coined the terms exploration and exploitation to describe the fundamental tension that arises in learning through experimentation. Endogenous growth theory stresses the relation between innovation, growth and the production of knowledge. Romer (1986) is an early contribution to this literature. More recent papers, such as Jovanovic and Nyarko (1996), develop quite explicit models of innovation as the result of learning from the exploration of new technologies. Bolton and Harris (1999) study strategic experimentation in a setting where multiple players face the same experimentation problem. Each agent learns from the experimentation of the other players. Since information is a public good, there is free riding and under-experimentation in equilibrium. In contrast to these papers, which take the payoffs of the players as given, I study the optimal compensation schemes that motivate exploration and exploitation.

Other papers have studied principal-agent models in which the choice of the agent is not limited to the level of effort. Lambert (1986) analyzes the provision of incentives when the agent selects among risky projects. Holmstrom and Milgrom (1991) develop a multi-task principal-agent model in which the agent allocates effort among multiple tasks and the principal observes a performance measure for each of these tasks. Dewatripont and Maskin (1995) and Von Thadden (1995) analyze incentives for the agent to select between short-term and long-term investments. In these models, the distribution of payoffs is known. Moreover, the agent takes an action in the first period and cannot change it later on. Therefore, these models do not incorporate experimentation, learning, and adaptation, important features of innovation.

Some papers study the incentives for innovation. Holmstrom (1989) assumes that performance measures for innovative activities are noisier, and therefore to motivate innovation the principal should use compensation schemes that are less sensitive to performance. In an incomplete contracting framework, Aghion and Tirole (1994) derive the optimal allocation of control rights that motivates innovation. In contrast to these papers, I model innovation as the result of learning through experimentation and study the central trade-off that arises in this setting, the trade-off between exploration and exploitation.

Recent research in organizational economics has studied how economic organization allows agents to most effectively use and acquire knowledge. Garicano (2000) shows that a hierarchical structure with lower-knowledge agents passing problems they cannot solve to higher-knowledge agents is the most efficient way to organize problem solving. Lee and Van den Steen (2006) study what type of knowledge an organization should collect and disseminate among its employees. The current paper is complementary to the above papers, as it studies how to provide incentives for employees to produce new knowledge, or innovate.

The results obtained here may help us understand issues in the management of creative workers. Business consultants usually say that to motivate innovation it is essential to have a corporate culture that tolerates failure, allows freedom to experiment, and rewards successful innovations in the long-run. Promises made in the form of a corporate culture can be enforced through reputation. In addition to a corporate culture that supports innovation, organizations may also use explicit contracts to encourage exploration. For example, business and academic organizations often protect researchers against failure by granting them tenure. Knowing they will not lose their jobs, these researchers are willing to take risks on research directions that are likely to fail, but may lead to breakthroughs.

This paper also sheds light on issues in executive compensation and corporate governance. Due to the recent scandals in the American corporate sector, executive compensation has been criticized as excessive and not related to performance. These attacks generate pressure for regulatory changes that limit the use of stock options, option repricing, golden parachutes, and entrenchment. However, this paper shows that the optimal contract that motivates exploration can be implemented through combinations of stock options, option repricing, golden parachutes, and entrenchment. Protecting the manager against failure is important to encourage exploration of new business models and organizational structures. Attempts to restrict the use of these instruments may thus have adverse effects on innovation, harming the long-term growth and performance of corporations and of the whole economy.

Finally, this paper helps us understanding the relationship between bankruptcy laws and entrepreneurship. To motivate entrepreneurship and reduce the productivity gap with the United States, the European Union is attempting to make the bankruptcy laws of its member countries less severe towards insolvent debtors. In contrast, after seeing a dramatic increase in the number of personal bankruptcies, the U.S. Congress recently passed a new creditor-friendly bankruptcy law. Debtor-friendly bankruptcy laws are optimal to implement exploration, because they protect entrepreneurs in case of failure, while creditor-friendly bankruptcy laws are optimal to implement exploitation, because they punish the entrepreneur in case of failure. One natural question to ask is why governments do not offer a menu of bankruptcy laws so that, upon contracting, creditors and debtors can choose the bankruptcy law that best fit their needs. Due to knowledge spillovers and imperfect protection of intellectual property rights, however, individuals pursuing exploratory activities cannot fully appropriate the economic value generated by the knowledge they create. This leads to underinvestment in exploration. By imposing debtor-friendly bankruptcy laws instead of offering a menu of bankruptcy laws, governments may alleviate the problem of underinvestment in exploration.

The paper is organized as follows. Section 2 discusses the tension between exploration and exploitaton in a single-agent decision problem. Section 3 introduces the tension between exploration and exploitation into a principal-agent model. Section 4 studies incentives for exploration and exploitation. Section 5 investigates the situation in which the principal cannot commit to a long-term contract. Section 6 studies what happens if the principal can terminate the agent after bad performance. Section 7 examines the optimal provision of feedback when the principal is better able than the agent to evaluate performance. Section 9 discusses applications of the model. Section 10 contains additional discussion and Section 11 concludes. All proofs are in the Appendix.

2 The Single-Agent Decision Problem

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

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

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 by:

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

The agent is risk-neutral and has a discount factor normalized to one. The agent thus chooses an action plan $\langle {}_{i}\,{}^{j}_{k} \rangle$ to maximize his total expected payoff

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

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

Two action plans need to be considered. Action plan $\langle {}_{1}\,{}^{1}_{1} \rangle$, which I call exploitation, is just the repetition of the conventional work method. Action plan $\langle {}_{2}\,{}^{2}_{1} \rangle$, which I call exploration, is to initially try the new work method, stick to the new work method in case of success in the first period, and revert to the conventional work method in case of failure in the first period. The total payoff $R(\langle {}_{2}\,{}^{2}_{1} \rangle)$ from exploration is higher than the total payoff $R(\langle {}_{1}\,{}^{1}_{1} \rangle)$ 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 in case 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. In case the agent discovers that the new work method is better than the conventional work method, the whole team benefits from his discovery.

3 The Principal-Agent Problem

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

The principal hires an agent to perform the task described in the previous section. The principal does not observe the actions taken by the agent. In each period, the agent incurs private costs $c_1 \geq 0$ if he takes action 1, the conventional work method, 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. Shirking has a lower probability of success than either of the two work methods, so that

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

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 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 1. When the principal offers the agent a contract $\vec{w}$ and the agent takes action plan $\langle i \, {}^{j}_{k} \rangle$, the total expected payments from the principal to the agent are given by

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

When the agent takes action plan $\langle i \, {}^{j}_{k} \rangle$, the total expected costs incured by the agent are given by

$$C(\langle i \, {}^{j}_{k} \rangle) = c_i + E[p_i]c_j + (1 - E[p_i])c_k.$$

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

$$W(\vec{w}, \langle i \, {}^{j}_{k} \rangle)$$

subject to to the incentive compatibility constraints,[3]

$$W(\vec{w}, \langle i\,{}^{j}_{k} \rangle) - C(\langle i\,{}^{j}_{k} \rangle) \geq W(\vec{w}, \langle l\,{}^{m}_{n} \rangle) - C(\langle l\,{}^{m}_{n} \rangle). \qquad (\mathrm{IC}_{\langle l\,{}^{m}_{n} \rangle})$$

This is a linear program with 6 unknowns and 27 constraints. When there is more than one contract that solves this program, I restrict attention to the contract that pays the agent earlier.[4]

The principal's expected profit $\Pi(\langle i\,{}^{j}_{k} \rangle)$ from implementing action plan $\langle i\,{}^{j}_{k} \rangle$ is given by

$$\Pi(\langle i\,{}^{j}_{k} \rangle) = R(\langle i\,{}^{j}_{k} \rangle) - W(\vec{w}(\langle i\,{}^{j}_{k} \rangle), \langle i\,{}^{j}_{k} \rangle). \tag{5}$$

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

Both the classical two-armed bandit problem and the standard work-shirk principal-agent model are special cases of this model. On 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 the principal just solves the two-armed bandit problem described in Section 2. On the 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.

4 Incentives for Exploration and Exploitation

In this section I study the optimal contracts that implement exploration and exploitation respectively. Given this focus, I do not study whether it is exploration or exploitation that yields a higher total expected profit for the principal. It is easy though to find situations in which exploration is optimal and in which exploitation is optimal. For example, if the principal expands production substantially after a success in the first period, then exploration is optimal.[5]

The relative costs c_2/c_1 between the new and the conventional work methods will be important in determining which incentive compatibility constraints are

[3] For simplicity, I assume that the agent has zero reservation utility. The participation constraints is thus not binding, since the agent has limited liability.

[4] 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 this.

[5] The expected revenue function

$$\begin{aligned} R(\langle i\,{}^{j}_{k} \rangle) = \{E[p_i]S + (1 - E[p_i])F\} & \\ + \gamma_S E[p_i]\,\{E[p_j|S,i]S + (1 - E[p_j|S,i])F\} & \\ + \gamma_F (1 - E[p_i])\,\{E[p_k|F,i]S + (1 - E[p_k|F,i])F\}\,, & \quad (6) \end{aligned}$$

allows for situations in which the principal changes the scale of production according to first-period performance. For example, if γ_S is high the principal expands production after a success in the first period.

binding, and consequently the form of the optimal contract. When c_2/c_1 is high, the agent is more inclined to conform to the conventional work method than to experiment with the new work method. This could be due to the extra effort incurred by the agent when searching and implementing a new work method. When c_2/c_1 is low, the agent is more inclined to experiment with the new work method than to conform to the conventional work method. This could be due to private benefits of learning a new work method. For clarity of exposition, I will restrict attention to

$$c_2/c_1 \geq (E[p_2] - p_0)/(p_1 - p_0). \tag{7}$$

Since the right-hand side of equation (7) is lower than 1, restricting attention to (7) 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.

4.1 Incentives for Exploitation

Proposition 1 derives the optimal contract that implements exploitation. Recalling from Section 2 exploitation is given by the action plan $\langle 1\,{}^{1}_{1} \rangle$. The following definitions are 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 $\vec{w}_1$ that implements exploitation is such that*

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

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

$$w_S = \alpha_1 + \frac{c_1}{(p_1 - p_0)(p_1 - E[p_2])}\left(\beta_1 - \frac{c_2}{c_1}\right)^+,$$

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

The formal proofs of all the propositions are in the Appendix. Here is the main intuition behind Proposition 1. To implement exploitation, the principal must prevent the agent from shirking and from exploring. The principal does not make payments to the agent after failures, since this only gives incentives for the agent to shirk or explore. Although there are 27 incentive compatibility constraints, it is easy to see that only a few may bind. The relevant incentive compatibility constraints are

$$(p_1 - p_0)w_{SS} \geq c_1 \tag{$\text{IC}_{\langle 1\,{}^{0}_{1}\rangle}$}$$

$$(p_1 - p_0)w_{FS} \geq c_1 \qquad (\text{IC}_{\langle 1\,{}^{1}_{0}\rangle})$$

$$(p_1 - p_0)w_S + (p_1^2 - p_0 p_1)w_{SS} - (p_1^2 - p_0 p_1)w_{FS} \geq c_1 \qquad (\text{IC}_{\langle 0\,{}^{1}_{1}\rangle})$$

$$(p_1 - E[p_2])w_S + (p_1^2 - E[p_2]E[p_2|S,2])w_{SS} - (p_1^2 - E[p_2]p_1)w_{FS} \geq$$
$$c_1 - c_2 + E[p_2](c_1 - c_2) \qquad (\text{IC}_{\langle 2\,{}^{2}_{1}\rangle})$$

The first three incentive compatibility constraints are associated with shirking. The last incentive compatibility constraint is associated with exploration.

If $c_2/c_1 \geq \beta_1$, then exploration is too costly for the agent. Only the incentive compatibility constraints associated with shirking are binding. The optimal contract that implements exploitation is identical to the optimal contract used to induce repeated effort in a standard work-shirk principal-agent model, so that $w_S = w_{SS} = w_{FS} = \alpha_1$.

