Delegating Search to Agents *

Tracy R. Lewis†

Revised February 2009

Abstract

Searching for a solution or for the best alternative is an important activity, one that is often delegated to an agent to perform. This paper examines the properties of delegated search to determine how it differs from the process of search that an individual undertakes on his own behalf. The analyses predict that as compared to self directed search, (a) delegated search is slower and takes longer to complete, (b) under delegation, the rate of search varies with time, with most of the discoveries occuring early in the search and (c) multiple agents may be employed to search where one agent would be sufficient if the search was carried out efficiently. These departures from the efficient decision theoretic search process arise as a contractual response to dynamic agency costs that naturally occur when search decisions are delegated to an agent.

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

*I am grateful to Jim Bergin, Alessandro Lizzeri, Jim Peck, Jano Zabojnik and to seminar participants at Ohio State University, and Queens University for their helpful comments on earlier drafts of this paper. I am particularly thankful to Marco Ottaviani for the extensive advice and encouragement he has provided to me in preparing this paper.

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

1 Introduction

The wealth and income of most people derives in part from their ability to *search*—that is by their ability to find a job, to discover low prices and to solve problems. Consequently economists have spent much time observing and modeling search. The work horse of modern search theory is the elegant decision theoretic model developed nearly 40 years ago by McCall(1970) and Mortenson(1970) to study job markets. This simple paradigm has spawned a surprisingly general and useful theory of search. The basic model has been extended to consider multiple stages and types of search, to incorporate learning and to allow for search in two sided markets. By combining decision models with theories of market equilibrium, economists have generated a rich set of predictions for a variety of different markets. For instance, theories of consumer search have been employed to explain equilibrium price dispersion and sales in product and service markets. Theories of job search have been used to predict wage distributions, employment duration and vacancy rates in labor markets.[1]

The decision theoretic model of search views the market as a means for generating random offers and prices. Individuals conduct their own search by sequentially sampling from the distribution of offers to find their preferred one. The theory predicts that search is governed by a simple stopping rule that balances the benefits of further search against the costs of additional sampling. The stopping rule results in a process whereby people draw samples at a uniform rate until an offer exceeding a predetermined value is discovered. While this model and existing search literature have generated many important insights and predictions, the extant theory overlooks one important detail. *Individuals often do not search themselves, but instead delegate search to independent agents to perform.* This paper explores the importance of delegation in search and develops a model of agency to explain how agents search. Specifically I derive a number of results regarding the relative efficiency of delegated search, and how agents are contracted to search.

Among my results, I find as compared with self directed search,

(a) Delegated search is slower to produce results,

(b) Under delegation the search rate varies with time, and the likelihood of discovery

[1] See Rogerson, Shimer and Wright (2006) and the references cited therein for an excellent overview of the theory and applications of modern search theory .

declines the longer the search lasts,

(c) *A sequence of short term agents may be employed to search under delegation even though it is less costly to use one long-lived agent.*

These results are shown to be a natural consequence of contractual arrangements to mitigate dynamic agency problems. For example, consider an owner who contracts with a risk neutral agent to search for a new process or technology. The contract is a long term agreement that the parties commit to that stipulates how the agent is evaluated and paid as well as requirements for his continued employment. The agent, whose effort and discoveries can't be observed, is also wealth constrained and therefore unable to post a large bond or to acquire the project to insure his performance. Hence this is a setting containing all of the essential ingredients, including moral hazard, hidden information, and liquidity constraints, that differentiate delegated search from self directed search. In the analysis to follow I show that in any setting with these ingredients, delegating search to agents will result in one or more of the distortions (a)-(c) described above.

The rationale for these findings stems from two important insights that are revealed by my analysis in sections 2 and 4. The first is that agents have a tendency to procrastinate when given a long time to search. Agents delay searching to spread their effort over time in order to reduce costs. Hence preventing this delay is a *dynamic cost of agency* that the owner incurs in addition to the usual agency costs caused by hidden action.[2] As a result delegated search is slower than if the agent were searching directly for himself. This explains prediction (a) above as well as why clients are sometimes dissatisfied with the service received from outsourcing search to an agent.

At the same time the analysis yields a second insight into how contracts for search are designed to reduce agent delay. One remedy for delay is to tilt the reward schedule towards the present away from the future. Offering a reward for early completion and a penalty for delay forces the agent to search harder in the present to avoid delay. The effect of this is to make the agent search faster at the start and slower towards the end as the search wears on. This explains prediction (b) above. In addition, a second remedy for delay is to employ multiple agents to search in sequence, until one finishes the job.

