

Search Agency ^{*}

Tracy R. Lewis[†]

Marco Ottaviani[‡]

July 2008

Preliminary Draft

Abstract

When delegating a sequential search activity to an agent, a principal faces two information problems: the agent's discovery effort cannot be observed (hidden action) and the agent's progress is private knowledge (hidden information). We characterize optimal incentive provision under monitored search (wherein the principal monitors the agent's progress during the search process) and delegated search (wherein the principal delegates decisions to an agent). We demonstrate how the organization of search activities responds to variations in the agent's costs and private benefits as search progresses to completion. Our findings provide a rationale for contracts that are observed in practice.

JEL Classification: D83 (Search; Learning; Information and Knowledge), M52 (Compensation and Compensation Methods and Their Effects).

^{*}We thank Patrick Bolton, Wouter Dessein, Luis Garicano, Ulrich Hege, Rocco Macchiavello, Gustavo Manso, Andrea Prat, John Sutton, Jeroen Swinkels, Tim Van Zandt, Peter Zemsky, and seminar participants at Berkeley, Duke, ESSET 2008, HEC, LSE, Montreal, Queens, USC, and Washington University for helpful comments.

[†]Fuqua School of Business, Duke University. E-mail: tracyl@duke.edu.

[‡]Kellogg School of Management, Northwestern University. E-mail: m-ottaviani@northwestern.edu.

1 Introduction

When searching for a new technology or manufacturing process, firms frequently hire researchers to search for the best alternative. The development of a new drug to treat hypertension, the design of a new security to hedge currency risks, or the development of a new technology to improve production or service are all examples of tasks that are often delegated to agents specializing in doing research and finding solutions. This paper studies how best to contract with agents to search and do research to discover new technology and methods for production.

A principal who hires an agent to search must overcome two information obstacles. First, there is a hidden action problem in so far as the agent's effort in discovery is nearly impossible to observe. Second, the principal is unlikely to know what discoveries have actually resulted from the agent's search, thus leading to a hidden information problem. Hence, the principal must acquire information about the agent's progress to know what additional work remains to be done as well as how much further effort is needed to complete this work. These information problems frequently lead to cost overruns, holdup, and delays in delivery of desired products and services. The procurer's ability to solve these agency problems turns on whether he can observe and monitor the agent's actions and commit to incentive payments and terms that affect the agent's performance.

Logically, there are two approaches to organize an agency search for the best alternative. First, when hiring an agent to work in a supervised setting, the principal is able to monitor the agent's progress. By appropriately rewarding the agent's successes, the principal can effectively control the agent's amount of effort. While monitored search may work well in principle, it may be costly to implement if discoveries are difficult to evaluate. A second alternative is for the principal to delegate search decisions to an independent agent. Under delegated search the agent works at her own rate to discover new solutions, and decides if and when to disclose her findings to the principal. Delegated search minimizes the principal's cost of acquiring information, but makes it more difficult to manage the agent in the process.

In this paper we characterize the optimal dynamic agency contract and structure for implementing a search for the best alternative. Our analysis considers how the agent is evaluated and rewarded, the duration of her employment, how the agent is managed as the

search progresses and how to tailor contracts to different settings where search is monitored or delegated. For our analysis we develop a continuous-time, sequential-search model for a setting in which a risk-neutral agent or (agents) are employed by a risk-neutral principal to discover a new technology or process. The principal offers the agent a sequence of short term contracts, which may be renegotiated to the mutual advantage of both parties after each period. A contract stipulates how the agent is evaluated and paid and the conditions for termination due to poor performance. The agent is wealth constrained and is therefore unable to post a bond or acquire an equity share in the project to insure her performance. Hence the setting we examine is one of a long term relationship governed by a series of short term agreements in which the parties can only commit to actions that are individually rational at each stage going forward.

At each instant, the agent chooses a level of search effort. Higher effort levels increase the chance of drawing a discovery from a known distribution. The agent is allowed to recall past discoveries and at each point in time she retains the highest discovery achieved so far. Discoveries have two potential effects. First, they may provide surplus to the principal conditional on adoption. Second, discoveries provide surplus to the agent, who can search at lower marginal cost in the future and enjoy non-pecuniary and private benefits from the knowledge he has acquired¹.

The presence of these benefits that accumulate with each discovery is fundamental to the agency relationship and how it changes as the search progresses. Whenever a discovery decreases the cost of effort for the future, the agent's incentives to search in the future to complete the search increase. ~~Since~~ discovery increasing incentives reduce the rewards required to induce the agent to perform. On the other hand, when discovery increases the agent's accumulated private benefits from the search, his incentives to search further are reduced, since the remaining benefits from making another discovery are decreased. Such discovery decreasing incentives increase the rewards that the agent must receive to induce him to continue searching. While these benefits would appear to be desirable by-products of agency search, they may complicate the management of the agent for some types of search.

For instance under monitored search, the principal can measure and reward the agent's progress to induce a desired level of effort. The compensation consists of a share of the

¹See Stern (2004) for evidence on the importance of such benefits for researchers.

surplus created after the search is finished. However the principal's inability to commit long term constrains the contract provisions he can credibly offer. We find that contract terms can only depend only on the current state of the search and not on the agent's previous performance. A threat to terminate the agent for non performance or to reduce payment for late discovery is not credible. Moreover after observing an increasing incentives discovery, the principal is forced to ratchet down future discovery payments, realizing the agent searches more intensely as his effort costs decline. In contrast the principal increase rewards for further discoveries after an decreasing incentives discovery, because he knows the agent is less interested in searching further. The ability of the principal to monitor enables the principal to address any conflicts of interest that might arise as the search progresses. Nonetheless, the agency costs of monitored search will cause the principal to search slower and less extensively as compared to the first-best search he undertakes working for himself.

When the principal is unable to monitor results he must delegate search decisions to the privately informed agent who observes the state of the search at each instant. The delegated search contract differs in important ways from the agreements described above. Under delegation, the contract is based on the agent's disclosures and the principal's evaluation of her progress. The agent may disclose all or a part of a discovery. The inability to track the agent's progress limits the adjustments in payments the principal can implement. Surprisingly, however, the principal can implement the monitored search program through delegation when discoveries increase incentives for further search. In this setting the principal pays the agent for any disclosure that will cause future payment to decrease. The agent is paid after each discovery, unlike monitored search where he receives a share of the total search surplus. Although the timing of payments differs, the agent's expected compensation is the same under monitored and delegated search. Hence the rate and extent of discovery is identical for both searches.

In contrast it is not possible to replicate monitored search through delegation when discoveries decrease incentives. Under delegation the agent increases future payments to the agent after a discovery since her incentives to perform decrease. Hence the agent has an incentive to delay the disclosure of search-ending discovery, to compel the principal to increase search incentives. As a result we find that the agent will gradually disclose his discoveries to gain higher rewards. In turn the principal will wait to increase payments until he is convinced that greater incentives are required. A form of Coasian conjecturing

about the agent's type exists here. The result of this two-sided holdup is that the disclosure of discoveries is delayed and the agent's payment for discoveries is driven up.

Our analysis is related to a well-established and growing literature on information and agency theory. This literature, including Freixas, Laffont, and Tirole (1988), Hart and Tirole (1988), and Laffont and Tirole (1993), attempts to trace out the impacts of information flows on optimal contracts to manage dynamic adverse selection problems. Our analysis demonstrates how ratcheting and gradual disclosure of progress that are common in agency relationships with limited commitment also exist in agency search. Moreover the "type" of agent is shown to evolve endogenously in our analysis of sequential search.

The advantage of monitoring in a dynamic setting that we address is related to earlier analyses of monitoring and auditing that include Baron and Besanko (1984), and Mookherjee and Png (1989). The distinction between monitoring and self disclosure that we emphasize is also addressed in the law and economics papers by Kaplow and Shavell (1994), and more generally by Shin (1994) and Lewis and Poitevin (1997). The issue of delegating decisions to better informed agents has been studied in various contexts by Lewis and Sappington (1997), Aghion and Tirole (1997), and Cremer, Khalil, and Rochet (1998). Our analysis extends these studies to a dynamic setting where an agent's expertise pertains to what future actions should be taken.

This paper is also related to a growing literature on optimal dynamic contracting in recursive programming models of agency relationships (Green, 1987, Spear and Srivastava, 1987, and Atkeson, 1991). Recent analyses in nance of the optimal capital structure in dynamic agency models including Fishman and DeMarzo (2007), as well as dynamic agency models of the firm including Albuquerque and Hopenhayn (2004), Clementi and Hopenhayn (2006), focus on the impact of agency relationships on capital structure and firm performance.

With respect to optimal search processes virtually all of the extensive literature abstracts from the agency problem on which we focus, with a few exceptions². This work on optimal unemployment insurance (see Shavell and Weiss, 1979, and Hopenhayn and Nicolini, 1997) concentrates on the tradeoff between risk sharing and incentives for search

²In the real estate literature, some papers have recognized the importance of adding agency considerations to the sequential search model, but have not characterized the solution to the problem. The closest contribution in that literature is Arnold (1989), who argues that the first best outcome is not a solution to the problem once agency considerations are introduced. Here, we provide a full characterization of the second best solution.

in a repeated moral hazard setting. In contrast our analysis is focused on the interaction of the agent's incentives to exert search effort with agent's incentives to report the private information the agent acquires during the search process. Also, in our model the conflict of interest between principal and agent evolves as the relationship progresses, due to the private benefits that result from new discoveries. These dynamic aspects are absent in the optimal unemployment insurance literature.

Our analysis is also related to the literature on learning in the presence of agency distortions. For instance Neher (1999) and Bergemann and Hege (2005) focus on the problems created by the limited commitment generated by inability of the principal to commit not to renegotiate ex-post suboptimal contractual arrangements.³ These issues are addressed by our analysis in the context of sequential search.

The paper proceeds as follows. Section 2 formulates the model and reports the first-best solution in the absence of agency problems. Section 3 analyzes monitored search. Section 4 characterizes the solution under delegated search. Section 5 concludes.

2 Direct Search

This section introduces our basic model of search. The model is first described in discrete time and subsequently transformed into continuous time for easier computations. As a benchmark for that analysis we characterize optimal direct search that is carried out by the principal.

2.1 Model of Sequential Search

A risk neutral principal seeks to discover a technology represented by $v \in [\underline{v}, \bar{v}] \subset R_+$ where $0 \leq \underline{v} < \bar{v} < \infty$. The principal adopts the technology after discovery to produce a good yielding a flow surplus of $f(v)$. This surplus represents the value of a good or process embedded with technology. The principal may also derive an interim benefit of $w(v)$ from the technology as it is adopted to improve his production process as the time of discovery. The final and interim surpluses are assumed to be smooth increasing functions of v , with $\omega_v, w_v \geq 0$ for all v .