If $c_2/c_1 < \beta_1$, exploration is not too costly for the agent. Instead of the incentive compatility constraint associated with shirking in the first period, it is the incentive compatibility constraint associated with exploration that is binding. 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, the principal must pay the agent an extra premium for a success in the first period to prevent exploration. Therefore, $w_S > w_{SS} = w_{FS} = \alpha_1$.

Figure 1 shows the optimal contract $\vec{w}_1$ that implements exploitation for different values of c_2/c_1 under the base case parameters.[6] 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, except that, if c_2/c_1 is low, the principal pays the agent an extra premium in case of success in the first period to prevent exploration. This extra premium is decreasing in c_2/c_1, since as c_2/c_1 increases the agent becomes less inclined to explore.

4.2 Incentives for Exploration

Proposition 2 derives the optimal contract that implements exploration. Recalling from Section 2, exploration is given by action plan $\langle 2\,{}^{2}_{1}\rangle$. 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.*

[6]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' rule, $E[p_2|F,2] = 0.129$.

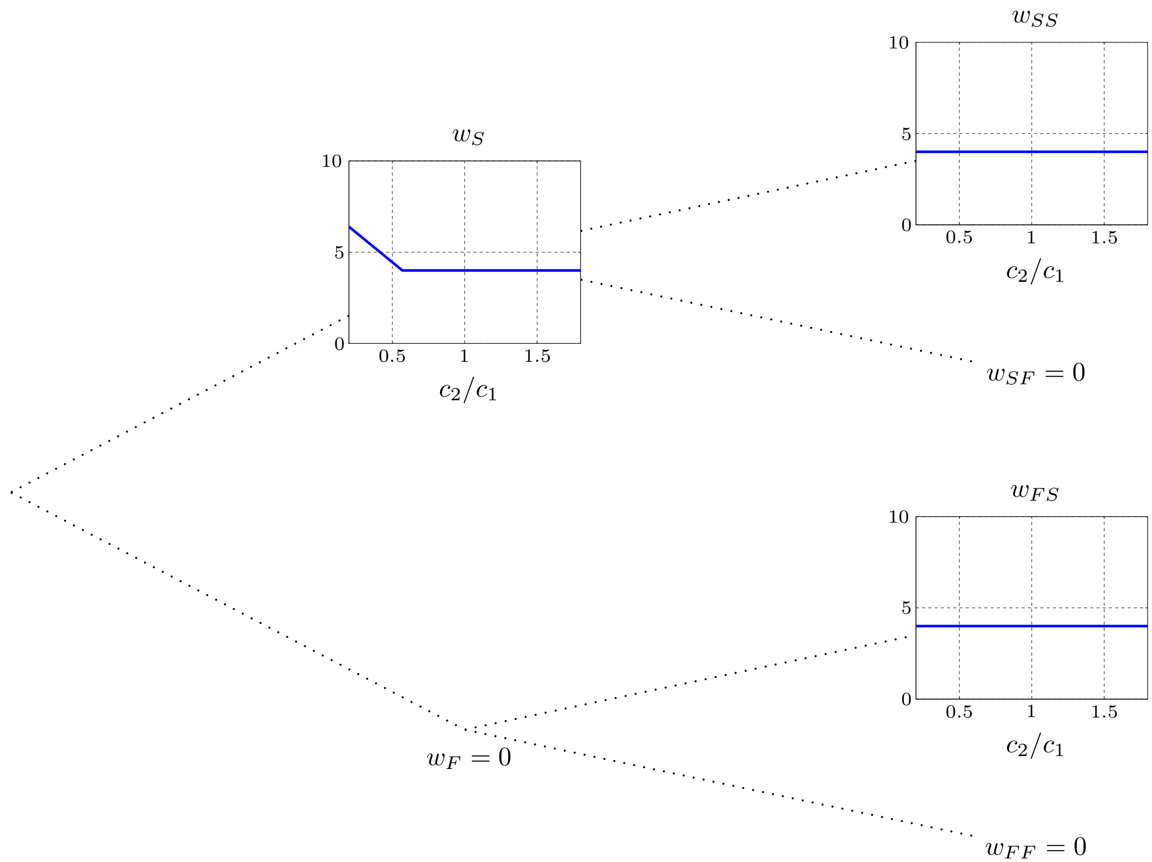

Figure 1: The optimal contract that implements exploitation under the base case parameters.

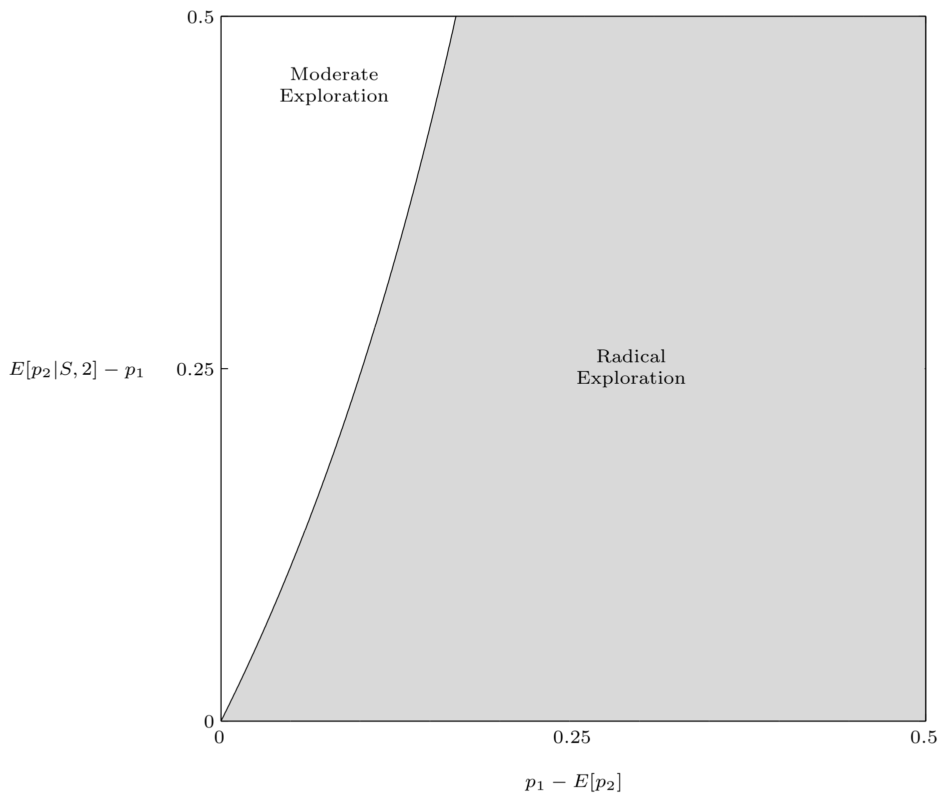

Figure 2: Moderate exploration (plain region) versus radical exploration (shaded region) when $p_1 = 1/2$.

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. Figure 2 illustrates the regions of moderate and radical exploration when $p_1 = 1/2$.

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

$$\alpha_2 = \max_{\tilde{\jmath} \in \{0,1\}} \frac{(1 + E[p_2])c_2 - p_0 c_{\tilde{\jmath}}}{E[p_2]E[p_2|S,2] - p_0 E[p_{\tilde{\jmath}}]} + \frac{(E[p_2] - p_0)p_0 \alpha_1}{E[p_2]E[p_2|S,2] - p_0 E[p_{\tilde{\jmath}}]}.$$

$$\beta_2 = \frac{(E[p_2]E[p_2|S,2] - p_0 p_1) + E[p_2](p_1 E[p_2|S,2] - p_0 p_1)}{(p_1^2 - p_0) + E[p_2](p_1^2 - p_0)}$$

Proposition 2 *The optimal contract $\vec{w}_2$ that implements exploration is such that*

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

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

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

If exploration is radical, then

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

and

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

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 case of a 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

$$(p_1 - p_0)w_{FS} \geq c_1 \tag{$\text{IC}_{\langle 2\,{}^{2}_{0}\rangle}$}$$

$$\begin{aligned}&(E[p_2]E[p_2|S,2] - p_0^2)w_{SS} - (E[p_2] - p_0)w_F - (E[p_2] - p_0)p_1 w_{FS} \\ &\qquad \geq c_2 + E[p_2](c_2 - c_1) + p_0 c_1\end{aligned} \tag{$\text{IC}_{\langle 0\,{}^{0}_{1}\rangle}$}$$

$$\begin{aligned}&(E[p_2]E[p_2|S,2] - p_0 p_1)w_{SS} - (E[p_2] - p_0)w_F - (E[p_2] - p_0)p_1 w_{FS} \\ &\qquad \geq c_2 + E[p_2](c_2 - c_1)\end{aligned} \tag{$\text{IC}_{\langle 0\,{}^{1}_{1}\rangle}$}$$

$$\begin{aligned}&(E[p_2]E[p_2|S,2] - p_1^2)w_{SS} + (p_1 - E[p_2])w_F + (p_1 - E[p_2])p_1 w_{FS} \\ &\qquad \geq (1 + E[p_2])(c_2 - c_1)\end{aligned} \tag{$\text{IC}_{\langle 1\,{}^{1}_{1}\rangle}$}$$

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 $w_{FS} = \alpha_1$ to the agent 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 $w_{SS} = \alpha_2$ to the agent.

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 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 success 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 3 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 failures, then the agent may prefer to exploit in order to avoid 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 a 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 that fails twice may have a higher compensation than an agent that obtains short-lived success. Because of the risky nature of exploration, failing twice may be a

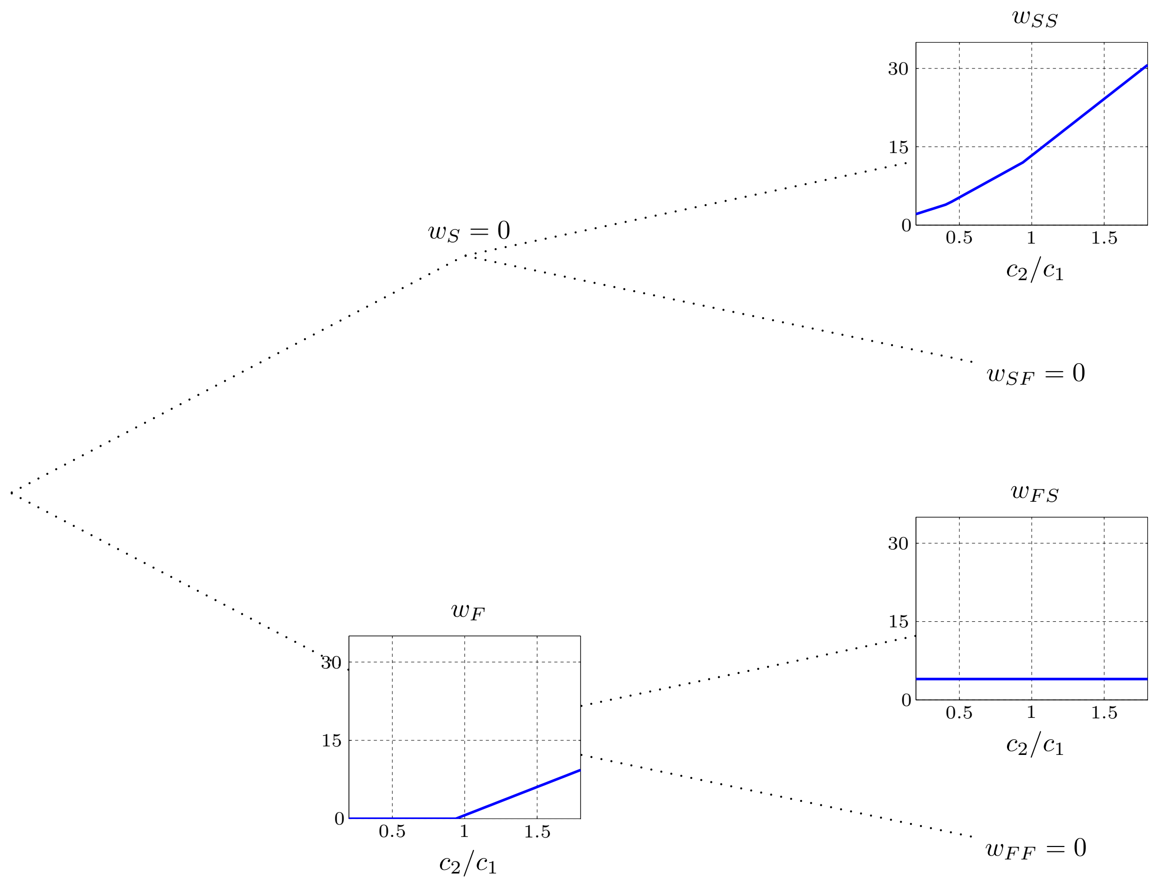

Figure 3: The optimal contract that implements exploration under the base parameters.

stronger signal for the principal that the agent explored and not exploited than obtaining a short-lived success.