[2] This *dynamic agency cost* has largely gone unnoticed in the literature. Notable exceptions are Toxvaerd(2006), Mason and Valimaki(2008) and Lewis and Ottaviani(2008), which are reviewed below.

Although replacing agents is costly, it may be an effective way to motivate agents to work faster, to avoid termination. This explains prediction (c) above.

With regards to contracting for search the analyses predict:

(d) Fast performing agents receive a progress payment contract along with an ownership share in the project.

(e) Slow performing agents receive a time of completion contract with a single payment once the search is complete.

The explanation for these findings is developed in section 4. There I find that the timing of discovery determines the form and the level of the agent's rewards. Agents who make early discoveries are awarded a progress payment and the ownership of the remaining project. The principal prefers this type of payment, since it induces the agent to work hard to complete the search remaining. In contrast, agents who are unlucky and progress slowly are paid a smaller amount and only at the time of completion. This leads to diminished and less flexible incentives for the agent to perform that may not be adjusted as the search evolves with each new discovery. These different processes for payment that depend on the timing of success are consistent with *progress payment* contracts and *time of completion* agreements that are often observed in practice.

Virtually all of the extensive search literature, to which this paper belongs, ignores the agency problem on which I focus.[3] One notable exception is the companion paper by Lewis and Ottaviani(2008) that characterizes the effect of delegation on search; it does so, however, for a setting where contracts are short term and subject to renegotiation. A comparison of these analyses reveals the importance of contract commitment in delegating search, an issue that is addressed in the analysis to follow. Another exception is the interesting work on optimal unemployment insurance (see Shavell and Weiss(1979) and Hopenhayn and Nicolini(1997)) that addresses the trade-off between risk sharing and incentives to search in a repeated moral hazard setting. Our analysis, in contrast, is focused on the agent's incentives to search and to disclose the information acquired during the

[3]In 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.

search process. A different form of delegation is analyzed in recent papers by Albrecht, Anderson and Vroman(2008) and by Compte and Jehiel(2008) on search by committee. These analyses compare stopping rules for search by committee against the rules an individual searcher would follow.

This paper is also related to a fast growing literature on dynamic agency. Our analysis heavily relies on recursive programming techniques of solving for optimal dynamic contracts that were pioneered by (Green(1987), Spear and Srivastava(1987), Atkeson(1991) and Phelan and Townsend(1991)). Our application of these techniques to study delegated search is new. Our paper is related to other recent applications including Sannikov(2007)'s continuous time principal-agent problem, Toxvaerd(2006)'s analysis of the importance of time in agency and Mason and Valimaki(2008)'s analysis of repeated moral hazard. These papers all solve a repeated moral hazard problem of some type. Our analysis, by contrast, addresses moral hazard and adverse selection in one unified model. The papers closest to ours from a formal perspective are the dynamic agency models of the firm by Fishman and DeMarzo(2007) and Clementi and Hopenhayn(2006). These analyses focus on the impact of agency problems including moral hazard and hidden information on investment and the capital structure of the firm. Our analysis, in contrast, pertains to the effect of agency on the collection and dissemination of information.

The plan for the remainder of the paper is the following. Section 2 presents an example of simple sequential search to fix ideas and to establish some basic insights about the impact of agency on search. Section 3 extends the example to a continuous time, multi-stage search setting. The formal results on the properties of delegated search including a characterization of optimal contracts for agency search are contained in Section 4. Section 5 reviews the principal predictions of the model and indicates some directions for future research. The Appendix contains proofs of the formal results reported in Sections 3 and 4.

2 Example of Agency Search

Consider the firm owner who wishes to replace an inefficient manufacturing process. Having determined the cause of the inefficiency, the owner has only to discover a remedy to the problem in order to increase her operating profits by the amount of $V_S > 0$. The owner agrees to delegate the search for a solution to a risk neutral agent. By expending effort

$e > 0$ at a cost of ce the agent can discover the solution with probability $p_S(e) = p_S e^{\gamma}$ for $\gamma \in (0,1)$ and $p_S > 0$ in each period[4]. The owner agrees to pay the agent a reward of $\tau > 0$ once the search is complete. For convenience divide time into two periods, *today* denoted by time T, and the *future* denoted by time F, which consists of all periods after *today*. Let $B \in (0,1)$ be the discount factor. Then the agent's expected profit today from this arrangement, denoted by $\pi_T(\tau)$, is the solution to the following dynamic program.