We model discovery as a sequential search process. Technologies, drawn with

³See also Dewatripont and Maskin (1995) and Cornelli and Yosha (2003).

recall and replacement from a known and stationary distribution function with density $f(v) = F'(v) > 0$ for $v \in [\underline{v}, \bar{v}]$.⁴ The principal exerts effort $e \geq 0$, to make random draws from $F(v)$ at the rate of $\phi(e)$. We assume that $\phi(e) = e^\gamma$ for $\gamma \in (0, 1)$. Although this specification is not necessary for our results, it does simplify the analysis to follow. Given $\phi(e)$, the probability of drawing a technology exceeding a time period $\Delta > 0$ is $\phi(e) \tilde{F}(v) \Delta$ where $\tilde{F}(v) = 1 - F(v)$ is the probability of drawing a technology exceeding

There is a variable cost of searching $c(v)e$. We assume that unit cost $c(v)$ declines at the rate of $c'_v \leq 0$ with improved technology. This arises from process innovation that makes search more efficient. Aside from search costs, there is a benefit that accrues to the searcher as a result of the search process. This benefit reflects value derived from learning and discovering the technology that the searcher may use in other related tasks. This benefit is assumed to be increasing with the technology at the rate θ . Define $C(e, v) = c(v)e - g(v)$.

2.2 Direct Search by Principal

Let $W_t^p(v)$ denote the principals expected search surplus starting at technology v at time t . The discount factor is $\delta = \frac{1}{1+r\Delta}$ where $\Delta > 0$ is the length between periods and $r > 0$ is the rate of discount. At each instant, the principal may terminate search and receive a perpetual flow of surplus valued $\frac{\Delta\omega(v)}{1-\delta}$, or continue searching to find a better technology. The principals surplus, is then recursively defined in the following dynamic program

$$W_t^p(v) = \max_{e_t^p} \left\{ \begin{array}{l} \delta W_{t+\Delta}(v) + \Delta (g(v) - w(v) - c(v) e_t^p) \\ + \delta \Delta \phi(e_t^p) \int_{\underline{v}}^{\bar{v}} [W_{t+\Delta}^p(v') - W_{t+\Delta}^p(v)] dF(v') \end{array} \right\} \quad (1)$$

The continuation surplus consists of next periods surplus, plus the net surplus flow multiplied by the period duration Δ plus the expected appreciation in surplus.

Under standard conditions, McCall (1970) shows that the solution to (1) is characterized by a unique technology level $\hat{v} \in (\underline{v}, \bar{v})$ such that

$$W_t^p(v) \begin{cases} > \\ \leq \end{cases} \frac{\Delta\omega(\hat{v})}{1-\delta} \text{ for } \begin{cases} v < \hat{v} \\ v \geq \hat{v}, \end{cases} \quad (2)$$

⁴We abstract from issues of learning about the distribution of values. While such issues are important for some applications, they are not critical for our analysis.

⁵Costs may also decline with learning by doing as one learns more about the technology.

so that it is optimal to terminate search for \hat{v}^p .⁶ To incorporate the stopping condition (2) in the problem above and to allow for a continuous time transformation of the analysis, we rewrite (1) in terms of Δ so that

$$W_t^p(v) = \max_{e_t^p} \left\{ \frac{1}{1+r\Delta} W_{t+\Delta}(v) + \Delta (g(v) + w(v) - c(v) e_t^p(v)) \right. \\ \left. + \frac{1}{1+r\Delta} \Delta \phi(e_t^p) \int_v^{\hat{v}^p} [W_{t+\Delta}^p(v') - W_{t+\Delta}^p(v)] dF \right\} \\ + \frac{1}{1+r\Delta} \Delta \phi(e_t^p) \int_{\hat{v}^p}^{\bar{v}} \left(\frac{\Delta \omega(v')}{r\Delta} - W_{t+\Delta}^p(v) \right) dF(v') \quad (3)$$

After multiplying both sides of (3) by $1+r\Delta$ and dividing both sides by Δ and rearranging terms we obtain

$$rW_t^p(v) = \frac{W_{t+\Delta}^p(v) - W_t^p(v)}{\Delta} + \max_{e_t^p} \left\{ (1+r\Delta) (g(v) + w(v) - c(v) e_t^p(v)) \right. \\ \left. + \phi(e_t^p) \int_v^{\hat{v}^p} [W_{t+\Delta}^p(v') - W_{t+\Delta}^p(v)] dF \right\} \\ + \phi(e_t^p) \int_{\hat{v}^p}^{\bar{v}} \left(\frac{\omega(v')}{r} - W_{t+\Delta}^p(v) \right) dF(v'). \quad (4)$$

Notice that $W_{t+\Delta}^p(v) - W_t^p(v) = 0$ as W_t^p does not depend explicitly on time. Hence taking the limit as $\Delta \rightarrow 0$ in (4) we obtain:

$$rW^p(v) = \max_{e^p} (w(v) + g(v) - c(v) e^p(v)) + \phi(e^p(v)) \left(\int_v^{\hat{v}^p} [W^p(v') - W^p(v)] dF \right. \\ \left. + \int_{\hat{v}^p}^{\bar{v}} \left(\frac{\omega(v')}{r} - W^p(v) \right) dF(v') \right). \quad (5)$$

Equation (5) makes clear that the principals search problem in continuous time (as well as discrete time) is stationary and depends only on the current technology. While search is ongoing, optimal effort $e^p(v)$ is characterized by

$$-c(v) + \phi'(e^p) \left(\int_v^{\hat{v}^p} [W^p(v') - W^p(v)] dF \right. \\ \left. + \int_{\hat{v}^p}^{\bar{v}} \left[\frac{\omega(v')}{r} - W^p(v) \right] dF \right) = 0.$$

⁶The condition is that there are diminishing returns from search for sufficiently high technologies. For our setting this requires that at

$$W^p(\hat{v}^p) = \frac{\omega(\hat{v}^p)}{r}$$

and

$$W_v^p(\hat{v}^p) < \frac{\omega_v(\hat{v}^p)}{r}$$

These conditions are satisfied whenever v is sufficiently large. This implies that terminal benefits from discovery are increasing sufficiently fast with technology.

In general optimal effort will vary as the costs and benefits from search change with the current technology. The principal searches optimally until a critical discovery is achieved whereby gains from future search are exhausted. This occurs at the critical technology where

$$W^p(\hat{v}^p) = \frac{\omega(\hat{v}^p)}{r}.$$

These results are summarized in the following:

Proposition 0: Under direct search discovery proceeds at the optimal rate until technology level \hat{v}^p is discovered:

(i) Under optimal discovery: $e^p(v) = \begin{cases} \arg \max_e W^p(v) & \text{for } v \leq \hat{v}^p \\ 0 & \text{for } v > \hat{v}^p. \end{cases}$

(ii) Stopping point exhaust gains from search: $W^p(\hat{v}^p) - \frac{\omega(\hat{v}^p)}{r} = 0.$

(iii) Effort changes with v at the rate $e_v = \frac{-c_v e^{\frac{(r+\phi\tilde{F}-\gamma\phi\tilde{F})}{\gamma}} - (g_v+w_v)\phi\tilde{F}}{c(v)e(v)(r+\phi\tilde{F})}$

Proof: The proof follows from standard arguments and is thus omitted. All other formal results not proved in the body of the paper appear in the Appendix.

Proposition 0 outlines some important properties of optimal search with implications for the analysis to follow. The first implication follows from property (iii) which indicates that search intensity is adjusted with each new innovation. When search costs decline, the rate of effort increases. This suggests under monitored search that the agent's progress should be checked regularly to make desired adjustments in his incentives to discover. Moreover under delegated search, the agent must be induced to make his own adjustments based on his private knowledge of discovery. A second implication is the agent may benefit from intermediate discoveries that reduce his cost of future search and increase his private benefits from search. This may cause dynamic conflicts of interest between the agent and principal who may value innovation differently. This conflict may present a problem for the principal particularly when search is delegated to an agent.

3 Monitored Search

In this section we analyze contracting for search where the agent's progress can be monitored and publicly observed. We refer to this as monitored search and compare it with

delegated search where monitoring is not possible in Section 4. The monitoring contract provisions are first described in Section 3.1. Section 3.2 defines the renegotiation proof subgame perfect and Markov perfect equilibrium for our setting. We conjecture that the two equilibria are equivalent so that we may restrict attention to Markov Equilibria in our analysis. Accordingly we solve for the unique Markov Perfect equilibrium and describe the properties of the agency search contract in Section 3.3. Finally we confirm the equivalence between the Markov and weakly renegotiation proof equilibrium in Section 3.4

3.1 Monitored Search Contracts

Under monitored search the principal is unable or unwilling to conduct the search. Instead he hires one or several workers from an infinite homogenous pool of agents to search on her behalf.⁷ Each agent has the same search technology, and the same search flow costs and benefits $c(v)$ and $g(v)$ as the principal. Hence, monitored search differs only from direct search in that the work is performed by an agent whose effort can not be observed or contracted on directly.

Agents are risk neutral and have an outside opportunity wage that is normalized to zero. Moreover agents are wealth constrained and can not make positive payments to the principal either to post a performance bond or to purchase an equity share in the firm. The agents' effort at searching can't be observed or verified. Therefore the agent must be induced to search by receiving positive rewards from the principal for discoveries that she makes.

The principal hires one agent at a time to search. We assume the parties are unable to commit to long term agreements. Rather the agency relationship is governed by a sequence of short term contracts that are revised after each period. The inability of the principal to commit arises because firm managers rotate frequently between different assignments within the firm or between different firms and consequently are not able to oversee and enforce long term agreements that may bind their successors to a possibly non optimal policy.

The interaction between the agent and principal proceeds as follows. At the beginning of each period, the principal decides to continue searching or to quit. If he continues,

⁷We assume the principal is only capable of monitoring one agent at a time. The analysis can be easily be extended to consider the simultaneous monitoring of numerous agents.

he offers a one period contract denoted by $\{\tau_t, \mu_t\}$ to the incumbent agent. Under the contract, the agent is promised a reward $\tau_t \geq 0$ and a probability $\mu_t \in [0, 1]$ of being terminated at the end of the period, conditional on the technology she discovers. Assuming the agent accepts the contract (which she always does) she allocates $e_t \geq 0$ to search. The results of search are then realized and the agent is paid and either terminated or retained depending on the discovery she makes. If the agent is terminated, the principal pays $S > 0$ to locate another agent for next period.