5 Commitment

In contrast to the previous sections, I now assume that the principal cannot commit to a long-term contract. In each period, the principal can only offer the agent a short-term contract specifying the agent's wages contingent on the current period performance. The problem is similar to the one proposed in Section 3, 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 $\vec{w}_1$ 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 $\vec{w}_2$ 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 $\vec{w}_4$ 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, on the contrary, 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.[7] 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 $W(\vec{w}_4, \langle {}_2{}^2_1 \rangle)$ of implementing exploration with short-term contracts is higher than the cost $W(\vec{w}_2, \langle {}_2{}^2_1 \rangle)$ 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.

Figure 4 compares the optimal contracts when the principal can and cannot commit to a long-term contract for different values of c_2/c_1 under the base case parameters. Under the base case parameters, exploration is implementable via a sequence of short-term contracts only if c_2/c_1 is very low. In this case, the optimal sequence of short-term contracts relies on short-term incentives to implement exploration, while the optimal long-term contract relies on long-term incentives to implement exploration. If c_2/c_1 is high, it is not possible to implement exploration with a sequence of short-term contracts. Figure 5 shows that even if exploration is implementable via a sequence of short-term contracts, the cost $W(\vec{w}_4, \langle {}_2{}^2_1 \rangle)$ of implementing exploration with short-term contracts can be significantly higher than the cost $W(\vec{w}_2, \langle {}_2{}^2_1 \rangle)$ of implementing exploration with a long-term contract.

6 Termination

In this section, I allow the principal to terminate the agent after a failure in the first period. 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.

[7]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 outcome 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 outcome as action 2 and costs less to the agent.

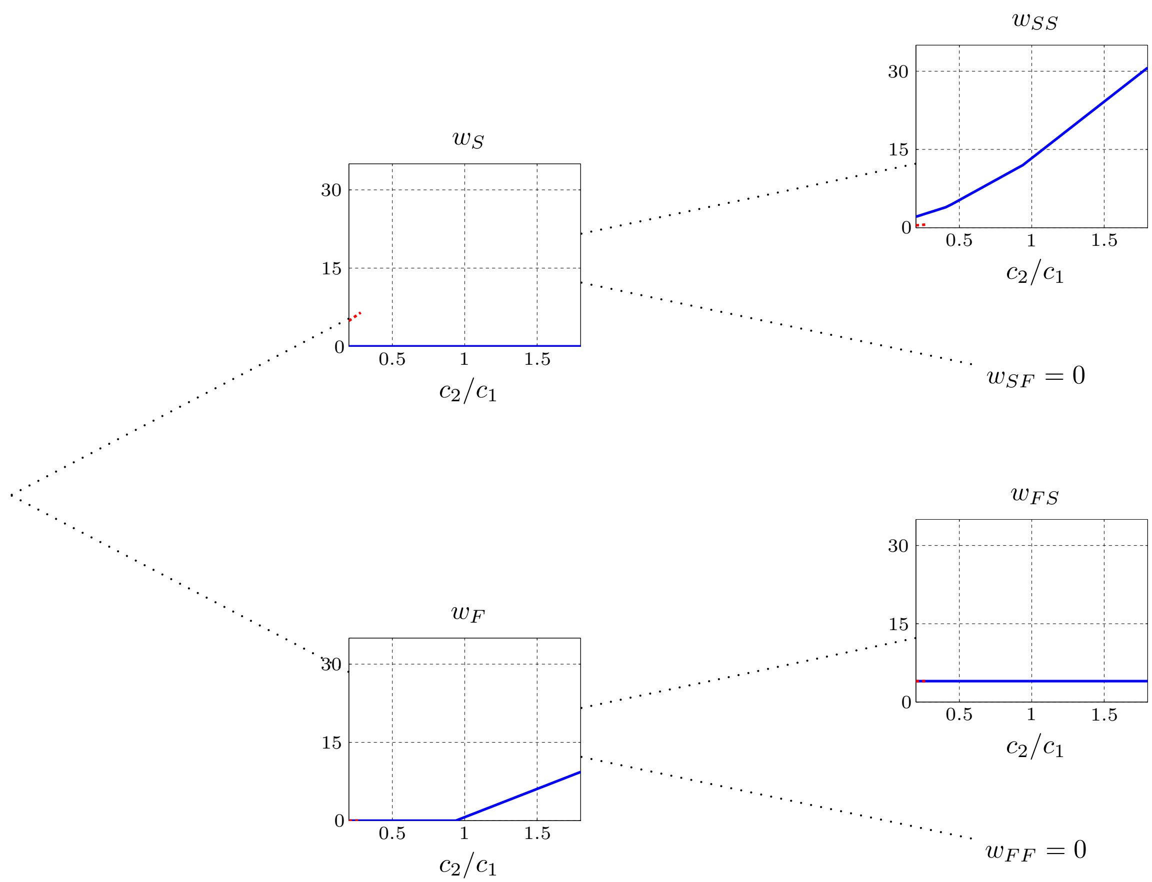

Figure 4: The optimal contract that implements exploration with (solid line) and without (dashed line) commitment under the base parameters.

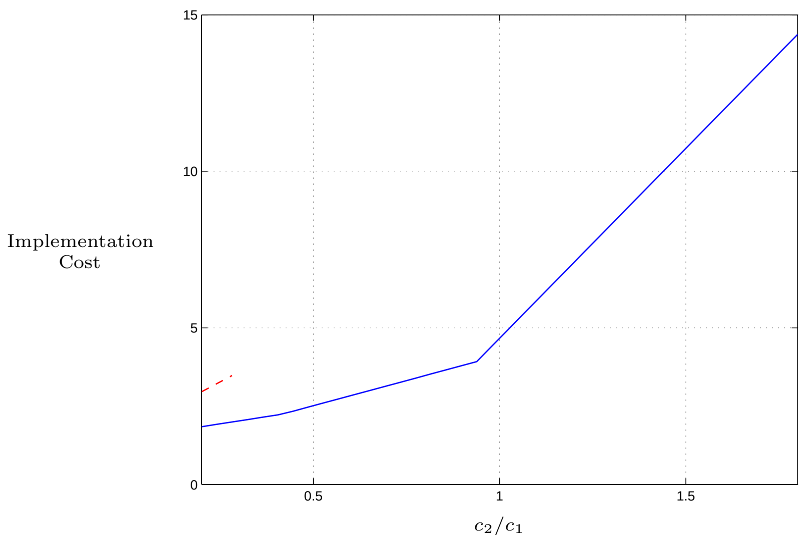

Figure 5: Cost of implementing exploration when the principal can (solid line) and cannot (dashed line) commit to a long-term contract under the base parameters.

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 $\langle 1\,{}^{1}_{t} \rangle$, and exploration with termination is represented by action plan $\langle 2\,{}^{2}_{t} \rangle$, 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.[8]

Proposition 5 derives the optimal contract that implements exploitation with termination.

Proposition 5 *The optimal contract $\vec{w}_5$ that implements exploitation with termination is such that*

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

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

and

$$w_S = (1-p_0)\alpha_1 + \frac{c_1}{(p_1-p_0)(p_1-E[p_2])}\left(\beta_1 - \frac{c_2}{c_1}\right)^+ .$$

The agent is more likely to fail in the first period if he shirks or employs the new work method than if he employs the conventional work method. To avoid

[8]In this context, implementing action plan $\langle i\,{}^{j}_{t} \rangle$ is the same as implementing action plan $\langle i\,{}^{j}_{0} \rangle$ with $w_{FF} = w_{FS} = 0$.

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 help inducing the agent to exert effort. Figure 6 compares the optimal contracts that implement exploitation and exploitation with termination for different values of c_2/c_1 under the base case parameters.

I now compare the total expected profits of the principal when he implements exploitation with the total expected profits of the principal when he implements exploitation with termination. It is optimal for the principal to implement exploitation with termination instead of exploitation if

$$R(\langle 1\,{}^{1}_{1}\rangle) - R(\langle 1\,{}^{1}_{t}\rangle) < W(\vec{w}_1, \langle 1\,{}^{1}_{1}\rangle) - W(\vec{w}_5, \langle 1\,{}^{1}_{t}\rangle). \tag{8}$$

To prevent the agent from shirking 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

$$R(\langle 1\,{}^{1}_{1}\rangle) - R(\langle 1\,{}^{1}_{t}\rangle) < (1-p_1)p_1\alpha_1. \tag{9}$$

When (9) holds, the benefits from inducing the agent not to shirk 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 prevent the agent from shirking in the second period.

Definition 2 *There is inefficient termination with exploitation if*

$$W(\vec{w}_1, \langle 1\,{}^{1}_{1}\rangle) - W(\vec{w}_5, \langle 1\,{}^{1}_{t}\rangle) > (1-p_1)p_1\alpha_1.$$

and there is inefficient continuation with exploitation if

$$W(\vec{w}_1, \langle 1\,{}^{1}_{1}\rangle) - W(\vec{w}_5, \langle 1\,{}^{1}_{t}\rangle) < (1-p_1)p_1\alpha_1.$$

There is inefficient termination with exploitation if the actual threshold for termination is higher than the efficient threshold for termination. There is inefficient continuation with exploitation if the actual threshold for termination is lower than the efficient threshold for termination. Inefficient continuation or termination may arise because the termination policy affects the incentives for the agent's first-period action.

Corollary 1 *There is inefficient termination with exploitation.*

As shown in Proposition 5, termination acts as a disciplinary device so that to implement exploitation with termination the principal needs to pay the agent

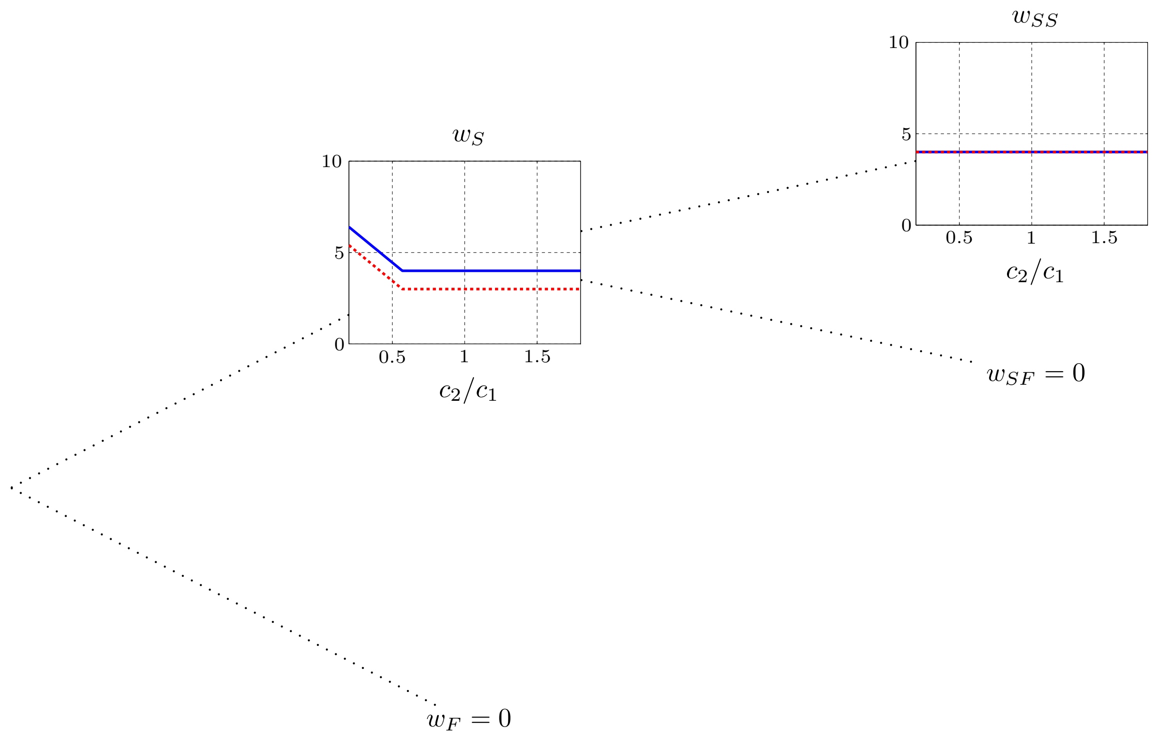

Figure 6: The optimal contracts that implement exploitation (solid line) and exploitation with termination (dashed line) under the base parameters.

lower first-period wages than to implement exploitation. There is inefficient termination with exploitation because the lower wages paid to the agent offset the losses from inefficient termination.