Formally the strategies of the players are described as follows. The set consists of all public histories h_t , that includes the set of events and decisions that are observed by both players prior to time t , where a history h_t is given by

$$h_t = \{\kappa_1, \dots, \kappa_{t-1}; v_1, \dots, v_{t-1}; \tilde{\mu}_1, \dots, \tilde{\mu}_{t-1}\}.$$

The history consists of the sequence of previous contract terms $\{\kappa_{t'}, \mu_{t'}\}$, the realization of technologies $v_{t'}$ and $\tilde{\mu}_{t'}$ the termination or retention decision following discovery. By convention, the technology $v_{t'}$ is set equal to 0 whenever the agent fails to draw a sample from $F(v)$. Note that the quitting decision is omitted from the history as it may be inferred by whether the search is still on going or not.

The principals contract strategy for payment is specified by

$$\tau_t : H_t \times v_t \rightarrow R_+$$

The termination strategy is a randomized rule for releasing the current agent that maps from the history and discovery outcome into a decision to retain (retain) or terminate ($\mu_t = 1$),

$$\mu_t : H_t \times v_t \rightarrow \Delta\{0, 1\},$$

where $\mu_t(h_t, v_t) \in [0, 1]$ signifies the probability of termination. The agent's effort allocation is given by

$$e_t : H_t \times \kappa_t \rightarrow R_+.$$

Finally, after the realization of e_t and the agent is either retained or terminated, the principal decides whether to quit or continue searching. This quitting strategy is represented

⁸The agent can also observe the effort she allocates each period. The principal is unable to observe effort. This difference in observable decisions is immaterial to our analysis, provided the outcome of the search is commonly observed.

by

$$q_{t+1} : H_t \times v_t \times \tilde{\mu}_t \rightarrow \{0, 1\},$$

where $q_t = 0$ indicates the search is terminated and 1 means the search continues.

3.2 Equilibrium

In the setting we focus on, the principal and agent repeatedly negotiate contract terms to govern the agents search. The contracts are of short duration and they are renegotiated after each period. Under monitored search, information about the outcome of the previous periods search is public whereas information on the effort devoted to discovery is privately known by the agent. For a given set of strategies $\{v_t, e_t, q_t\}_{t=0}^{\infty}$, we denote the value function for the principal at time the beginning of period t by $W^m(h_t)$ and the surplus function for the incumbent agent by $\Pi^m(h_t)$. In equilibrium, the game progresses through a sequence of decisions with corresponding continuation values for the principal denoted by $W^m(\kappa_t | h_t)$, and $W^m(q_{t+1}, \tilde{\mu}_t, v_t, \kappa_t | h_t)$ and $\Pi^m(e_t, \kappa_t | h_t)$ for the agent.

Definition: A subgame perfect equilibrium (SPE) is a strategy profile

$$\{\kappa_t^m, e_t^m, \rho_t^m, q_t^m\}_{t=1}^{\infty}$$

such that for all h_t

$$\begin{aligned} W^m(\kappa_t^m | h_t) &\geq W^m(\kappa_t | h_t) && \text{for all } \kappa_t \\ \Pi(e_t^m, \kappa_t | h_t) &\geq \Pi(e_t, \kappa_t | h_t) && \text{for all } e_t \\ W^m(q_{t+1}^m, \tilde{\mu}_t, v_t, \kappa_t | h_t) &\geq W^m(q_{t+1}, \tilde{\mu}_t, v_t, \kappa_t | h_t) && \text{for all } q_{t+1}. \end{aligned}$$

The three sets of optimality conditions reflect the contract decision, the effort allocation and the quit or continue decision that occur in sequence within and across each period. A particularly appealing subset of subgame perfect equilibrium, are those that are mutually preferable to the parties at each stage of the game. Such equilibrium are renegotiation proof in the sense that there does not exist an alternative pair of subgame perfect continuation strategies that the players would mutually prefer to adopt. This is an important characteristic of equilibrium for our setting where parties are able to renegotiate any arrangement to their mutual advantage in between periods. A particular refinement of renegotiation proof equilibria, due to Farrell and Maskin,(89) requires that renegotiations be time consistent in the sense that any renegotiation that was feasible at one subgame is

also feasible at another subgame that begins from the same (payo/relevant) history. For our setting, the payo/relevant history of the game is the highest technology discovered prior to t . More formally we define the renegotiation refinement by,

Definition: A subgame perfect equilibrium $\{\kappa_t^m, e_t^m, \rho_t^m, q_t^m\}_{t=1}^\infty$ is weakly renegotiation proof provided that for all histories h_t and h'_t such that $v_{h_t} = v_{h'_t}$ there is no continuation equilibrium such that $(W^m(h_t), \Pi^m(h_t)) \geq (W^m(h'_t), \Pi^m(h'_t))$ with at least one holding with strict inequality.

One particular SPE is the Markov perfect equilibrium where strategies only depend on the payo/relevant history of the game. This equilibrium is both appealing to employ because it conditions strategies on simple (and payo/relevant) histories of play and because it is simpler to compute than other subgame perfect equilibria with strategies depending on arbitrary histories of play. Formally Markov Perfect Equilibrium are characterized by,

Definition: A Markov perfect equilibrium (MPE) is a subgame perfect equilibrium

$$\{\kappa_t^m, e_t^m, \mu_t^m, q_t^m\}_{t=1}^\infty$$

such that for all different histories h_t and h'_t where $v_{h_t} = v_{h'_t}$ and all κ_t, q_t ,

$$\begin{aligned} \kappa_t^m(h_t) &= \kappa_t^m(h'_t) \\ e_t^m(h_t, \kappa_t) &= e_t^m(h'_t, \kappa_t) \\ q_{t+1}^m(h_t, \max[v_{h_t}, v], \tilde{\mu}_t) &= q_{t+1}^m(h'_t, \max[v_{h'_t}, v], \tilde{\mu}_t). \end{aligned}$$

In the analysis to follow we shall show Markov perfect equilibria are equivalent to weakly renegotiation proof subgame perfect equilibrium in our setting. This allows us to restrict attention to MPE in which strategies depend only on the payo/relevant history of play. Such equilibria are easy to compute and straightforward to describe.

3.3 Characterization of Monitored Search

In this section we provide an informal derivation of the optimal agency contract for the setting in which search discoveries are monitored. Following the Grossman and Hart (1983) methodology for analyzing moral hazard, it is useful to first calculate the cost of inducing agents to supply a given level of effort. Once the agency cost of effort supply is known we can then determine what effort maximizes the principals expected surplus under monitored search.

⁹Grossman and Hart's (1983) analysis of static moral hazard is conveniently separated into two stages. In the first stage they calculate the cost of inducing any desired effort. In stage two they select the

3.3.1 Costs of Inducing Agents to Search

Suppose that at the current technology the principal wishes to induce effort $e(v)$ with the contract $\kappa_t(v)$. The contract provides for a payment $\tau_t(v', v)$ and a termination probability of $\mu_t(v', v)$ for a discovery of v' when the agent starts at v . Given $\kappa_t(v)$ the agent solves

$$\begin{aligned} \Pi_t^m(v) = & \max_{e(v)} -\Delta C(e(v), v) + & (6) \\ & \delta \Delta \phi(e(v)) \int_v^{\bar{v}} \Pi_{t+\Delta}^m(\max(v', v)(1 - \mu_t(v', v)) + \tau_t(v', v)) dF(v') \\ & \delta(1 - \phi(e(v), \Delta))(1 - \mu_t(0, v)) \Pi_{t+\Delta}^m(v). \end{aligned}$$

It is apparent from (6) that the agent's reward can without loss of generality be paid only upon completion of the search. This follows because the agent is risk neutral and therefore cares only about the expected value of the payment she receives. When payment is made only upon completion the principal can reduce rewards at interim stages without violating wealth constraints, provided that the terminal payment is positive. Moreover the principal avoids any unnecessary payments to agents who may eventually be terminated. Assuming the principal stops search once \hat{v}^m we can represent the terminal payment by

$$F(\hat{v}^m) \tau_t(v) = \int_v^{\bar{v}} \tau_t(v', v) dF(v').$$

In this case the agent's effort choice satisfies

$$c(v) = \delta \phi'(e) \left(\begin{array}{l} \int_v^{\hat{v}^m} \Pi^m(\max(v', v)(1 - \mu(v', v))] dF(v') \\ - (1 - \mu(0, v)) \Pi^m(v) + \tilde{F}(\hat{v}) \tau(v) \end{array} \right). \quad (7)$$

where we have dropped the subscript from all expressions as the solution is independent of time. Solving for the payment $\tau(v)$ that induces $e(v)$ from (7) we obtain,

$$\tau(v) = \frac{1}{\tilde{F}(\hat{v}^m)} \left(\begin{array}{l} (1 - \mu(0, v)) \Pi^m(v) - \int_v^{\hat{v}^m} \Pi^m(\max(v', v)(1 - \mu(v', v)) dF \\ + \frac{c(v)e(v)}{\gamma \delta \phi(e)} \end{array} \right). \quad (8)$$

The principal's total cost of inducing $e(v)$ is

$$C^p(e(v)) = \min_{\{\mu(v', v)\}} \left(\begin{array}{l} S \Delta \phi(e) \int_v^{\hat{v}^m} \mu(v', v) dF(v') \\ + (1 - \Delta \phi(e)) \mu(0, v) + \tau(v) \Delta \phi(e) \tilde{F}(\hat{v}^m) \end{array} \right).$$

optimal effort that maximizes the principal's surplus given the costs of effort. We apply this approach to our analysis of dynamic moral hazard.

This cost is strictly increasing in (v', v) , which implies it is never optimal to terminate an agent who has just made a discovery. The rationale is that payments required by the agent to perform are decreasing in the likelihood she will be retained once a discovery is made. In effect extending the agent's tenure is a reward for discovery that substitutes for a monetary payment. Hence an agent's tenure should be most secure after a discovery.

In contrast, the principal's payment to the agent declines with (v) . This is evident from totally differentiating (8) to obtain $d\mu(0, v)/d\mu(0, v) = \frac{-\Pi(v)}{\bar{F}(\hat{v}^m)} < 0$. Intuitively, the threat of termination is a penalty the agent incurs whenever she fails to perform. This penalty enables the principal to maintain the same incentives to perform with smaller rewards. On balance an increase in (v) changes the principal's total cost at the rate of

$$\frac{dC^p}{d\mu(0, v)} = -\Delta\phi(e)\Pi + (1 - \Delta\phi(e))S.$$

If $-\Delta\phi(e)\Pi + (1 - \Delta\phi(e))S < 0$ costs are reduced by threatening termination. Note however, as we approach continuous time with $\Delta \rightarrow 0$ that $\frac{dC^p}{d\mu(0, v)} \rightarrow S > 0$ so that termination increases total costs. The rationale for this finding is that as the length between periods shrinks, the likelihood of a discovery within one period becomes arbitrarily small. The threat of termination for non performance can not induce the agent to work harder since she is unlikely to discover a new technology in the next instant no matter how much she searches.

Summarizing our findings to this point and providing a characterization of the cost minimizing search process in continuous time we have:

Proposition 1: Any feasible effort sequence $\{e(v)\}$ is implemented at least cost to the principal by a contract: $\kappa^m \equiv \{\tau^m(v), \mu^m(v', v)\}$ with these features:

- (i) The contract is stationary depending only on the current technology v .
- (ii) Agents are paid a terminal fee $\tau^m(v)$ that may be an equity share of the principal's terminal surplus once search is completed.
- (iii) Agents are never terminated.
- (iv) Agents earn rent $\Pi(v) = \frac{c(v)e(v)(1-\gamma)+\gamma g(v)}{\gamma r}$ for $v \leq \hat{v}^m$.

3.3.2 Optimal Monitored Search

The principal selects an effort sequence $\{e^m(v)\}$ for $v \leq \hat{v}^m$, and a terminal technology $\{a^m(v)\}$ subject to the constraint that $\{a^m(v)\}$ be implementable by $\{\tau^m(v)\}$, the least cost payment schedule, in order to solve the following continuous-time dynamic program

$$rW^m(v) = \max_{e^m(v), \hat{v}^m} w(v) + \phi(e^m) \left[\int_v^{\hat{v}^m} (W^m(v') - W^m(v)) dF + \int_{\hat{v}^m}^{\bar{v}} \left(\frac{\omega(v)}{r} - W^m(v) - \tau^m(v) \right) dF \right]. \quad (9)$$

The principal's expected surplus appreciates at the rate at which new intermediate and terminal technologies are discovered. The optimal payment $\tau^m(v)$ for completing the search is set by the condition

$$\frac{dW^m}{de^m} \frac{de^m}{d\tau^m} + \frac{dW^m}{d\tau^m} = 0. \quad (10)$$

Thus, we have

$$\frac{dW^m}{de^m} \frac{de^m}{d\tau^m} = \phi(e^m) \tilde{F}(\hat{v}^m) > 0.$$

Equation (10) implies that the amount of effort induced is distorted below the surplus maximizing level, at which the principal's surplus is maximized. This is done to reduce the agency rents accruing to the privately informed agent, similar to the distortions that arise in static agency settings. This suggests that the rate of search is likely to be slower under monitored search than under direct search where the principal is unconcerned about reducing agency costs.

The stopping rule for terminating search is given by the condition

$$\frac{\omega(\hat{v}^m)}{r} - W^m(\hat{v}^m) = 0.$$

The benefits from further search are exhausted at \hat{v}^m . Searching beyond \hat{v}^m yields negative returns.¹⁰ It seems intuitive that the search will be stopped at a lower technology under monitored search as compared to direct search. This is because the surplus generated under monitored search is less than under direct search because of the agency costs of search. Hence the returns from continuing to search are smaller when an agent conducts the search.

¹⁰It is easily verified that there is a unique e^m satisfying the stopping rule whenever $\omega(v)$ is sufficiently large

As the search progresses and new technologies are discovered the incentives for the principal and agent to continue searching will change in predictable ways. To gain some insight about this process, consider how a discovery affects the principals behavior on the margin. As the technology improves, the principals optimal $e^m(v)$ payment as characterized by (10) adjusts so as to maintain a constant level of variable search costs such that

$$\frac{dc(v) e^m(v)}{dv} = c_v e^m(v) + e_v^m c(v) = 0.$$

Hence if marginal search costs falls with a new discovery so that the principal can induce a moderate increase in effort to speed the rate of discovery, while keeping the agents rent $\Pi(v) = \frac{(1-\gamma)c(v)e^m(v)+\gamma g(v)}{\gamma}$ at a constant level. However when search costs fall the agent will want to increase effort above the principals desired level. This will call for a reduction in the payment the agent receives to dampen his incentives to search. Hence in situations where discovery increases the agent's incentives to search, the principal will respond by reducing payments for new discovery.

In contrast, suppose that a new discovery has no affect on marginal cost, but rather it increases the private benefit $w(v)$ that the agent enjoys from searching. This arises for instance when the agent can benefit more from working with a more advanced technology. As a result of this the agents incentives to search further will decrease as the remaining benefits to be captured fall with each new discovery. In this situation discovery decreases the agent's incentives to search the principal will respond by increasing payments for new discovery in order keep the agents effort at the same level.

The following Proposition summarizes the preceding arguments and characterizes the optimal search when agents can be monitored.

Proposition 2: There is a unique MPE for monitored search. The optimal program is implemented by terminal payments, $\tau^m(v) \geq 0$ with these properties:

- (i) Discovery proceeds at the optimal rate $e^m(v) = \arg \max_e W^m(v)$ until technology level \hat{v}^m is discovered: At \hat{v}^m further gains from search are exhausted with $\frac{\omega(\hat{v}^m)}{r} = W^m(\hat{v}^m)$.
- (ii) Payments $\tau^m(v)$ decrease when when $c_v < 0$ or $w_v > 0$ such that discovery increases incentives to search.

(iii) Payments $\tau^m(v)$ increase when $g_v > 0$ such that discovery decreases incentives to search.

(iv) As compared to direct search, effort is less ($e^m(v) < e^p(v)$) and search is less extensive ($\hat{v}^m < \hat{v}^p$) under monitored search.

A clear implication of these findings is the necessity of monitoring for implementing search with terminal payments. It is only because the agent's current technology can be monitored that it is possible to support search with terminal payments that change according to the initial technology. Without monitoring, each agent could claim to have begun from an initial technology that would maximize her payments for discovery. If for instance $\tau^m(v)$ is decreasing with v , an agent is better off not reporting interim discoveries because she will receive a smaller reward if she reaches the terminal state from a more advanced technology. This is illustrated in Figure 2 where a hypothetical payment schedule is depicted. In regions where terminal payments are increasing, the agent has an incentive to disclose small discoveries, whereas this isn't the case in regions where discoveries are decreasing.

The prediction of Proposition 2 regarding the effects of agency on search is also interesting for its implication about relative performance in different settings. Recent empirical findings by Levitt and Syverson (forthcoming), Bernheim and Meer (2007) and Hendel, Nevo, and Ortalo-Magné (2007) are consistent with our predictions on the effects of agency on search. These studies compare the prices at which houses are sold when the houses are marketed by an agent who represents a client seller as compared to when the agent owns the house herself. The studies find that agent-owned houses are typically sold for a higher price and that the time required to sell the house is shorter (correcting for the sales price) than a client-owned house. These findings are predicted by our model which shows that search will proceed slower and will be terminated sooner when search is conducted by an agent who represents a client rather than representing herself.

3.4 Renegotiation Proof Equilibrium

The monitored search contract that we've derived above displays some distinctive features. These include provisions for stationary terms that don't depend on time, performance payments that are awarded only after search is completed, permanent tenure for agents

and performance payments that are calibrated to the current technology. Moreover the contract is inefficient, as it induces an incomplete search that proceeds too slowly. The fact that the contract is not efficient and that it is special in some respects calls into question how robust the provisions are to different environments and which of our assumptions about agency setting are important for rationalizing the equilibrium.

Perhaps the stationarity of the contract is the most striking feature of the agency equilibrium. Provisions for terminating a non performing agent after a given period of time or paying greater rewards for discoveries that are made earlier are ruled out in equilibrium. It seems likely that the Markov Perfect equilibrium impose consistency requirements on the contract that preclude non stationary payments or terminations. After all a contract that was optimal for a given payoff/relevant state at one time must still be optimal later on provided the state is the same and there is no new information.

An important insight due to Bergemann and Hege(06) is to suggest that Markov equilibrium strategies may be equivalent to renegotiation proof-subgame perfect equilibrium in some settings. The intuition for this is that Markovian equilibrium impose consistency in behavior, in that the best reply strategies that are feasible for one subgame must continue to be feasible at another (payoff/equivalent) subgame. This is close, and in some cases identical, to the requirements for best replies strategies to be weakly renegotiation proof.

To illustrate this equivalence in our setting, consider the following simple example of a monitored search that has reached the terminal technology. At this stage the interaction between principal and agent becomes a repeated game, which terminates after the next discovery $v > \hat{v}^m$. If $\{\tilde{\kappa}_t, \tilde{e}_t, \tilde{q}_t\}$ is a weakly renegotiation proof SPE it follows that for all payoff/relevant equivalent histories h_t and $h_{t'}$ that

$$\{W^m(h_t), \Pi^m(h_t)\} = \{W^m(h_{t'}), \Pi^m(h_{t'})\}. \quad (11)$$

Otherwise if (11) didn't hold renegotiation proofness would be violated. Moreover it's easy to demonstrate that there is a unique weakly renegotiation proof SPE that is described by the strategies: the principal offers $\{\tau^m(\hat{v})\}$ as long as there has been no discovery and the agent has previously accepted all contracts. The principal reverts to the MPE offering $\{\tau^m(\hat{v})\}$ if the agent deviates and does not accept the contract. The agent accepts any contract that she expects to break even on and she selects an effort which maximizes her expected payoff. Notice finally that (11) is trivially satisfied by MPE continuation

strategies as well. Moreover there is also a unique MPE for this game. Hence it follows the two equilibrium are the same.

More generally we have:

Proposition 3: The unique MPE and weakly renegotiation proof SPE are equivalent.

In view of Proposition 4 its important to explain how agency search can be inefficient when it is generated through a bargaining process which is renegotiation proof. Inefficiencies persist because the agency contract can only be renegotiated after each period. This means that an inefficient allocation of effort or an inefficient termination of an agent can not be renegotiated before it occurs. Only renegotiation of contract provisions, and not the parties performance under the contract are subject to renegotiation. This implies that the parties can and do commit to behave efficiently during the period covered by the current contract.