Proposition 6 derives the optimal contract that implements exploration with termination. The following definitions will be useful in stating Proposition 6:

$$\alpha_6 = \max_{\tilde{j} \in \{0,1\}} \frac{(1+E[p_2])c_2 - p_0 c_{\tilde{j}}}{E[p_2]E[p_2|S,2] - p_0 E[p_{\tilde{j}}]},$$

$$\beta_6 = \frac{(E[p_2]E[p_2|S,2] - p_0 p_1) + E[p_2](p_1 E[p_2|S,2] - p_0 E[p_2|S,2])}{(p_1^2 - p_0) + E[p_2](p_1^2 - p_0)}.$$

Proposition 6 *The optimal contract $\vec{w}_6$ that implements exploration with termination is such that*

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

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

$$w_{SS} = \alpha_6 + \frac{p_1 - p_0}{E[p_2]E[p_2|S,2] - p_0 p_1} \frac{p_1(1+E[p_2])c_1}{E[p_2]E[p_2|S,2] - p_1^2} \left(\frac{c_2}{c_1} - \beta_6\right)^+ .$$

If exploration is radical, then

$$w_F = \frac{p_1(1+E[p_2])c_1}{E[p_2]E[p_2|S,2] - p_1 E[p_2]} \left(\frac{c_2}{c_1} - \beta_6\right)^+$$

and

$$w_{SS} = \alpha_6 + \frac{E[p_2] - p_0}{E[p_2]E[p_2|S,2] - p_0 p_1} \frac{p_1(1+E[p_2])c_1}{E[p_2]E[p_2|S,2] - p_1 E[p_2]} \left(\frac{c_2}{c_1} - \beta_6\right)^+ .$$

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 $c_2/c_1 \geq \beta_6$, 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 $c_2/c_1 < \beta_6$, 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 7 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.

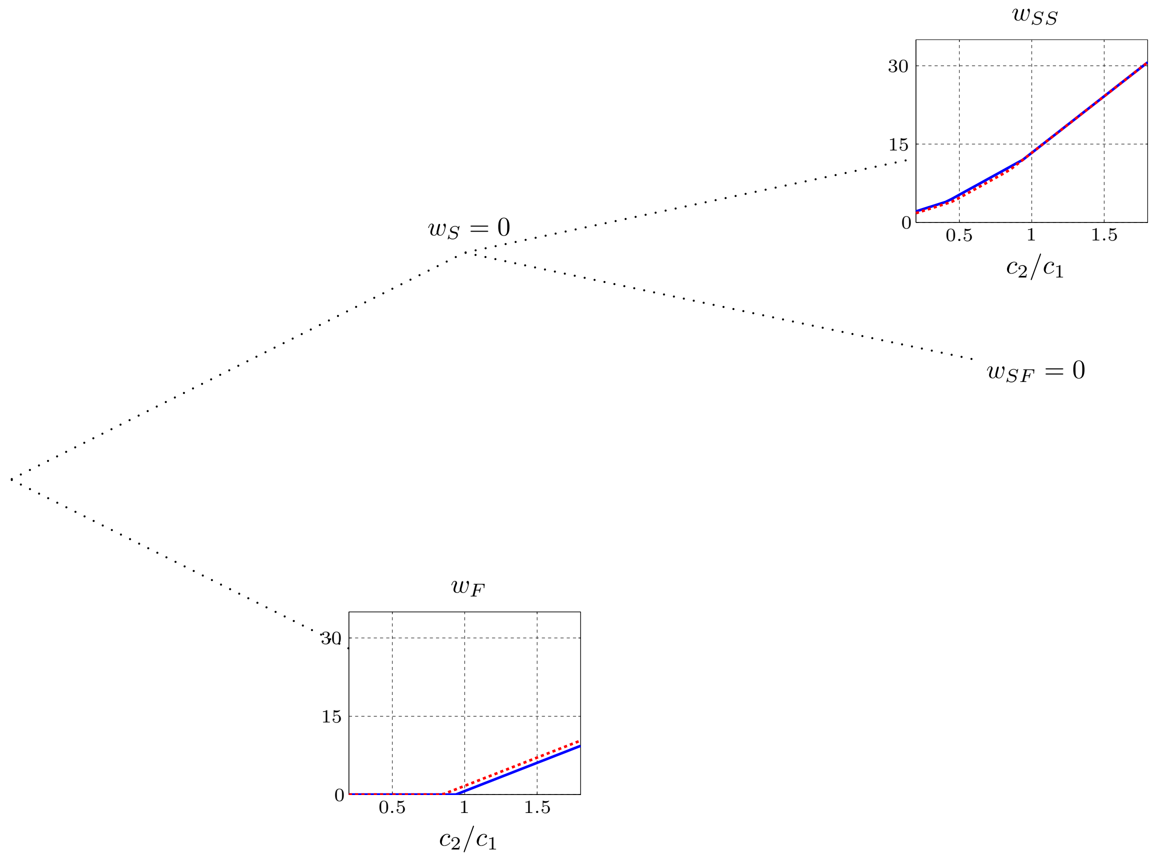

Figure 7: The optimal contract that implements exploration (solid line) and exploration with termination (dashed line) under the base parameters.

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

$$R(\langle {}_{2}\,{}^{2}_{1} \rangle) - R(\langle {}_{2}\,{}^{2}_{t} \rangle) < W(\vec{w}_2, \langle {}_{2}\,{}^{2}_{1} \rangle) - W(\vec{w}_6, \langle {}_{2}\,{}^{2}_{t} \rangle)$$

To prevent the agent from shirking 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

$$R(\langle {}_{2}\,{}^{2}_{1} \rangle) - R(\langle {}_{2}\,{}^{2}_{t} \rangle) < (1 - E[p_2])p_1\alpha_1. \tag{10}$$

When (10) holds, the benefits from inducing the agent not to shirk 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 prevent the agent from shirking in the second period.

Definition 3 *There is inefficient termination with exploration if*

$$W(\vec{w}_2, \langle {}_{2}\,{}^{2}_{1} \rangle) - W(\vec{w}_6, \langle {}_{2}\,{}^{2}_{t} \rangle) > (1 - E[p_2])p_1\alpha_1,$$

and there is inefficient continuation with exploration if

$$W(\vec{w}_2, \langle {}_{2}\,{}^{2}_{1} \rangle) - W(\vec{w}_6, \langle {}_{2}\,{}^{2}_{t} \rangle) < (1 - E[p_2])p_1\alpha_1.$$

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

Corollary 2 investigates the conditions under which there is inefficient termination with exploration or inefficient continuation with exploration. The following definitions will be useful in stating Corollary 2:

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

$$\kappa_e \equiv \frac{(1 - E[p_2])(E[p_2]E[p_2|S,2] - p_0 p_1)}{(1 - E[p_2])(E[p_2]E[p_2|S,2] - p_0 p_1) + E[p_2]E[p_2|S,2](E[p_2] - p_0)}.$$

Corollary 2 *If $c_2/c_1 < \max(\kappa_m, \kappa_e)\beta_2 + (1 - \max(\kappa_m, \kappa_e))\beta_6$, then there is inefficient termination with exploration. If $c_2/c_1 > \max(\kappa_m, \kappa_e)\beta_2 + (1 - \max(\kappa_m, \kappa_e))\beta_6$, then there is inefficient continuation with exploration.*

As shown in Proposition 6, 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 inefficient termination. For high values of c_2/c_1, the agent is inclined to exploit. Inefficient continuation allows the principal to pay the agent lower wages in the first period, offsetting the losses from inefficient continuation.

As shown in Corollary 2, there is inefficient 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.

7 Feedback

In this section, I study what happens if the principal is better able than the agent to evaluate performance. This is likely to be the case if there is a significant subjective component in the evaluation of performance or if the agent has little experience with the task. The focus of the section is on whether the principal should provide feedback on performance to the agent.

Feedback, or interim performance evaluation, has received little attention in the economics literature. In a setting where the principal's problem is to induce the agent to exert effort, Lizzeri, Meyer, and Persico (2002) and Fuchs (2004) find that it is optimal for the principal not to reveal information about performance to the agent.[9] Outside a principal-agent setting, Ray (2004) develops a model in which interim performance evaluation serves the purpose of screening bad projects. Here, the optimal provision of feedback will depend on whether the principal wants to implement exploration or exploitation.

I now 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 $\mathrm{IC}_{\langle i\,{}^{j}_{k}\rangle}$ 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 $\langle i\,{}^{j}_{k}\rangle$ 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.

[9]This result is closely related to the results of Abreu, Milgrom, and Pearce (1991) on the reusability of punishments.

The following definitions will be useful in stating Proposition 7:

$$\alpha_7 = \frac{2c_1}{p_1^2 - p_0^2},$$

$$\beta_7 = \frac{E[p_2]E[p_2|S,2] - p_0^2}{p_1^2 - p_0^2}.$$

Proposition 7 *To implement exploitation, it is optimal for the principal not to provide feedback on performance to the agent. The optimal contract that implements exploitation without feedback is such that*

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

$$w_{SF} = w_{FS} = \frac{(p_1 + p_0)c_1}{p_0(p_1 - E[p_2]) + E[p_2](E[p_2|S,2] - p_1)} \left(\beta_7 - \frac{c_2}{c_1}\right)^+,$$

and

$$w_{SS} = \alpha_7 - \frac{2(1 - p_1 - p_0)c_1}{p_0(p_1 - E[p_2]) + E[p_2](E[p_2|S,2] - p_1)} \left(\beta_7 - \frac{c_2}{c_1}\right)^+.$$

If the principal does not provide feedback, then incentive compatibility constraints associated with exploration, shirking in the second period in case of a success in the first period, and shirking in the second period in case of failure in the first period, which are binding when interim performance is publicly observable, can be ignored. Therefore, it is less costly to implement exploitation if information about interim performance is not revealed to the agent. The relevant incentive compatibility constraints are:

$$(p_1^2 - p_0^2)w_{SS} + (p_1(1-p_1) - p_0(1-p_0))w_{SF} + ((1-p_1)p_1 + (1-p_0)p_0)w_{FS} \geq 2c_1 \quad (\mathrm{IC}_{\langle 0\,{}^0_0 \rangle})$$

$$(p_1^2 - E[p_2]E[p_2|S,2])w_{SS} + (p_1(1-p_1) - E[p_2](1 - E[p_2|S,2]))w_{SF} + ((1-p_1)p_1 + (1-E[p_2])E[p_2|F,2])w_{FS} \geq 2(c_1 - c_2) \quad (\mathrm{IC}_{\langle 2\,{}^2_2 \rangle})$$

The optimal contract that implements exploitation without feedback has $w_{SS} \geq w_{SF} = w_{FS} \geq w_{FF} = 0$. If $c_2/c_1 > \beta_7$, then $\mathrm{IC}_{\langle 0\,{}^0_0 \rangle}$ is binding, and incentives are provided through w_{SS} only. If $c_2/c_1 < \beta_7$, then $\mathrm{IC}_{\langle 2\,{}^2_2 \rangle}$ is binding and $w_{SS} > w_{SF} = w_{FS} > 0$, since providing incentives only through w_{SS} could induce the agent to try the new work method. Figures 8 and 9 compare the optimal contracts and the costs to implement exploitation with and without feedback under the base case parameters.