4 Delegated Search

It is more difficult to motivate the agent to search when his progress cant be observed. Under monitored search the agents incentives to discover may be adjusted by the principal according to the agents previous success. Under delegated search it is the agent who determines incentives for future discovery by what he decides to disclose. Whereas the agents progress is automatically observed under monitoring, under delegation the agent must be induced to disclose what he has discovered. Hence the agent may conceal whether she succeeded in making any discovery at all or whether she has discovered a break through technology that would cause the project to be terminated, for instance. The agent may prefer to conceal or make incremental disclosures of her progress to extend the project and receive greater rewards for her work. This motive to conceal progress is particularly strong when the agent derives private benefit from working on the project, which declines as the project nears completion.

4.1 Equilibrium under Delegated Search

The sequence of events that transpires under delegated search is the same as before with the exception that (a) after each search the agent privately observes the discovery realization and (b) she makes a disclosure of her discovery to the principal. Since the agents

progress is not publicly observed the histories of observable events for the principal and agent must diverge. The agent continues to observe all of the previous strategic decisions, as well all of the actual discoveries she has made, so that her private history in continuous time is given by

$$h(t) = \{\tau(t'); e(t'); v(t'); \tilde{v}(t')\}_{t' < t}.$$

The principal observes only his own decisions and the discoveries disclosed by the agent, so his private history is

$$\tilde{h}(t) = \{\tau(t'); e(t'); \tilde{v}(t')\}_{t' < t}.$$

The principal forms a posterior belief of the distribution $\tilde{G}(v) := G(v | \tilde{h}(t))$ of the highest discovered technology based on his private history and beliefs about the agents strategies, This distribution is Bayesian updated after each disclosure based on the principals beliefs about the agents effort and disclosure strategies.

The principals strategies are:

$$\tau : \tilde{H}(t) \times \tilde{v}(t) \rightarrow R_+$$

$$q : \tilde{H}(t) \times \tilde{v}(t) \rightarrow \Delta\{0, 1\}.$$

The principal offers payments for disclosed (rather than) actual discoveries, and the decision to quit searching depends on the principals personal history. The agents strategies are

$$e : H(t) \times \tau \rightarrow R_+$$

$$\sigma : H(t) \times v(t) \rightarrow \{0, [\underline{v}, \max[v_h, v(t) v(t)]]\}.$$

The agents effort choice is conditioned on his private history and the current period contract payment. The agents disclosure choice depends on his private history and his current period discovery. The agent may disclose no discovery, represented by, or she may disclose any technology that does not exceed the highest technology actually discovered by the end of period. This restriction on feasible disclosures reflects the idea that an agent may hide progress but they may not claim more than they have achieved.

The relevant equilibrium for this setting of delegated search is a Bayesian Perfect Equilibrium which is specified below by:

Definition A sequence of strategies $\{q^d, \tau^d, e^d, \sigma^d\}_{t=0}^{t=\infty}$ comprises a Bayesian Perfect

Equilibrium provided for all histories $h(t)$ and $\tilde{h}(t)$ that

$$\begin{aligned} W^d(q^d | \tilde{h}(t), G(v | \tilde{h}(t))) &\geq W^d(q | \tilde{h}(t), G(v | \tilde{h}(t))) \quad \text{for all } q \\ W^d(\tau^d | \tilde{h}(t), G(v | \tilde{h}(t))) &\geq W^d(\tau | \tilde{h}(t), G(v | \tilde{h}(t))) \quad \text{for all } \tau \\ \Pi^d(e^d, \tau | h(t)) &\geq \Pi^d(e, \tau | h(t)) \quad \text{for all } e, \tau \\ \Pi^d(\sigma^d, \tau | h(t)) &\geq \Pi^d(\sigma, \tau | h(t)) \quad \text{for all } \sigma, \tau \end{aligned}$$

and $G(v)$ is Bayesian updated by the principal's correct beliefs about the agent's strategies wherever possible

The four optimality conditions insure the sequential optimality of the equilibrium strategy decisions. The restriction on the distribution function ensures the principal is updating his beliefs based on his conjectures about the agent's strategies, wherever possible.

4.2 Delegated Search: Discovery Increasing Incentives

We begin our characterization of delegated search by inquiring whether monitored search can be implemented under delegation. In principle this would require the agent to faithfully disclose each new discovery as it occurs. But under delegation the agent will only disclose a discovery if her expected surplus increases as a result. Consider the incentives for disclosure under a program in which the principal offers the agent a reward $\tau^d(v')$ whenever she discloses a new discovery v' . A necessary condition for implementing monitored search is that the agent receives the same expected payoff from discovery under delegation than she receives under monitored search. This requires

$$\tilde{F}(\hat{v}) \tau^m(\tilde{v}) = \int_{\tilde{v}}^{\tilde{v}'} \tau^d(\tilde{v}') dF. \quad (12)$$

Differentiating (12) with respect to \tilde{v} one can solve for $\tau^d(\tilde{v})$ as $\tau^d(\tilde{v}) = \frac{-\tilde{F}(\hat{v})\tau_v^m(\tilde{v})}{f(\tilde{v})}$ and the agent receives payment

$$\tau(\tilde{v}', \tilde{v}) = \int_{\tilde{v}}^{\tilde{v}'} \frac{-\tilde{F}(\hat{v})\tau_v^m(\tilde{v})}{f(\tilde{v})} d\tilde{v} \quad (13)$$

for disclosing a discovery v' from \tilde{v} . Denote by $\Pi^d(v, v)$ the continuation surplus for the agent who has previously discovered and disclosed v . Given rewards $\{\tau^d(\tilde{v})\}$ the agent's

expected surplus under this program is represented by

$$r\Pi^d(v, v) = \max_{e, \sigma} -c(v)e(v) + g(v) \quad (14)$$

$$+ \phi(e) \int_v^{\hat{v}} (\Pi^d(v', \sigma(v', v)) - \tau(\sigma(v', v), v) - \Pi^d(v, v)) dF,$$

where the agent selects a disclosure $e(v)$ which is less than or equal to the actual discovery v' . The agent prefers full disclosure provided $\tau^d(\tilde{v})$ is increasing in \tilde{v} . This requires that $\tau_v^m(\tilde{v}) \leq 0$ according to (13). This suggests that the monitored search program where rewards are decreasing in technology can be implemented under delegated search by offering payments $\tau^d(\tilde{v})$.¹¹ Moreover, monitored search can not be implemented under delegation when the monitored search program requires payments which are strictly increasing for some

Summarizing our findings we have,

Proposition 4: Monitored search programs with decreasing payments $\tau^m(v)$ can be implemented in PBE such that $(W^d(v), \Pi^d(v, v)) = (W^m(v), \Pi^m(v))$ with the following strategies:

$$(i) \tau^d(\tilde{v} | h(t)) = -\frac{\tau_v^m(\tilde{v})f(\tilde{v})}{F(\hat{v}^m)}; q(\tilde{v}) = \begin{cases} 0 & \text{if } \tilde{v} \leq \hat{v}^m \\ 1 & \text{if } \tilde{v} > \hat{v}^m; \end{cases}$$

$$(ii) e^d(\tau | h(t)) = e^m(v_{h(t)}, \tau^m); \sigma^d(v, \tilde{v}) = v;$$

$$(iii) G(v | \tilde{h}_t) = \begin{cases} 0 & \text{if } v \leq \tilde{v}_h \\ 1 & \text{if } \tilde{v} > \tilde{v}_h. \end{cases}$$

The Bayesian Perfect Equilibrium described in Proposition 4 is constructed by replacing the actual discovered technology by the disclosed technology in the principals optimal contract. Under delegation the agent is rewarded with a positive payment equal to the decrease in rewards she expects to receive in the future due to her discovery. The agent is not harmed by her disclosure and thus is willing to reveal any improvements she makes. Delegation differs from monitoring in that the agent receives interim payments for any progress she reports whereas she is paid only upon completion under monitored search.

¹¹The principal must prefer to implement the monitored search under these conditions. This is insured if $w_v > 0$ so that the principal prefers to know what progress has been achieved in order to implement the improvements. When $w_v = 0$, it's possible the principal prefers non disclosure, if it enables him to implement search at a lower cost.

Nonetheless the agent receives the expected payments under both and thus she supplies the same effort in each state under either monitoring or delegation. Hence monitoring is not necessary, the search may be entirely delegated to the agent.

It is important to understand why search can be delegated only in those settings where discovery increases incentives for future search. The reason is that when a discovery occurs the principal wishes to reduce future rewards because the agent requires less payment to supply effort. The agent is content to reveal his discovery so long as he is compensated for the reduced rewards he will receive in the future. The principal is happy to provide this compensation, because it allows him to offer lower payments in the future. Hence the incentives of the two parties to induce disclosure are well aligned in this case, so that it is not necessary for the principal to monitor the agent since the agent will voluntarily disclose any discovery to the principal.

4.3 Delegated Search: Discovery Decreasing Incentives

In contrast when discovery decreases incentives to continue search the interests of the agent and principal can not be aligned. The principal wants to increase payments after a small discovery, realizing the incentives for the agent to search are reduced. But if the agent makes a large discovery that terminates search, she will only disclose part of the discovery in order to receive a greater payment for the terminal discovery in the future. Hence there is no way to align the incentives of the agent and principal to disclose in this setting.

4.3.1 Model

This calls into question what search programs may be implemented under delegation when discovery decreases incentives for further search? To analyze this we turn to a special case of our model with the following features. Assume for simplicity there exists three relevant technology intervals $I_0 = (\underline{v}, v_1]$, $I_1 = (v_1, \hat{v}]$, $I_2 = (\hat{v}, \bar{v}]$, where $\underline{v} < v_1 < \hat{v} < \bar{v}$. The search begins with $\theta \in I_0$. The principal receives benefits $\frac{v(\hat{v})}{r}$ once a technology $\theta \in I_2$ has been discovered. The probability of drawing $\theta \in I_0, I_1, I_2$ is respectively p_0, p_1, p_2 where $p_i > 0$ for all i and $p_0 + p_1 + p_2 = 1$. We assume the unit cost of effort is constant but that the agents' expected benefits increase as the technology increases from I_0 to I_1 so that $g_0 < g_1$. This insures the incentives to search further are reduced after a discovery of I_1 . For the agent in intervals I_0 and I_1 by g_0 and g_1 respectively with \cdot .