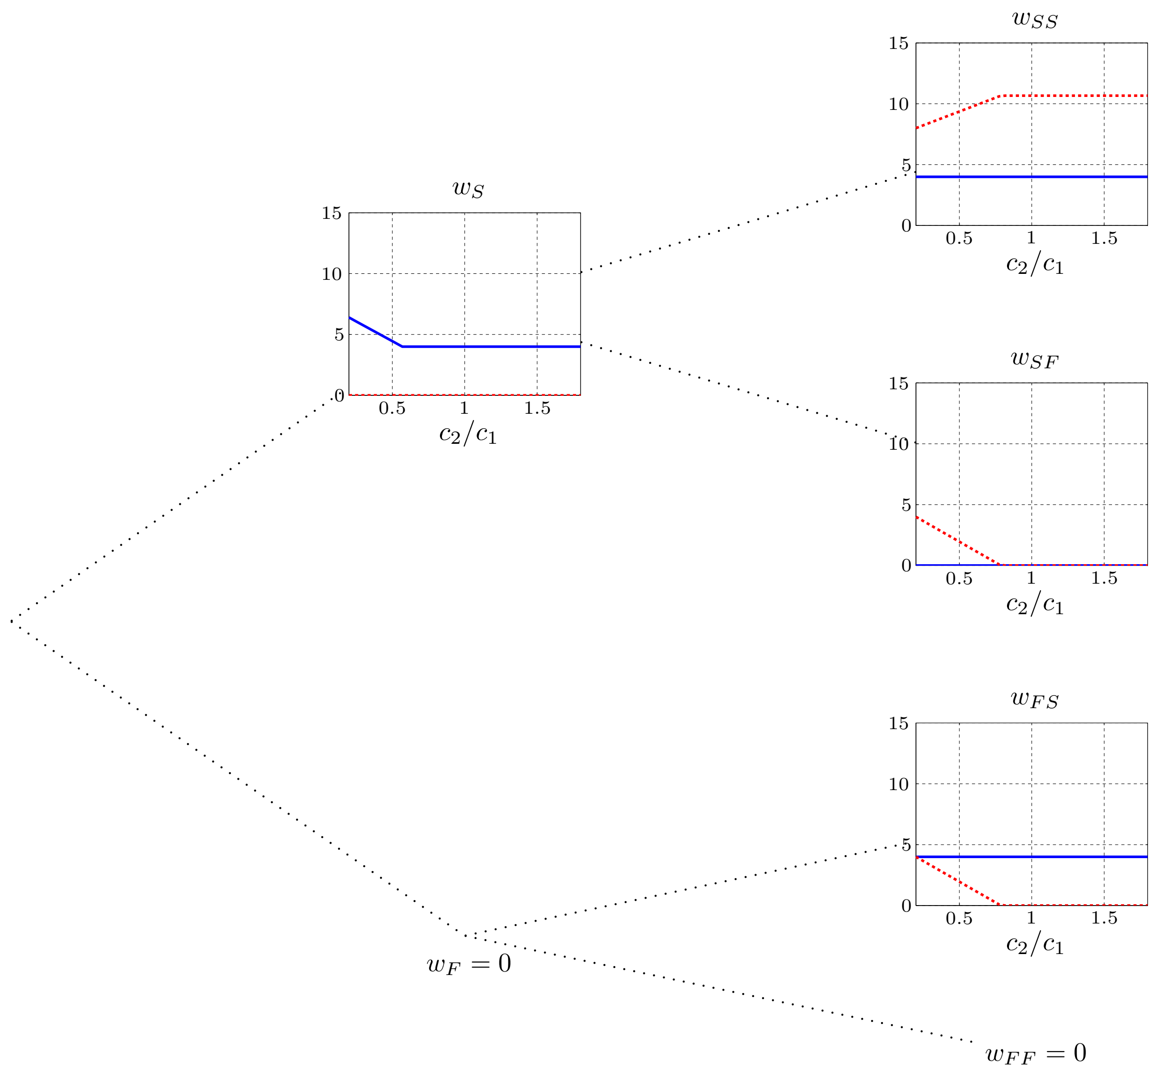

Figure 8: The optimal contract that implements exploitation with (solid line) and without (dashed line) feedback under the base parameters.

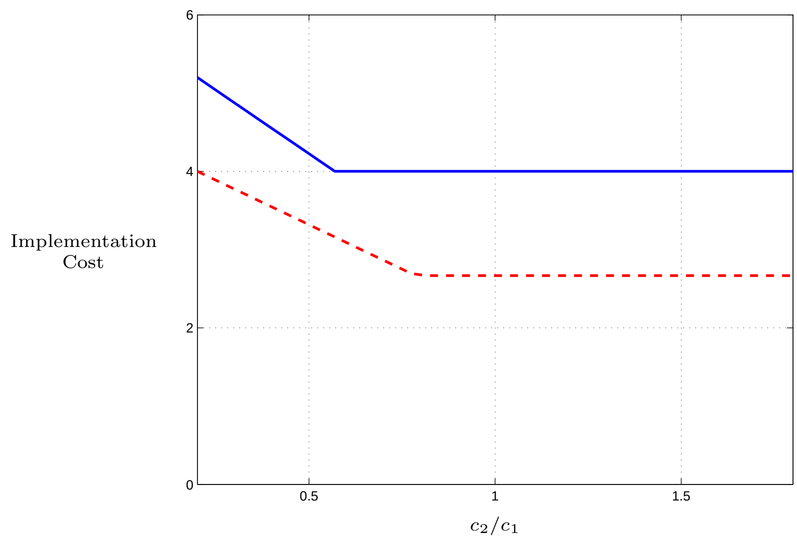

Figure 9: Cost of implementing exploitation with (solid line) and without (dashed line) feedback under the base parameters.

Proposition 8 *To implement exploration, the principal must provide feedback to the agent. The optimal contract $\vec{w}_8$ that implements exploration when the principal is better able than the agent to evaluate performance is the same as the optimal contract $\vec{w}_2$ that implements exploration derived in Proposition* 2.

Feedback is essential to implement exploration. It allows for rapid and efficient experimentation. 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 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.

8 Experimental Evidence in Psychology

A central tenet of economics is that pay-for-performance motivates the agent to exert more effort and consequently improves performance. Some empirical studies investigate the effect of pay-for-performance on effort provision. In a study with agricultural workers in the Philippines, Foster and Rosenzweig (1994) show that conditional on calories intake, workers on piece-rates lose more weight than workers on fixed salaries. Other empirical studies investigate the effect of pay-for-performance on performance. Lazear (2000) shows that the productivity of windshield installers of a large auto-glass company increased when their compensation structure changed from fixed-wage to piece-rates. Paarsch and Shearer (1999) find similar results for Canadian tree planters and Ehrenberg and Bognanno (1990) find similar results for professional golf players. There is thus empirical evidence that pay-for-performance motivates the agent to exert effort and consequently improves performance. The available evidence, however, is mainly for simple routine tasks, in which effort is the key input of the worker and there is not much room for exploration.

In the last decades, researchers in psychology started to challenge this view, providing experimental evidence that pay-for-performance may actually undermine performance in tasks that involve creativity. For example, Glucksberg (1962) asked subjects to mount a candle on a vertical screen, using only the screen, the candle, a book of matches, and a box of thumbtacks. To solve the problem, subjects must have figured out that they could use the box of thumbtacks not only as a container but also as a platform for the candle. Subjects who were offered no reward solved the problem significantly faster than subjects who were offered rewards.

McCullers (1978), McGraw (1978), Kohn (1993) and Amabile (1996) survey the subsequent research, which replicated the findings of Glucksberg (1962) for different tasks. The main conclusion of this line of research is that pay-for-performance improves performance in tasks that require simple, routine, unchanging responses, but can undermine performance in tasks that require flexibility, conceptual and perceptual openness, and creativity. Paying the agent in terms of current performance may make the agent inclined to exploit, because the agent thinks that the likelihood of success with a conventional work method is higher than the likelihood of success with a new work method. However, if the task requires exploration, such as the task in Glucksberg's experiment, then pay-for-performance may undermine performance, as pay-for-performance induces the agent to exploit and not to explore.

Future research may try to revisit these experiments with different compensation and feedback structures. It would be interesting to investigate the effects of rewards for early failure, long-term incentives, feedback and termination on creativity and performance in these experiments.

9 Applications

9.1 Management of Creative Workers

Modern corporations and business consultants have recognized the potential hazards of standard pay-for-performance contracts to innovation, and have devised alternative incentive schemes to motivate innovation.[10] For example, considered to be one of the most innovative corporations, 3M is well-known for having a corporate culture that provides freedom to experiment, tolerance for early failures, and rewards for successful innovations. William L. McKnight, who joined 3M in 1907 as an assistant bookkeeper and served as a chairman of

[10] Business books, such as Farson and Keyes (2002) and Sutton (2002), contain several examples of the alternative incentive schemes adopted by innovative corporations.

the board from 1949 to 1966, crafted the management principles that are still in place in the corporation. His basic rule of management was that:[11]

> Mistakes will be made. But if a person is essentially right, the mistakes he or she makes are not as serious in the long run as the mistakes management will make if it undertakes to tell those in authority exactly how they must do their jobs. Management that is destructively critical when mistakes are made kills initiative. And it's essential that we have many people with initiative if we are to continue to grow.

As shown in Proposition 4, the ability to commit to a long-term contract is essential to encourage exploration. Modern corporations often rely on their corporate culture to overcome this commitment problem and encourage exploration. Promises made in the form of a corporate culture are enforced through reputation. From Propositions 2 and 6, and Corollary 2, a corporate culture that tolerates failure and reward long-term success is optimal to motivate exploration.

Innovative organizations 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 wiling to explore new research directions that are likely to fail, but may lead to breakthroughs. From Corollary 2, by committing not to terminate researchers, research departments are able to motivate exploration.

Sometimes a corporation designates a few of its units for exploration. Thomke (2002) discusses the case of Bank of America, which assigned 25 of its branches to being 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 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." To make it clearer that failures are welcome, the top executives of Bank of America were considering an increase in the target failure rate from 30% to 40%.

In sum, anecdotal evidence suggests that business consultants and innovative corporations have realized that incentive schemes that motivate innovation are different than the standard pay-for-performance incentive schemes. It would be interesting to do a careful study comparing incentive schemes of different organizations or different divisions inside an organization. The model developed in this paper predicts that incentive schemes will depend on whether the assigned task requires more exploration or exploitation.

[11] `http://solutions.3m.com/wps/portal/_l/en_US/_s.155/123521/_s.155/123521`

9.2 Executive Compensation and Corporate Governance

In the context of the recent scandals in the American corporate sector, executive compensation has increasingly been criticized as excessive and not related to performance. This public outcry creates pressure for regulations that limit the use of stock options, golden parachutes, entrenchment, and option repricing.[12]

In an article on the state of U.S. corporate governance, Holmstrom and Kaplan (2003) warn of the risk of a regulatory overreaction:

> . . . an effort to regulate the system so that such outrage will never again occur would be overly costly and counterproductive. It would lead to inflexibility and fear of experimentation. In today's uncertain climate, we probably need more organizational experimentation than ever. The New Economy is moving forward and, in order to exploit the potential efficiencies inherent in the new information technologies, new business models and new organizational structures are likely to be desirable and valuable. Enron was an experiment that failed. We should take advantage of its lessons not by withdrawing into a shell, but rather by improving control structures and corporate governance so that other promising experiments can be undertaken.