The principal selects a payment strategy $(\tilde{h}(t))$ to reward a disclosure of $\in I_2$ conditional on his private history $\tilde{h}(t) = \{\tau(t'); \tilde{I}(t')\}_{t' \leq t}$ where $\tilde{I}(t')$ is the sequence of previous disclosures. Offering just a reward for the final discovery is sufficient to implement the principals preferred search for this setting. The agent with current technology noted by A_i , selects an effort allocation $(h(t), \tau)$ and a disclosure policy $\sigma_i(h(t), \tau, v')$. The strategies are conditional on the current payment, on the current discovery and on the agents private history $\tilde{h}(t) = \{\tau(t'); e(t'); I(t'); \tilde{I}(t')\}_{t' \leq t}$ consisting of previous strategic decisions as well as the sequence of the agents actual disclosures. The disclosure policy σ_i is a mixed strategy mapping from the set of histories and new discoveries into the set of feasible disclosures $\{j | j \leq i\}$. σ_i^j is the probability that A_i discloses technology $j \leq i$

The principals posterior of the different agent types is denoted by λ_i for $i = 0, 1, 2$ and $\lambda = (\lambda_0, \lambda_1, \lambda_2)$ is the corresponding probability distribution. Since A_i may only disclose technologies less or equal to i , a lower bound for the distribution is $\tilde{I}_{\tilde{h}(t)}$ the highest technology disclosed in history $\tilde{h}(t)$. The principal Bayesian updates his priors λ after each period, based on \tilde{h} and his knowledge of the agents disclosure and effort allocation strategies.

Let $W^d(\lambda | \tilde{h}(t))$ and $\Pi_i^d(h(t))$ represent the surplus functions for the principal and agent type A_i agent respectively. The principal selects $\{\tau^d(\tilde{h}(t))\}$ to maximize $W^d(\lambda | \tilde{h}(t))$ and A_i selects $\{e_i, \sigma_i\}$ to maximize $\Pi_i^d(h(t))$

4.3.2 Delegated Search Equilibrium

To gain some insight for how the parties behave in this setting, lets first conjecture what the equilibrium disclosure policy for the agent might look like. It seems intuitively clear that in equilibrium, disclosure of the terminal discovery must be gradual. Otherwise if disclosure were immediate, then the principal would infer that any agent who had not yet disclosed should receive a greater payment for final discovery to induce her to search more. Hence an agent with a terminal discovery would benefit by delaying disclosure to receive greater compensation in the future.

In order for there to be gradual disclosure, it must be indifferent as to when she discloses. This requires that the payment is increasing at the rate of during periods of no disclosure. Otherwise she would immediately disclose or wait until the discounted

payment had reached a maximum. Suppose for instance that $\tilde{h}(t) = I_1$ so that the principal expects the agent is either type A_1 or A_2 with probabilities λ_1 and λ_2 respectively. In order for the payment to grow at the rate r during a period of no disclosure, must be falling at a certain rate $\frac{d\lambda_2}{dt}$, to satisfy the condition $(\frac{d\tau}{dt} = r\tau)$ that real payments are constant, where $\tau(\lambda_1, \lambda_2)$ is the payment that maximizes $W^d(\lambda_1, \lambda_2)$ and

$$\frac{d\tau}{dt} = \left(-\frac{d\tau}{d\lambda_1} + \frac{d\tau}{d\lambda_2} \right) \frac{d\lambda_2}{dt}. \quad (15)$$

For each instant of no disclosures, the prior adjusts at the rate at

$$\frac{d\lambda_2}{dt} = \lambda_1 \phi(e(\tau_1)) p_2 - \lambda_1 \lambda_2 \sigma_2^2(\lambda_2), \quad (16)$$

where $\lambda_1 \phi(e(\tau_1)) p_2$ is the rate at which A_1 types become A_2 types through new discovery and $-\lambda_1 \lambda_2 \sigma_2^2(\lambda_2)$ is the rate at which the likelihood of A_1 is Bayesian downgraded because of no disclosure. In order for to adjust at the required rate $\frac{d\lambda_2}{dt}$ must disclose at the rate $\sigma_2^2(\lambda_2)$ that satisfies (16) such that

$$\sigma_2^2(\lambda_2) = \frac{\frac{d\lambda_2}{dt} - \lambda_1 \phi(e(\tau_1)) p_2}{\lambda_1 \lambda_2} \quad (17)$$

But since A_2 is indifferent to when he discloses, the rate of disclosure required to maintain the reward at the required level to support mixing is an optimal response for the agent.

It turns out these conjectures are supported by the following:

Proposition 5: A PBE $\{\tau^d, e^d, \sigma^d\}$ for delegated search with discovery decreasing incentives exists with strategies and beliefs λ that satisfy

- (i) $\tau(\tilde{h}(t))$ maximizes $W^d(\tilde{h}(t))$ for all $\tilde{h}(t)$
- $e^d(h(t), \tau), \sigma^d(h(t), \tau)$ maximize $\Pi^d(h(t))$ for all $h(t)$ and τ

Along the equilibrium path:

- (ii) For all $\tilde{h}(t)$ such that $\tilde{I}_{h(t)} = I_o, \tau^d(\tilde{h}(t)) = \tau^d(1, 0, 0)$

The principal presumes there have been no discoveries until the first disclosure;

- (iii) For all $\tilde{h}(t)$ such that $\tilde{I}_h = I_1, \tau^d(\tilde{h}(t)) = \tau^d(0, \lambda_1(\tilde{h}(t)), \lambda_2(\tilde{h}(t)))$

The principal selects rewards under the presumption that there has been partial disclosure of discovery;

(iv) For all $h(t)$ such that $I_h = 1$ or $\sigma_1^d(h(t)) = 1$

Type 1 agents fully disclose their discovery;

(v) For all $h(t)$ such that $I_h = 2$, then $\sigma_2^2(h(t), \tau^d), \sigma_2^1(h(t), \tau^d) > 0$

Type 2 mixes between partial and complete disclosure;

(vi) For all $\tilde{h}(t)$ such that $I_h = 1$ and $\lambda_2(\tilde{h}(t)) > 0$, $\frac{d\tau(0, \lambda_1, \lambda_2)}{dt} = r\tau(0, \lambda_1, \lambda_2)$

The real reward remains constant until there is terminal disclosure or the likelihood of a type 2 agent is zero.

Along the equilibrium path the agent begins in I_1 . The agent is offered a constant payment $\tau(1, 0, 0)$ which maximizes $V^d(1, 0, 0)$ until she discloses a discovery $v \in I_1$. The agent allocates effort t_0 which maximizes her profit π_0^d until she makes a discovery v . If $v \in I_1$ the agent discloses v and continues her search until she makes another discovery. If the initial discovery is $v \in I_2$ the agent mixes between disclosing and I_1 . If the agent discloses v she receives payment $\tau(1, 0, 0)$ and the search is terminated. Otherwise the agent discloses v and stops searching and waits until a later time to disclose I_2 . In order for A_2 to postpone disclosure, the real payment must be constant so that must be rising at the rate r . In equilibrium the principal believes λ_2 is falling as long as no disclosure has yet occurred. A_2 is indifferent to disclosure so she discloses at a rate to support the principal's updated beliefs about types. Eventually if there is no disclosure of I_2 after some period, the principal correctly concludes that the agent must be starting in I_1 . The payment for terminal success then converges to $\tau(1, 0, 0)$ the optimal payment for an agent starting in I_1 .

The delegated equilibrium portrayed in Proposition 5 has the property we anticipated that disclosure is gradual. Agents delay disclosing terminal discovery in order to secure greater rewards for their search. The principal can not commit to keeping the payment for terminal discovery low, if he knows the agent has made intermediate progress. Therefore an agent will not immediately disclose the discovery because he could receive greater rewards by

waiting. The principal correctly infers that agents who have disclosed an intermediate discovery, may have already finished their search. Therefore the principal only raises rewards gradually, only after he becomes convinced that the agent is still searching and therefore needs an inducement to continue.

The equilibrium exhibits the often observed time delays and cost escalation that seem to plague agency contracts. It appears that when discovery decreases incentives to finish the search that such deterioration in performance is inevitable for the reasons we've discussed here. In effect, the agency relationship deteriorates over time into a type of "two-sided holdup.". The agent delays completing the search to receive greater payment. while the principal delays increasing payment until he is convinced it is necessary to induce the agent to finishing searching.

Since the contract terms change over time, until terminal disclosure, it appears that the contract is non stationary and that it varies with time. This arise even though the costs and benefits from search do vary over time. In reality, the contract is exhibiting a type of pseudo non stationarity. The perceived state of the relationship is changing systematically overtime, so that contract terms that track those variations in states appear to change with time rather than with changing perceptions.

Finally it is apparent that monitoring is most valuable in these setting where inducing completion of a project grows more difficult with each discovery. In those case the incentives of the principal and agent for disclosure can not be aligned. The result is a dysfunctional search process where the agent holds out to receive greater payments and the principal resists increasing rewards until he learns that the search has not been completed.

5 Conclusion

This paper derives the optimal dynamic contract and organization structure for a principal who hires an agent to search for innovations. The principal does not know how diligently the agent works or what discoveries she obtains unless he monitors her progress. The principal directs the agent in one of two ways in order to resolve these information problems. The first is through monitored search whereby the principal monitors the agent's progress during the search process. Under monitoring the agent is rewarded with a final payment upon completion of the project. The payment varies according to size of the terminal discovery. By observing the agent's progress the principal can adjust payments to the

agent to either increase or decrease her rate of discovery as desired. To minimize the cost of agency search the principal induces the agent to search at a lower intensity level and to terminate search at a lower discovery level than is optimal.

The second process for managing the agent is through delegated search. In this instance, the principal is unable to observe discoveries directly. He must therefore rely on agent to disclose her progress. Payments for discoveries must be structured in such a way as to induce the agent to disclose her progress. Since delaying disclosure or only reporting a portion of discovery is always possible, the principal must reward larger discoveries with greater payments if she is to induce the agent to fully reveal her progress. This constrains the types of search programs the principal can implement under delegation. For instance programs in which incentives to search increase as the project progresses can not be implemented because this requires that smaller terminal discoveries receive greater rewards than larger discoveries which is not incentive compatible. In such instances this means that payments will be not respond to interim discoveries, and that agents of different types (at different stages of the project) will receive the same rewards so that bunching will arise.