In fact, it is easy to see from Propositions 2 and 6 that the optimal contracts that provide incentives for exploration and exploration with termination can be implemented by combinations of stock options, option repricing, golden parachutes, and entrenchment. Stock options provide long-term incentives while golden parachutes protect the agent against early failures. Option repricing not only induces the agent to exert effort in case his stock options goes out of the money, but also induces the agent to employ the new work method in the first period as it protects the agent against early failures. Finally, as shown in Corollary 2, inefficient continuation with exploration is sometimes optimal. One can interpret this as entrenchment, since the manager may keep his job even though it is ex post efficient for the firm to terminate the manager. These results suggest that regulations that limit the use of stock options, option repricing, golden parachutes, and entrenchment may just make it more difficult to motivate exploration, and can potentially have negative effects on innovation, and long-term growth.

Previous studies, such as Lambert (1986), Feltham and Wu (2001) and Lambert and Larcker (2004), have developed static models in which the optimal compensation that encourages risk-taking is convex, resembling a stock option. Other studies derived optimal contracts that, for different reasons than the one 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 (2005) show that stock options and golden parachutes may be optimal to induce the manager to reveal information

[12]See, for example, Bebchuk and Fried (2004) and "Rewards for Failure," *British DTI consultation*, June 2003.

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 the 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.

Executive compensation has recently been under criticism. The main criticisms are that compensation is too high and is not related to performance. These attacks create pressure for regulations that limit the use of golden parachutes, option repricing and entrenchment. In this paper, I show that combinations of stock options, option repricing, golden parachutes, and entrenchment are optimal to implement exploration. Limiting the use of these instruments may thus have adverse effects on innovation.

9.3 Bankruptcy Laws and Entrepreneurship

A recent debate involves bankruptcy laws in Europe and in the United States. On one hand, to encourage entrepreneurial activity, the European Council issued in June of 2000 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 the Member States and that new bankruptcy laws should determine a less severe treatment towards insolvent debtors. On the other hand, after eight years of discussion, the U.S. Congress passed in April of 2005 a new creditor-friendly bankruptcy law, the "Bankruptcy Abuse and Consumer Protection Act." The new law makes it more difficult for insolvent debtors to obtain exemptions and discharge of obligations. The main objective was to hold Americans more responsible for paying off their personal debt. Since many entrepreneurs finance their startups with personal debt, the new law may have unintended consequences for entrepreneurship.

The model developed in this paper provides a framework to analyze 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. 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. The creditor may allow the entrepreneur to keep working after default if it is still profitable to do so, but the entrepreneur earns just enough money so that he does not shirk. To sum up, a bankruptcy law that is based on the principle of fresh start is optimal to motivate exploration, while a bankruptcy law that is based on the principle of absolute priority is optimal to motivate exploitation.

A natural question to ask is why governments impose a single mandatory bankruptcy law. By considering the incentives for exploration, this paper provides a potential explanation for this question. Due to knowledge spillovers and imperfect intellectual-property-rights (IPR) protection, individuals involved in exploratory activities cannot fully appropriate the economic value generated by the knowledge they produce. As argued by Nelson (1959), this leads to under-exploration when compared to the socially efficient level of exploration. To alleviate the under-exploration problem, governments must impose a debtor-friendly bankruptcy laws instead of a menu of bankruptcy laws. Debtor-friendly bankruptcy laws protect the agent against failure and therefore induce exploration.

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 have found beneficial effects of deviations from absolute priority. For example, Bebchuk and Picker (1993), and Berkovitch, Israel, and Zender (1997, 1998) show that deviations of absolute priority may encourage investments in firm-specific versus general human capital. Baird (1991), Heinkel and Zechner (1993), Povel (1999), and Berkovitch and Israel (1998, 1999) show that deviations of absolute priority induce the entrepreneur to reveal private information to creditors when bad outcomes occur.[13] 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.

This subsection's interpretation of the model focuses on the effects of bankruptcy laws on entrepreneurship, innovation, and growth. It argues that social gains can be made from a bankruptcy law based on the principle of fresh start. By protecting entrepreneurs against failures, a bankruptcy law based on the principle of a fresh start encourages exploration, generating knowledge that may benefit society. Since entrepreneurship, innovation, and growth are major concerns for most countries, they should play an important role in the design of bankruptcy laws.

[13]Landier (2002) develops a model with multiple equilibria in which the stigma of failure may prevent entrepreneurs from abandoning bad projects. Gromb and Scharfstein (2003) and Hellman (2003) study regulatory and labor markets conditions under which innovation is developed inside or outside the firm.

10 Extensions

I assumed throughout the paper that the agent is risk-neutral and has limited liability. Results similar to the ones obtained here hold if the agent is risk-averse. The critical elements influencing the optimal contracts are the likelihood ratios between the different action plans, and 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 restricted 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 discussed in Section 2, 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 the high rate of failure in innovative projects as the fundamental difference between innovative and traditional projects. If this is indeed the case, 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, we will often see more often in practice rewards for failure than rewards for big successes.

I assumed throughout the paper that the agent could choose between the conventional work method, with known probability of success, and the new work method, with unknown probability of success. The results in this paper still hold if the agent can choose between a less innovative work method or a more innovative work method, as long as the agent does not learn too much when he employs the less innovative work method.

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 assumed in this paper that the principal observes only the output produced by the agent, and consequently, can use only output-based pay. In many situations, however, a noisy signal about the action taken by the agent is publicly observable. One can show that if the principal observes a noisy signal about the agent's action, 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.

I also assumed that the reservation utility of the agent is zero. It is possible, however, to imagine situations in which after a success with a new work method the worker has the option to leave the organization for a higher wage. Through deferred compensation, the optimal contract that implements exploration reduces the incentives for the agent to leave the organization.

11 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. Because of that, 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, inefficient continuation, commitment to a long-term incentive plan, and timely feedback on performance are all important ingredients to motivate exploration.

Practices such as the institution of tenure, golden parachutes, 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, limiting their use may have adverse effects on 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 a new issue: what is the optimal menu of contracts that separates creative workers from shirkers and conventional workers? An additional extension to consider is the case of multiple agents. This raises another issue: what is the optimal balance between individual and team incentives to motivate exploration? I leave these issues for future research.

The ideas presented here can also be tested empirically. One natural place to start is the experimental evidence in psychology on the detrimental effects of reward on performance. Revisiting these experiments with different reward structures seems to be a promising direction of research. It would be interesting to see the effects of reward for early failure, long-term incentives, feedback and termination on creativity and performance in these experiments.

12 Appendix

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

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

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

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

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

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

I now argue that some incentive compatibility constraints are redundant. If $(i,j) \neq (1,1)$, then it follows from $\text{IC}_{\langle {}_1{}^1_0 \rangle}$ and $\text{IC}_{\langle {}_i{}^j_1 \rangle}$ that $\text{IC}_{\langle {}_i{}^j_0 \rangle}$ are redundant. If $(i,k) \neq (1,1)$, then it follows from $\text{IC}_{\langle {}_1{}^0_1 \rangle}$ and $\text{IC}_{\langle {}_i{}^1_k \rangle}$ that $\text{IC}_{\langle {}_i{}^0_k \rangle}$ are redundant. If $\langle {}_i{}^j_k \rangle \neq \langle {}_2{}^2_1 \rangle$ 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 $\text{IC}_{\langle {}_i{}^j_k \rangle}$ is redundant. Rewriting the incentive compatibility constraints that are not redundant:

$$(p_1 - p_0)w_{SS} \geq c_1 \quad (\text{IC}_{\langle {}_1{}^0_1 \rangle})$$

$$(p_1 - p_0)w_{FS} \geq c_1 \quad (\text{IC}_{\langle {}_1{}^1_0 \rangle})$$

$$(p_1 - p_0)w_S + (p_1^2 - p_0p_1)w_{SS} - (p_1^2 - p_0p_1)w_{FS} \geq c_1 \quad (\text{IC}_{\langle {}_0{}^1_1 \rangle})$$

$$\begin{aligned}(p_1 - E[p_2])w_S + (p_1^2 - E[p_2]E[p_2|S,2])w_{SS} - (p_1^2 - E[p_2]p_1)w_{FS} \geq& \\ c_1 - c_2 + E[p_2](c_1 - c_2)& \quad (\text{IC}_{\langle {}_2{}^2_1 \rangle})\end{aligned}$$

I now show that $\text{IC}_{\langle {}_1{}^0_1 \rangle}$ and $\text{IC}_{\langle {}_1{}^1_0 \rangle}$ 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 $\vec{w}'$ be the same as $\vec{w}$ except that $w'_{SS} = w_{SS} - \Delta_1$, $w'_S = w_S + p_1\Delta_1$, $w'_{FS} = w_{FS} - \Delta_2$, and $w'_F = w_F + p_1\Delta_2$. The contract $\vec{w}'$ satisfies the above constraints, $C(\vec{w}', \langle 1\,{}^{1}_{1}\rangle) = C(\vec{w}, \langle 1\,{}^{1}_{1}\rangle)$, and $\vec{w}'$ pays the agent earlier than $\vec{w}$. The incentive compatibility constraints $\text{IC}_{\langle 2\,{}^{2}_{1}\rangle}$ and $\text{IC}_{\langle 0\,{}^{1}_{1}\rangle}$ become

$$(p_1 - p_0)w_S \geq c_1 \qquad (\text{IC}_{\langle 0\,{}^{1}_{1}\rangle})$$

$$\begin{aligned}(p_1 - E[p_2])w_S + E[p_2](E[p_2|S,2] - p_1)&\frac{c_1}{p_1 - p_0} \\ &\geq c_1 - c_2 + E[p_2](c_1 - c_2) \qquad (\text{IC}_{\langle 2\,{}^{2}_{1}\rangle})\end{aligned}$$

If $c_2/c_1 \geq \beta_1$ then $\text{IC}_{\langle 0\,{}^{1}_{1}\rangle}$ is binding. Otherwise, $\text{IC}_{\langle 2\,{}^{2}_{1}\rangle}$ is binding. ■

Proof of Proposition 2: The optimal contract $\vec{w}$ that implements action plan $\langle 1\,{}^{1}_{1}\rangle$ satisfies the following incentive compatibility constraints:

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

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 $\text{IC}_{\langle i\,{}^{j}_{k}\rangle}$ and $C(\vec{w}', \langle 1\,{}^{1}_{1}\rangle) < C(\vec{w}, \langle 1\,{}^{1}_{1}\rangle)$. 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 $\text{IC}_{\langle i\,{}^{j}_{k}\rangle}$ and $C(\vec{w}', \langle 2\,{}^{2}_{0}\rangle) < C(\vec{w}, \langle 2\,{}^{2}_{0}\rangle)$. 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 $\text{IC}_{\langle i\,{}^{j}_{k}\rangle}$ are still satisfied, $C(\vec{w}', \langle 2\,{}^{2}_{0}\rangle) = C(\vec{w}, \langle 2\,{}^{2}_{0}\rangle)$, and the contract $\vec{w}'$ pays the agent earlier than $\vec{w}$.

If follows from $\text{IC}_{\langle 2\,{}^{2}_{0}\rangle}$ and $\text{IC}_{\langle i\,{}^{j}_{1}\rangle}$ that $\text{IC}_{\langle i\,{}^{j}_{0}\rangle}$ and $\text{IC}_{\langle i\,{}^{j}_{2}\rangle}$ are redundant. From $\text{IC}_{\langle 2\,{}^{2}_{0}\rangle}$, we have that $w_{FS} \geq \frac{c_1}{p_1 - p_0}$ and $\text{IC}_{\langle i\,{}^{j}_{1}\rangle}$ implies $\text{IC}_{\langle i\,{}^{j}_{0}\rangle}$. Since $c_2 \geq \frac{E[p_2] - p_0}{p_1 - p_0} c_1$, $\text{IC}_{\langle i\,{}^{j}_{1}\rangle}$ implies $\text{IC}_{\langle i\,{}^{j}_{2}\rangle}$.