Our analysis is but a first attempt at analyzing dynamic agency relationships. Moreover the model of search we consider is special in several respects. Consequently there are several extensions of our analysis that one might pursue. For instance, other forms of search and innovation that require riskier and more creative approaches to research might be analyzed using our framework. Mansos (2007) study of how to induce innovation could be adopted to a dynamic agency setting like ours to discover what types of dynamic contracts and organizations are best suited to manage creativity. Another possibility is to study the optimal degree of concentration in research, and whether innovation is best done by one or sever independent research groups. This would address the question of how organization structure can address the problems of information asymmetries and alignment of dynamic incentives that are central to dynamic agency relationships. A final promising direction for further research is to consider how property rights over new discoveries can be allocated between the principal and agent through the stages of discovery to align incentives and facilitate the disclosure of information.




Figure 1: Compensation Payments

6 Appendix

Proof of Proposition 1 and 2: Propositions 1 and 2 are proved together through a sequence of steps. To begin we first consider the problem in discrete time with periods of duration $\Delta > 0$. and then consider the limiting case in which $\Delta \rightarrow 0$ to derive the continuous time formulation:

Agent's discrete time problem:

In what follows we assume the principal adopts a stopping rule in which he terminates search once a technology equal to or greater than (v, \bar{v}) has been discovered. Below we verify this is optimal. Under this assumption, the expected surplus for an agent who is offered contract $\{\tau_t, \mu_t\}$ given state $e_{ht} < \hat{v}^m$ is

$$\begin{aligned} \Pi_t^m(v) = & \max_{e_t} (g(v_t) - c(v_t) e_t(v_t)) \Delta \\ & + \delta \Delta \phi(e_t) \int_v^{\hat{v}^m} ((1 - \mu(v', v)) \Pi^m(\max(v', v_t))) dF \\ & + \delta \Delta \phi(e_t) \tilde{F}(\hat{v}^m) \tau(v) \\ & + \delta (1 - \Delta \phi(e_t)) \Pi^m(v) (1 - \mu(0, v)). \end{aligned} \quad (A1)$$

The agent selects $e(v)$ to maximize (A1). Hence the condition for the optimal effort allocation as $\Delta \rightarrow 0$ becomes

$$c(v) = \phi'(e) \left(\int_v^{\hat{v}^m} \Pi^m(v') - \Pi^m(v) dF(v') + \tilde{F}(\hat{v}^m) (\tau(v) - \Pi^m(v)) \right) = 0. \quad (A2)$$

For future reference we record the effect of greater payments on the optimal effort allocations as

$$\frac{de^m}{d\tau} = - \frac{\phi'(e^m)^2 \tilde{F}(\hat{v}^m)}{\phi''(e^m) c(v)} = \frac{(1 - \gamma) \tilde{F}(\hat{v}^m) \phi(e(v))}{\gamma c(v) e^m(v)}.$$

Note also that we may solve for the continuous time surplus function for the agent by multiplying (A1) by $(1 + r\Delta)$ and dividing by Δ and letting Δ go to zero to obtain

$$r\Pi^m(v) = \max_{e(v)} -C(e(v), v) + \phi(e) \left(\int_v^{\hat{v}^m} (\Pi^m(v') - \Pi^m(v)) dF(v') + \tilde{F}(\hat{v}^m) (\tau(v) - \Pi^m(v)) \right) \quad (A3)$$

It follows from (A3) and (A2) that

$$\Pi^m(v) = \frac{c(v) e^m(v) (1 - \gamma) + \gamma g(v)}{\gamma r} \text{ for } v \leq \hat{v}^m.$$

This completes the proof of Proposition 1.

Now consider the principals problem in discrete time,

$$W_t^m(v) = \delta W_{t+\Delta}^m(v) + \max_{\tau, \hat{v}^m} \delta \Delta \phi(e_t) \left(\int_{\underline{v}}^{\hat{v}^m} (W_{t+\Delta}^m(v') - W_{\tau+\Delta}^m(v)) dF + \int_{\hat{v}^m}^{\bar{v}} \left(\Delta \frac{\omega(v)}{\Delta r} - W_{t+\Delta}^m(v) - \tau^m(v) \right) dF \right) \quad (\text{A4})$$

Now by multiplying by $(1 + r\Delta)$ and dividing by Δ and letting $\Delta \rightarrow 0$ in (A4) we obtain the continuous time dynamic program is

$$rW^m(v) = \max_{\tau^m, \hat{v}^m} \phi(e_t) \left(\int_{\underline{v}}^{\hat{v}^m} (W^m(v') - W^m(v)) dF + \int_{\hat{v}^m}^{\bar{v}} \left(\frac{\omega(v)}{r} - W^m(v) - \tau^m(v) \right) dF \right). \quad (\text{A5})$$

The conditions for the optimal payments and the optimal stopping technology are given respectively by

$$0 = \frac{dW^m}{de^m} \frac{de^m}{d\tau^m} + \frac{dW^m}{d\tau^m} \quad (\text{A6})$$

$$= \phi'(e) \left(\int_{\underline{v}}^{\hat{v}^m} (W^m(v') - W^m(v)) dF + \int_{\hat{v}^m}^{\bar{v}} \left(\frac{\omega(v)}{r} - W^m(v) - \tau^m(v) \right) dF \right) \left(-\frac{\tilde{F}(\hat{v}^m) \phi'^2(e^m(v))}{c(v) e^m(v) \phi''} \right) - \phi(e) \tilde{F}(\hat{v}^m),$$

$$\frac{\omega(\hat{v})}{r} - W^m(\hat{v}^m) = 0. \quad (\text{A7})$$

By differentiating (A6) and (A7) totally with respect to v we obtain the following two differential equations involving g_v, c_v, g_v and τ_v :

$$c_v e(v) + c(v) e_v = 0 \quad (\text{A8})$$

$$\begin{aligned} \tau_v \left(\frac{\tilde{F}(\hat{v}^m) r \gamma \phi(e(v))}{r + \phi(e(v) \tilde{F}(v))} \right) &= -c_v(e(v)) \left(\frac{r + \phi(e) \tilde{F}(v) - \gamma \phi(e(v) \tilde{F}(v))}{r + \phi(e(v) \tilde{F}(v))} \right) \\ &+ e_v (-c(v) (1 - \gamma)) \\ &+ g_v \left(\frac{\gamma \phi(e(v))}{r + \phi(e(v) \tilde{F}(v))} \right). \end{aligned} \quad (\text{A9})$$

Combining these two equations were able to solve e_v as a function of v and g_v to obtain

$$\tau_v = c_v \left(\frac{e(v)}{\tilde{F}(\hat{v}^m) \phi(e(v))} \right) + \left(\frac{g_v}{\tilde{F}(\hat{v}^m) r} \right). \quad (\text{A10})$$

From (A10) we can verify that parts (i) and (iii) of Proposition 3 hold.

Now to obtain part (iv) note that from (A6) and (A8) we have

$$\begin{aligned} W^m(v) &= \frac{c(v) e(v) (1 - \gamma)}{r\gamma^2}, \\ W_v &= 0. \end{aligned}$$

This implies that there exists a unique \hat{v}^m that satisfies the stopping condition (A7). Further the stopping condition under direct search is given by

$$\frac{\omega(\hat{v}^p)}{r} - W^p(\hat{v}^p) = 0. \quad (\text{A11})$$

It follows that $W^p(v) > W^m(v)$ for all $v < \hat{v}^m, \hat{v}^p$ since the direct search can duplicate any monitored search process, without incurring any rents. Consequently it follows from (A11) and (A7) that $\hat{v}^p > \hat{v}^m$. Moreover, its easy to show that

$$\begin{aligned} W^p(v) &= \frac{(1 - \gamma) c(v) e^m(v) + \gamma g(v)}{\gamma r} \\ &> W^m(v) \\ &= \frac{(1 - \gamma) c(v) e^p(v)}{\gamma^2 r}, \end{aligned} \quad (\text{A12})$$

thus implying that $e^p(v) > e^m(v)$. This completes the proof of part (i) of Proposition 2.

Finally to prove that the MPE is unique, we first note that starting at the unique stopping point \hat{v}^m the equilibrium is unique from that state forward. At the unique values of the strategies $\{e(\hat{v}^m), e(\hat{v}^m)\}$ For $v < \hat{v}^m$ the effort strategy $e^m(v)$ is uniquely defined by

$$e^m(v) = \frac{e(\hat{v}^m) c(\hat{v}^m)}{c(v)}$$

and the payment strategy $\tau^m(v)$ is uniquely defined by

$$\begin{aligned} \tau(v) &= \Pi(v) + \frac{c(v) e(v)}{\gamma \phi(e(v)) \tilde{F}(\hat{v}^m)} \\ &= \frac{c(\hat{v}^m) e(\hat{v}^m) (1 - \gamma) + \gamma g(v)}{\gamma r} + \frac{c(\hat{v}^m) e(\hat{v}^m)}{\gamma \phi\left(\frac{e(\hat{v}^m) c(\hat{v}^m)}{c(v)}\right) \tilde{F}(\hat{v}^m)}. \end{aligned}$$

Hence the MPE is unique.

Proof of Proposition 3: (Sketch) The proof proceeds by demonstrating that any SBE with weakly renegotiation proof strategies must be stationary. It then follows that any stationary SBE must be a MPE. Therefore since there is a unique MPE, it must coincide with the weakly renegotiation proof SPE.

Suppose there exists a weakly renegotiation proof SPE with non stationary strategies. That implies that there is a history with a subsequence $(v_{t'}, h_{t'})$ at which $v_{h_{t'}} = v_{h_{t''}}$ such that strategies selected at t' and t'' are different for one or both players. It follows that the players value functions remain constant along this sequence by the definition of weakly renegotiation proof SBE. This implies there must be multiple SPE all of which yield the same values for both players. Since the players could pick the same continuation strategies along the path, there must therefore be multiple stationary SBE all of which yield the same payoffs to the players.

It follows that any stationary weakly renegotiation proof SPE must be an MPE. Since there is a unique MPE, there must therefore be unique stationary weakly negotiated SPE. Hence the two equilibrium must be equivalent.

Proof of Proposition 4: Given $\tau^d(\tilde{v} | \tilde{h}(t))$ the agents best effort and disclosure responses $\{e^d, \sigma^d\}$ solve the following dynamic program

$$r\Pi^d(v, v) = \max_{e, \sigma} -c(v)e(v) + g(v) + \phi(e) \int_v^{\tilde{v}} (\Pi^d(v', \sigma(v', v)) - \tau(\sigma(v', v), v) - \Pi^d(v, v)) dF.$$

The optimal effort response is $e^d(v, \tau) = e^m(v, \tau^m)$ and the optimal disclosure is $\sigma^d(v, \tilde{v}) = v$ since $\tau(\sigma(v', v), v)$ is increasing in σ given that $\tau_v^m \leq 0$ by construction of the disclosure payments. Given the agents effort and disclosure strategies $\{e^d, \sigma^d\}$ the principal optimal payment response is $q^d(\tilde{v}', \tilde{v}) = \int_{\tilde{v}}^{\tilde{v}'} \frac{-\tilde{F}(\tilde{v})\tau_v^m(\tilde{v})}{f(\tilde{v})} d\tilde{v}$. This strategy and the quitting strategy q^d implement the monitored search outcome, which is the optimal one among those where the agent can perfectly monitor the agents progress. Finally the Bayesian updated beliefs $G(v | \tilde{h}_t)$ given τ^d, e^d, q^d are given by part (iii) of the Proposition.