Rewriting the incentive compatibility constraints that are not redundant:

$$(p_1 - p_0)(w_{FS} - w_{FF}) \geq c_1 \qquad (\text{IC}_{\langle 2\,{}^{2}_{0}\rangle})$$

$$\begin{aligned}(E[p_2]E[p_2|S,2] - p_1 E[p_j])w_{SS} + (p_1 - E[p_2])w_F& \\ + ((1 - E[p_2])p_1 - (1 - p_1)p_0)w_{FS}& \\ \geq (c_2 + E[p_2]c_2 + (1 - E[p_2])c_1) - (c_1 + p_1 c_j)& \qquad (\text{IC}_{\langle 1\,{}^{j}_{1}\rangle})\end{aligned}$$

$$
\begin{aligned}
(E[p_2]E[p_2|S,2] - p_0E[p_j])w_{SS} - (E[p_2] - p_0)w_F& \\
+ ((1 - E[p_2])p_1 - (1 - p_0)p_0)w_{FS}& \\
\geq (c_2 + E[p_2]c_2 + (1 - E[p_2])c_1) - p_0c_j& \quad (\mathrm{IC}_{\langle 0\,{}^{j}_{1} \rangle})
\end{aligned}
$$

$$
(E[p_2|S,2] - E[p_j])w_{SS} \geq c_2 - c_j. \qquad (\mathrm{IC}_{\langle 2\,{}^{j}_{1} \rangle})
$$

The incentive compatibility constraint $\mathrm{IC}_{\langle 2\,{}^{2}_{0} \rangle}$ 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 $\mathrm{IC}_{\langle i\,{}^{j}_{1} \rangle}$ are still satisfied, $C(\vec{w}', \langle 2\,{}^{2}_{1} \rangle) = C(\vec{w}, \langle 2\,{}^{2}_{1} \rangle)$, and the contract $\vec{w}'$ pays the agent earlier than $\vec{w}$. On the other hand, the incentive compatibility constraints $\mathrm{IC}_{\langle 1\,{}^{2}_{1} \rangle}$, $\mathrm{IC}_{\langle 1\,{}^{0}_{1} \rangle}$, $\mathrm{IC}_{\langle 2\,{}^{1}_{1} \rangle}$, and $\mathrm{IC}_{\langle 2\,{}^{0}_{1} \rangle}$ are redundant. If $c_2 \geq c_1$, $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle}$ implies $\mathrm{IC}_{\langle 1\,{}^{2}_{1} \rangle}$, and if $c_2 < c_1$, $\mathrm{IC}_{\langle 1\,{}^{2}_{1} \rangle}$ is trivially satisfied. Also, $\mathrm{IC}_{\langle 0\,{}^{1}_{1} \rangle}$ and $\mathrm{IC}_{\langle 0\,{}^{1}_{2} \rangle}$ imply $\mathrm{IC}_{\langle 1\,{}^{0}_{1} \rangle}$. Moreover, $\mathrm{IC}_{\langle 0\,{}^{2}_{1} \rangle}$ implies $\mathrm{IC}_{\langle 2\,{}^{0}_{1} \rangle}$. Finally, $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle} + \frac{p_1 - E[p_2]}{E[p_2]-p_0}\mathrm{IC}_{\langle 0\,{}^{1}_{1} \rangle}$ implies $\mathrm{IC}_{\langle 2\,{}^{1}_{1} \rangle}$.

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, $\mathrm{IC}_{\langle 0\,{}^{1}_{1} \rangle}$ implies $\mathrm{IC}_{\langle 0\,{}^{0}_{1} \rangle}$ and $\mathrm{IC}_{\langle 0\,{}^{2}_{1} \rangle}$. Either $w_F > 0$, and $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle}$ and $\mathrm{IC}_{\langle 0\,{}^{1}_{1} \rangle}$ are binding or $w_F = 0$ and $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle}$ is binding. When $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle}$ and $\mathrm{IC}_{\langle 0\,{}^{1}_{1} \rangle}$ 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 $\mathrm{IC}_{\langle 0\,{}^{j}_{1} \rangle}$ and $\mathrm{IC}_{\langle 2\,{}^{2}_{0} \rangle}$ are binding, $w_{SS} = w_{SS}^{0j1}$, and $w_F = 0$, where

$$
\begin{aligned}
j \in \arg\max_{\tilde{j} \in \{0,1\}} \quad w_{SS}^{0\tilde{j}1} \equiv \frac{(1 + E[p_2])c_2 - p_0c_{\tilde{j}}}{(E[p_2]E[p_2|S,2] - p_0E[p_{\tilde{j}}])}& \\
+ \frac{(E[p_2] - p_0)p_0\frac{c_1}{p_1-p_0}}{(E[p_2]E[p_2|S,2] - p_0E[p_{\tilde{j}}])}&
\end{aligned}
$$

I first prove that the candidate contract is feasible. For that it suffices to show that $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle}$ is satisfied. If $E[p_2]E[p_2|S,2] \geq p_1^2$, then $\mathrm{IC}_{\langle 0\,{}^{1}_{1} \rangle}$ implies $\mathrm{IC}_{\langle 1\,{}^{1}_{1} \rangle}$. 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 to that, $\text{IC}_{\langle 0\,{}^{j}_{1}\rangle}$ is not satisfied for any $w_{SS} < w_{SS}^{0j1}$. Therefore, it is impossible to improve on the candidate contract. ■

Proof of Proposition 3: Follows from the fact that $\text{IC}_{\langle 1\,{}^{0}_{1}\rangle}$ and $\text{IC}_{\langle 1\,{}^{1}_{0}\rangle}$ are binding under the optimal long-term contract. ■

Proof of Proposition 4: In order to implement $\langle 2\,{}^{2}_{1}\rangle$, the following incentive compatibility constraints must be satisfied:

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

Moreover, for the contract to be renegotiation-proof, we must have $j,k\in\mathcal{I}$ such that $\text{IC}_{\langle 2\,{}^{j}_{1}\rangle}$ and $\text{IC}_{\langle 2\,{}^{2}_{k}\rangle}$ bind.

If $c_2 \geq \frac{E[p_2|S,2]-p_0}{p_1-p_0}c_1$, from $\text{IC}_{\langle 2\,{}^{1}_{1}\rangle}$ we have that

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

$$w_{SF}=0.$$

This contradicts $\text{IC}_{\langle 1\,{}^{1}_{1}\rangle}+\frac{p_1-E[p_2]}{E[p_2]-p_0}\text{IC}_{\langle 0\,{}^{1}_{1}\rangle}$. Therefore, $\langle 2\,{}^{2}_{1}\rangle$ 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 $\text{IC}_{\langle 2\,{}^{0}_{1}\rangle}$ we have that

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

$$w_{SF}=0.$$

From $\text{IC}_{\langle 2\,{}^{2}_{k}\rangle}, k\in\{0,2\}$ 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:

$$w_S \geq \frac{c_2(1-p_0)}{E[p_2]-p_0}+p_0\frac{c_1}{p_1-p_0} \qquad (\text{IC}_{\langle 0\,{}^{2}_{1}\rangle})$$

$$\begin{aligned}w_S \geq \frac{c_2(E[p_2|S,2]-p_0(1-(p_1-E[p_2])))}{(E[p_2|S,2]-p_0)(E[p_2]-p_0)}\\ -\frac{c_1p_0(E[p_2|S,2]-p_0)}{(E[p_2|S,2]-p_0)(E[p_2]-p_0)}+p_0\frac{c_1}{p_1-p_0}\quad (\text{IC}_{\langle 0\,{}^{1}_{1}\rangle})\end{aligned}$$

$$w_S \geq \frac{c_2(E[p_2|S,2] - p_0(1+E[p_2]) + p_0^2)}{(E[p_2|S,2]-p_0)(E[p_2]-p_0)} + p_0\frac{c_1}{p_1-p_0} \qquad (\mathrm{IC}_{\langle {}_{0}{}^{0}_{k}\rangle})$$

It is easy to show that, given $c_2 < \frac{E[p_2|S,2]-p_0}{p_1-p_0}c_1$, $\mathrm{IC}_{\langle {}_{0}{}^{0}_{1}\rangle}$ implies $\mathrm{IC}_{\langle {}_{0}{}^{1}_{1}\rangle}$ and $\mathrm{IC}_{\langle {}_{0}{}^{2}_{1}\rangle}$. Therefore, from $\mathrm{IC}_{\langle {}_{0}{}^{0}_{1}\rangle}$, 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

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

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

Proof of Proposition 5: Similar to the proof of Proposition 1. ■

Proof of Corollary 1: Comparing the costs of implementing exploitation and exploitation with termination from the contracts derived in Propositions 1 and 5, one obtains that: $W(\vec{w}_1, \langle {}_{1}{}^{1}_{1}\rangle) - W(\vec{w}_5, \langle {}_{1}{}^{1}_{t}\rangle) = (1-p_1+p_0)p_1\alpha_1 > (1-p_1)p_1\alpha_1$. There is inefficient termination with exploitation. ■

Proof of Proposition 6: Similar to the proof of Proposition 2. ■

Proof of Corollary 2: Follows from comparing the costs of implementing exploration and exploration with termination from the contracts derived in Propositions 2 and 6. If $c_2/c_1 > \max(\kappa_m, \kappa_e)\beta_2 + (1-\max(\kappa_m,\kappa_e))\beta_6$, then $W(\vec{w}_2, \langle {}_{1}{}^{1}_{1}\rangle) - W(\vec{w}_6, \langle {}_{1}{}^{1}_{t}\rangle) > (1-E[p_2])p_1\alpha_2$ and there is inefficient continuation with exploration. If $c_2/c_1 < \max(\kappa_m, \kappa_e)\beta_2 + (1-\max(\kappa_m,\kappa_e))\beta_6$, then $W(\vec{w}_2, \langle {}_{1}{}^{1}_{1}\rangle) - W(\vec{w}_6, \langle {}_{1}{}^{1}_{t}\rangle) < (1-E[p_2])p_1\alpha_2$ and there is inefficient termination with exploration. ■

Proof of Proposition 7: For the no feedback policy to have any effect the principal must set $w_S = w_F$, or otherwise the agent can infer the output in the first period from the first period wages. Therefore, the optimal contract $\vec{w}$ that implements action plan $\langle {}_{1}{}^{1}_{1}\rangle$ without feedback must have $w_S = w_F$ and satisfy the following incentive compatibility constraints:

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

for all $\langle {}_{i}{}^{j}_{k}\rangle \neq \langle {}_{1}{}^{1}_{1}\rangle$ with $j = k$.

First, I show that $w_S = w_F = 0$. Suppose $w_S = w_F > 0$. A contract $\vec{w}'$ that is the same as $\vec{w}$ but has $w_S = w_F = 0$ satisfies all the above constraints and has $C(\vec{w}', \langle 1\,{}^{1}_{1}\rangle) < C(\vec{w}, \langle 1\,{}^{1}_{1}\rangle)$. Next, I show that $w_{FF} = 0$. Suppose that $w_{FF} > 0$. A contract $\vec{w}'$ that is the same as $\vec{w}$ but has $w_{FF} = 0$ satisfies all the above constraints and has $C(\vec{w}', \langle 1\,{}^{1}_{1}\rangle) < C(\vec{w}, \langle 1\,{}^{1}_{1}\rangle)$.

Since $c_2/c_1 \geq (E[p_2]-p_0)/(p_1-p_0)$, $\text{IC}_{\langle 0\,{}^{1}_{1}\rangle}$ and $\text{IC}_{\langle 0\,{}^{0}_{0}\rangle}$ imply $\text{IC}_{\langle 0\,{}^{2}_{2}\rangle}$. Similar arguments can be used to show that $\text{IC}_{\langle 2\,{}^{1}_{1}\rangle}$, $\text{IC}_{\langle 1\,{}^{2}_{2}\rangle}$ and $\text{IC}_{\langle 2\,{}^{0}_{0}\rangle}$ are redundant.

The remaining incentive compatibility constraints can be written as

$$(p_1^2 - p_0^2)w_{SS} + (p_1(1-p_1) - p_0(1-p_0))w_{SF} + (p_1(1-p_1) - p_0(1-p_0))w_{FS} \geq 2c_1 \quad (\text{IC}_{\langle 0\,{}^{0}_{0}\rangle})$$

$$(p_1^2 - E[p_2]E[p_2|S,2])w_{SS} + (p_1(1-p_1) - E[p_2](1-E[p_2|S,2]))w_{SF} + (p_1(1-p_1) - (1-E[p_2])E[p_2|F,2])w_{FS} \geq 2(c_1 - c_2) \quad (\text{IC}_{\langle 2\,{}^{2}_{2}\rangle})$$

$$(p_1^2 - p_0p_1)w_{SS} + (1-p_1)(p_1-p_0)w_{SF} - p_1(p_1-p_0)w_{FS} \geq c_1 \quad (\text{IC}_{\langle 0\,{}^{1}_{1}\rangle})$$

$$(p_1^2 - p_0p_1)w_{SS} - p_1(p_1-p_0)w_{SF} + (1-p_1)(p_1-p_0)w_{FS} \geq c_1 \quad (\text{IC}_{\langle 1\,{}^{0}_{0}\rangle})$$

Using Bayes' rule the incentive compatibility constraint $\text{IC}_{\langle 2\,{}^{2}_{2}\rangle}$ can be written as

$$(p_1^2 - E[p_2]E[p_2|S,2])w_{SS} + (p_1(1-p_1) - E[p_2](1-E[p_2|S,2]))w_{SF} + (p_1(1-p_1) - E[p_2](1-E[p_2|S,2]))w_{FS} \geq 2(c_1 - c_2) \quad (\text{IC}_{\langle 2\,{}^{2}_{2}\rangle})$$

I now show that we can restrict attention to contracts that have $w_{SF} = w_{FS}$. Suppose $w_{SF} \neq w_{FS}$. Therefore, a contract $\vec{w}'$ that has $w'_{SF} = w'_{FS} = (w_{SF} + w_{FS})/2$ satisfies all of the above incentive compatibility constraints and has $C(\vec{w}, \langle 1\,{}^{1}_{1}\rangle) = C(\vec{w}', \langle 1\,{}^{1}_{1}\rangle)$.

The candidate for the optimal contract is the one in the statement of Proposition 7. If $c_2/c_1 \geq \beta_7$ then only $\text{IC}_{\langle 0\,{}^{0}_{0}\rangle}$ is binding. If $c_2/c_1 < \beta_7$, then both $\text{IC}_{\langle 0\,{}^{0}_{0}\rangle}$ and $\text{IC}_{\langle 2\,{}^{2}_{2}\rangle}$ are binding. One can check that $\text{IC}_{\langle 1\,{}^{0}_{0}\rangle}$ and $\text{IC}_{\langle 0\,{}^{1}_{1}\rangle}$ are satisfied under the optimal contract. ■

Proof of Proposition 8: Action plan $\langle 2\,{}^{2}_{1}\rangle$ can only be implemented if the principal provides feedback on performance to the agent. ■

References

Abreu, D., P. Milgrom, and D. Pearce, 1991, Information and Timing in Repeated Partnerships, *Econometrica* 59, 1712–1733.

Acharya, V., K. John, and R. Sundaram, 2000, On the Optimality of Resetting Executive Stock Options, *Journal of Financial Economics* 57, 65–101.

Aghion, P., O. Hart, and J. Moore, 1992, The Economics of Bankruptcy Reform, *Journal of Law, Economics, and Organizations* 8, 523–546.

Aghion, P., and J. Tirole, 1994, The Management of Innovation, *Quarterly Journal of Economics* 109, 1185–1209.

Almazan, A., and J. Suarez, 2003, Entrenchment and Severance Pay in Optimal Governance Structures, *Journal of Finance* 58, 519–548.

Amabile, T., 1996, *Creativity in Context: Update to the Social Psychology of Creativity.* (Westview Press Boulder, CO).

Arrow, K., 1969, Classificatory Notes on the Production and Diffusion of Knowledge, *American Economic Review* 59, 29–35.

Ayotte, K., 2007, Bankruptcy and Entrepreneurship: The Value of a Fresh Start, *Journal of Law, Economics and Organization* 23.

Baird, D., 1991, The Initiation Problem in Bankruptcy, *International Review of Law and Economics* 11, 223–232.

Bebchuk, L., and J. Fried, 2004, *Pay Without Performance.* (Harvard University Press Cambridge, MA).

Bebchuk, L., and R. 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.

Bennis, W., and B. Nanus, 1997, *Leaders: Strategies for Taking Charge.* (Harper Business New York).

Bergemann, D., and J. Valimaki, 2006, Bandit Problems, in S. Durlauf, and L. Blume, eds.: *The New Palgrave Dictionary of Economics* (Macmillan Press,).

Berkovitch, E., and R. Israel, 1998, The Bankruptcy Decision and Debt Contract Renegotiations, *European Finance Review* 2, 1–27.

Berkovitch, E., and R. Israel, 1999, Optimal Bankruptcy Laws Accross Different Economic Systems, *Review of Financial Studies* 12, 347–377.

Berkovitch, E., R. Israel, and F. Zender, 1997, An Optimal Bankruptcy Law and Firm-Specific Investments, *European Economic Review* 41, 487–497.

Berkovitch, E., R. Israel, and F. Zender, 1998, The Design of Bankruptcy Laws: A Case for Management Bias in Bankruptcy Reorganizations, *Journal of Financial and Quantitative Analysis* 33, 441–464.

Berry, D., and B. Fristedt, 1985, *Bandit Problems: Sequential Allocation of Experiments.* (Chapman and Hall).

Bolton, P., and C. Harris, 1999, Strategic Experimentation, *Econometrica* 67, 349–374.

Dewatripont, M., and E. Maskin, 1995, Credit and Efficiency in Centralized and Decentralized Economics, *Review of Economic Studies* 62, 541–555.

Ehrenberg, R., and M. Bognanno, 1990, Do Tournaments Have Incentive Effects?, *Journal of Political Economy* 98, 1307–1324.

Farson, R., and R. Keyes, 2002, *Whoever Makes the Most Mistakes Wins: The Paradox of Innovation.* (The Free Press New York).

Feltham, G., and M. Wu, 2001, Incentive Efficiency of Stock Versus Options, *Review of Accounting Studies* 6, 7–28.

Foster, A., and M. Rosenzweig, 1994, A Test for Moral Hazard in the Labor Market: Contractual Arrangements, Enforcement, and Health, *Review of Economics and Statistics* 76, 213–227.

Fuchs, W., 2004, Contracting with Repeated Moral Hazard and Private Evaluations, Working paper, Stanford Graduate School of Business.

Fudenberg, D., B. Holmstrom, and P. Milgrom, 1990, Short-Term Contracts and Long-Term Agency Relationships, *Journal of Economic Theory* 51, 1–31.

Garicano, L., 2000, Hierarchies and the Organization of Knowledge in Production, *Journal of Political Economy* 108, 874–904.

Glucksberg, S., 1962, The Influence of Strength of Drive on Functional Fixedness and Perceptual Recognition, *Journal of Experimental Psychology* 63, 36–41.

Gromb, D., and D. Scharfstein, 2003, Entrepreneurship in Equilibrium, Working paper, London Business School.

Harris, M., and A. Raviv, 1978, Some Results on Incentive Contracts with Applications to Education, Insurance, and Law Enforcement, *American Economic Review* 68, 20–30.

Heinkel, R., and J. Zechner, 1993, Financial Distress and Optimal Capital Structure Adjustments, *Journal of Economics, Management, and Strategy* 2, 531–565.

Hellman, T., 2003, When do employees become entrepreneurs, Working paper, Stanford Graduate School of Business.

Hermalin, B., and M. Katz, 1991, Moral Hazard and Verifiability: The Effects of Renegotiation in Agency, *Econometrica* 59, 1735–1753.

Holmstrom, B., 1979, Moral Hazard and Observability, *Bell Journal of Economics* 10, 74–91.

Holmstrom, B., 1989, Agency Costs and Innovation, *Journal of Economic Behavior and Organization* 12, 305–327.

Holmstrom, B., and S. Kaplan, 2003, The Dangers of Too Much Governance, *Sloan Management Review* 45, 96.

Holmstrom, B., and P. Milgrom, 1991, Multi–Task Principal–Agent Analysis, *Journal of Law, Economics, and Organizations* 7, 24–52.

Inderst, R., and H. Mueller, 2005, Keeping the Board in the Dark: CEO Compensation and Entrenchment, Working paper, New York University.

Jensen, M., 1991, Corporate Control and the Politics of Finance, *Journal of Applied Corporate Finance* 4, 13–33.

Jovanovic, B., and Y. Nyarko, 1996, Learning by Doing and the Choice of Technology, *Econometrica* 64, 1299–1310.

Kohn, A., 1993, *Punished by Rewards: The Trouble with Gold Stars, Incentive Plans, A's, Praise, and Other Bribes.* (Houghton Mifflin Company New York).

Lambert, R., 1986, Executive Effort and Selection of Risky Projects, *The Rand Journal of Economics* 17, 77–88.

Lambert, R., and D. Larcker, 2004, Stock Options, Restricted Stock, and Incentives, Working paper, The Wharton School, University of Pennsylvania.

Landier, A., 2002, Entrepreneurship and the Stigma of Failure, Working paper, Graduate School of Business, The University of Chicago.

Lazear, E., 1986, Salaries and Piece Rates, *The Journal of Business* 59, 405–431.

Lazear, E., 2000, Performance Pay and Productivity, *American Economic Review* 90, 1346–1361.

Lee, D., and E. Van den Steen, 2006, Managing Know-How, Working paper, MIT Sloan School of Management.

Lizzeri, A., M. Meyer, and N. Persico, 2002, The Incentive Effects of Interim Performance Evaluations, Working paper, CARESS.

March, J., 1991, Exploration and Exploitation in Organizational Learning, *Organization Science* 2, 71–87.

McCullers, J. C., 1978, Issues in Learning and Motivation, in M. R. Lepper, and D. Greene, eds.: *The Hidden Costs of Reward* (Lawrence Erlbaum Associates, Hillsdale, NJ).

McGraw, K. O., 1978, The Detrimental Effects of Reward on Performance: A Literature Review and a Prediction Model, in M. R. Lepper, and D. Greene, eds.: *The Hidden Costs of Reward* (Lawrence Erlbaum Associates, Hillsdale, NJ).

Nelson, R., 1959, The Simple Economics of Basic Scientific Research, *Journal of Political Economy* 67, 297–306.

Paarsch, H., and B. Shearer, 1999, The Response of Worker Effort to Piece Rate: Evidence from the British Columbia Tree-Planting Industry, *Journal of Human Resources* 34.

Povel, P., 1999, Optimal "Soft" or "Tough" Bankruptcy Procedures, *Journal of Law, Economics and Organization* 15, 659–684.

Ray, K., 2004, Performance Evaluation and Efficient Sorting, Working paper, Stanford Graduate School of Business.

Romer, P., 1986, Increasing Returns and Long-Run Growth, *Journal of Political Economy* 94, 1002–1037.

Schumpeter, J., 1934, *The Theory of Economic Development.* (Harvard University Press Cambridge, MA).

Stiglitz, J. E., and A. Weiss, 1983, Incentive Effects of Termination: Applications to the Credit and Labor Markets, *American Economic Review* 73, 912–927.

Sutton, R., 2002, *Weird Ideas that Work: 11 1/2 Practices for Promoting, Managing, and Sustaining Innovation.* (The Free Press New York).

Thomke, S., 2002, Bank of America, *Harvard Business Review* 9–603–022 and 9–603–023.

Von Thadden, E., 1995, Long–Term Contracts, Short–Term Investment and Monitoring, *Review of Economic Studies* 62, 557–575.