Proof of Proposition 5: The proof proceeds in three steps. First we verify that given the principals payment strategy $q^d(\tilde{h}(t))$ and beliefs about the agents responses,

the agents actual responses $\{\sigma^d(h(t), \tau), \tau^d(h(t), \tau)\}$ are optimal. The second step is to show that given $\{\sigma^d(h(t), \tau), \tau^d(h(t), \tau)\}$ the principals beliefs are correct and the payment strategy is an optimal response. Finally we verify that the principals beliefs about the agents type are Bayesian updated appropriately.

The principals payment strategy satisfies the following:

$$\begin{aligned} \tau^d(\tilde{h}(t)) &= \tau^d(1, 0, 0) \text{ for } h(t) \text{ such that } \tilde{h}_h = I_0 \\ \tau^d(\tilde{h}(t)) &= \tau^d(0, \lambda_1, \lambda_2) \text{ for } h(t) \text{ such that } \tilde{h}_h = I_1 \end{aligned}$$

where $\tau^d(h(t))$ is the payment that maximizes $\tau^d(h(t))$ and $\lambda(\tilde{h}(t))$ is the principals beliefs about the agents types given history $\tilde{h}(t)$. The principals corresponding beliefs about the agents type and her behavior are that

For all $\tilde{h}(t)$ such that $\tilde{h}_h = I_0 : \lambda(\tilde{h}(t)) = (1, 0, 0)$ and $(e_o, \sigma_o) = (e_o^d, \sigma_o^d)$

For all $\tilde{h}(t)$ such that $\tilde{h}_h = I_1 : \lambda(\tilde{h}(t)) = (0, \lambda_1(\tilde{h}(t)), \lambda_2(\tilde{h}(t)))$ and $(e_o, \sigma_o) = (e_o^d, \sigma_o^d)$

and the agents type and behavior is self evident for all histories such that $\tilde{h}_h = I_2$ as the search is terminated.

Its easy to verify that the agents actual effort and disclosure strategies $\{\sigma^d, \tau^d\}$ are an optimal response to the principals payments. In particular as long as the agent is type 0, his best response is to allocate effort that maximizes his expected profits. Moreover, once the agent makes a discovery his best response is to disclose if he discovers \tilde{I}_1 and to either disclose \tilde{I}_1 or \tilde{I}_2 if the discovers \tilde{I}_2 . The agent is indifferent as to what he discloses because in the continuation game, the agent expects to be receive a discounted payment of $\tau(1, 0, 0)$ in the future as the terminal discovery payment is expected to rise at the rate r .

Regarding the principal, his payment strategy $\tau^d(\tilde{h}(t))$ is a best response to the agents strategies, given the probabilities that she attaches to the agents type from observing the history $\tilde{h}(t)$. In particular the optimal payment for the principal is $(1, 0, 0)$ until the agent makes a disclosure. This is because the agent is a type 0 until she discloses and the payment $(1, 0, 0)$ is optimal given the agents type. Once a disclosure has been made the principals best response is $(0, \lambda_1(\tilde{h}(t)), \lambda_2(\tilde{h}(t)))$ until a disclosure of \tilde{I}_2 . The principal expects that a type 1 agent is indifferent between disclosing \tilde{I}_1 and \tilde{I}_2 . This expectation is correct in so far as the principal increases the terminal payment

$\tau(0, \lambda_1(\tilde{h}(t)), \lambda_2(\tilde{h}(t)))$ at the rate of $\sigma_2(\tilde{I}_2, \lambda_2)$. The rate of increase in the terminal payment is supported by the updating of the agent types that arises as long as a terminal discovery is not disclosed. This requires that

$$\begin{aligned} r\tau(0, 1 - \lambda_2, \lambda_2) &= \frac{d\tau}{d\lambda_2} \frac{d\lambda_2}{dt} \\ &= \frac{d\tau}{d\lambda_2} (1 - \lambda_2) \phi(e_1(\tau(\lambda_2))) p_2 - (1 - \lambda_2) \lambda_2 \sigma_2(I_1). \end{aligned}$$

Or solving for the disclosure rate $\sigma_2(\tilde{I}_2, \lambda_2)$, we obtain

$$\sigma_2(\tilde{I}_2, \lambda_2) = (1 - \lambda_2) \phi(e_1(\tau(\lambda_2))) p_2 - \frac{r\tau(\lambda_2)}{\frac{d\tau}{d\lambda_2}}.$$

The disclosure of the terminal discovery must proceed at the rate $\sigma_2(\tilde{I}_2, \lambda_2)$ to support the principals beliefs about the agents type. Since the type is indifferent to disclosing or not, it is optimal for him to disclose at the required rate $\sigma_2(\tilde{I}_2, \lambda_2)$. This completes the proof to Proposition 5.

References

- Albuquerque, Rui, and Hugo A. Hopenhayn. 2004. Optimal Lending Contracts and Firm Dynamics. *Review of Economic Studies*, 71(2): 285-315.
- Atkeson, Andrew. 1991. International Lending with Moral Hazard and Risk of Repudiation. *Econometrica*, 59(4): 1069-1089.
- Aghion, Philippe, and Jean Tirole. 1997. Formal and Real Authority in Organizations. *Journal of Political Economy*, 105(1): 1-29.
- Arnold, Michael A. 1989. The Principal-Agent Relationship in Real Estate Brokerage Services. *Journal of the American Real Estate and Urban Economics Association*, 20(1): 89-106.
- Baron, David P., and David Besanko. 1984. Regulation, Asymmetric Information, and Auditing. *RAND Journal of Economics*, 15(4): 447-470.
- Bergemann, Dirk, and Ulrich Hege. 2005. The Financing of Innovation: Learning and Stopping. *Rand Journal of Economics*, 36(4): 719-752.
- Bernheim, Douglas, and Jonathan Meer. 2007. How Much do Real Estate Brokers Add? A Case Study. Stanford Institute for Economic Policy Research, Discussion Paper No. 06-041.
- Clementi, Gian Luca, and Hugo A. Hopenhayn. 2006. A Theory of Financing Constraints and Firm Dynamics. *Quarterly Journal of Economics*, 121(1): 229-265.
- Cornelli, Francesca, and Oved Yosha. 2003. Stage Financing and the Role of Convertible Securities. *Review of Economic Studies*, 70(1): 1-34.
- Cremer, Jacques, Fahad Khalil, and Jean-Charles Rochet. 1998. Contracts and Productive Information Gathering Games and Economic Behavior, 25(2): 174-193.
- DeMarzo, Peter M., and Michael J. Fishman. 2007. Optimal Long-Term Financial Contracting. *Review of Financial Studies*, 20(5): 2079-2128.
- Dewatripont, Mathias, and Eric Maskin. 1995. Credit and Efficiency in Centralized and Decentralized Economies. *Review of Economic Studies*, 62(4): 541-555.
- Freixas, Xavier, Roger Guesnerie, and Jean Tirole. 1985. Planning under Incomplete Information and the Ratchet Effect. *Review of Economic Studies*, 52(2): 173-191.
- Green, Edward. 1987. Lending and the Smoothing of Uninsurable Income, in Edward Prescott and Neil Wallace, Eds., *Contractual Arrangements for Intertemporal Trade*, Minneapolis: University of Minnesota Press.
- Grossman, Sanford J. and Oliver D. Hart. 1983. An Analysis of the Principal-Agent Problem. *Econometrica*, 51(1): 7-45.

- Hart, Oliver D. and Jean Tirole. 1988. Contract Renegotiation and Coasian Dynamics. *Review of Economic Studies*, 55(4): 509-540.
- Hendel, Igal, Aviv Nevo, and François Ortalo-Magné. 2007. The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison.com. NBER Working Paper 13360.
- Hopenhayn, Hugo A., and Jean Pablo Nicolini. 1997. Optimal Unemployment Insurance. *Journal of Political Economy*, 105(2): 412-438.
- Kaplow, Louis, and Steven Shavell. 1994. Optimal Law Enforcement with Self-Reporting of Behavior. *Journal of Political Economy*, 102(3): 583-606.
- Laffont, Jean-Jacques, and Jean Tirole. 1993. *A Theory of Incentives in Procurement and Regulation*. Cambridge, MA: MIT Press.
- Levitt, Steven D., and Chad Syverson. Forthcoming. Market Distortions when Agents are Better Informed: The Value of Information in Real Estate Transactions. *Review of Economics and Statistics*.
- Lewis, Tracy R., and Michel Poitevin. 1997. Disclosure of Information in Regulatory Proceedings. *Journal of Law, Economics, and Organization*, 13(1): 50-73.
- Lewis, Tracy R., and David E.M. Sappington. 1997. Information Management in Incentive Problems. *Journal of Political Economy*, 105(4): 796-821.
- Manso, Gustavo. 2007. Motivating Innovation. MIT mimeo.
- McCall, John J. 1970. Economics of Information and Job Search. *Quarterly Journal of Economics*, 84(1): 113-126.
- Mookherjee, Dilip, and Ivan Png. 1989. Optimal Auditing, Insurance, and Redistribution. *Quarterly Journal of Economics*, 104(2): 399-415.
- Neher, Darwin V. 1999. Staged Financing: An Agency Perspective. *Review of Economic Studies*, 66(2): 255-274.
- Shavell, Steven, and Laurence Weiss. 1979. Optimal Payment of Unemployment-Insurance Benefits over Time. *Journal of Political Economy*, 87(6): 1347-1362.
- Shin, Hyun Song. 1994. The Burden of Proof in a Game of Persuasion. *Journal of Economic Theory*, 64(1): 253-264.
- Spear, Stephen E., and Sanjay Srivastava. 1987. On Repeated Moral Hazard with Discounting. *Review of Economic Studies*, 54(4): 599-617.
- Stern, Scott. 2004. Do Scientists Pay To Be Scientists? *Management Science*, 50(6): 835-853.