$$\pi_T(\tau) = \max_e p_S(e_T)\tau + (1 - p_S(e_T)) B\pi_F(\tau) - ce_T \tag{1}$$

where

$$\pi_F(\tau) = \frac{p_s(e_F)\tau - ce_F}{1 - B(1 - p_S(e_F))}$$

For a given reward, τ, the agent's optimal rate of effort is determined by the condition,

$$p_s'(e_T)(\tau - B\pi_F(\tau)) = c$$

Solving this condition for, $p_S(e_T)\tau(e)$, the expected reward required to induce effort e is

$$p_S(e_T)\tau(e) = \frac{ce_T}{\gamma} + B\pi_F(\tau) \tag{2}$$

The first term, $\frac{ce_T}{\gamma}$, measures the static agency cost of not being able to monitor the agent's effort. The second term, $B\pi_F(\tau)$, is the agent's opportunity cost of making a discovery today. A discovery today, eliminates the chance of continuing to search in the future. This opportunity that the agent foregoes by discovering a solution is a *dynamic agency cost* that the principal bears to induce the agent to exert effort today.

The owner's surplus denoted by $W_T(\tau)$ is the solution to the dynamic program,

$$W_T(\tau) = \max_{\tau} p_s(e_T)(V_S - \tau) + (1 - p_S(e(\tau))) BW_F(\tau) \tag{3}$$

subject to (2), where,

$$W_F(\tau) = \frac{p_S(e_F)(V_S - \tau)}{1 - B(1 - p_S(e_F))}$$

The condition for the optimal compensation is given by

$$\{p_S'(e)(V_S - B(W_F + \pi_F)) - c\} \frac{de}{d\tau} = \frac{p_S(e_F)}{1 - B(1 - p_S(e_F))} > 0 \tag{4}$$

[4] I assume p_S is sufficiently small so that the probability of discovering a solution is less than 1 over the relevant range of effort that is employed.

This indicates that the marginal benefit of additional search effort (given by the *LHS* of (4)) is strictly positive for the optimal contract.

The implication of this is that agents search too little and too slow under delegation. This occurs for the usual reason that it is too costly to induce the agent to search without monitoring his effort. But here in addition, the agent's incentives to work are further reduced by the opportunity he foregoes to earn future profits from search if he discovers a solution. Since the agent's compensation for finding the solution does not change with time, he can afford to delay his search in the present, knowing there is always another chance to succeed in the future. But what if the agent's compensation was based on the time the search is completed? Could this improve his performance?

Suppose for instance that compensation in period T is increased by $d\tau_T$ and compensation in period F is decreased by $d\tau_F$, holding today's profit constant, so that $d\pi(\tau) = 0$. The required variation $(d\tau_T, d\tau_F)$ must therefore satisfy,

$$\begin{aligned} d\pi_T(\tau) &= \frac{\partial \pi_T}{\partial \tau_T} d\tau_T + \frac{\partial \pi_T}{\partial \tau_F} d\tau_F + \frac{\partial \pi_T}{\partial e_T}\left(\frac{\partial e_T}{\partial \tau_T} d\tau_T + \frac{\partial e_T}{\partial \tau_F} d\tau_F\right) + \frac{\partial \pi_T}{\partial e_F}\frac{\partial e_F}{\partial a_F} d\tau_F \\ &= \frac{\partial \pi_T}{\partial \tau_T} d\tau_T + \frac{\partial \pi_T}{\partial \tau_F} d\tau_F = 0 \end{aligned} \tag{5}$$

where the second line of (5) follows from the envelope theorem. The effect of these variations on the owner's surplus evaluated at $(\tau_T = \tau_F = \tau)$ is[5],

$$dW_T(\tau) = \underbrace{\{p_S'(e)(V_S - B(W_F + \pi_F)) - c\}}_{\text{marginal benefit of effort}} \cdot \underbrace{\left\{\frac{\partial e_T}{\partial \tau_T} d\tau + \frac{\partial e_T}{\partial \tau_F} d\tau_F + B(1 - p_S(e))\frac{de_F}{d\tau_F} d\tau_F\right\}}_{\text{net shift in effort to present}} > 0 \tag{6}$$

That is, the owner increases her surplus, holding π_T constant, by shifting rewards to the present. The change in surplus (6) equals the present value increase in surplus from

[5] Substituting for $p_S(e)\tau_T$ from (1) into the (3) yields

$$W(\tau) = -\pi_T - ce_T + p_S(e_T) V_S + (1 - p_S(e)) B(W_F + \pi_F)$$

Applying the variation $(d\tau_T, d\tau_F)$ to $W_T(\tau)$ I obtain (6) in the text. Notice that $e_T = e_F = e$ since the expression for dW_T is evaluated at $\tau_T = \tau_F = \tau$. The second term in brackets in (6) reduces to

$$\frac{\partial e_T}{\partial \tau_T} d\tau_T + \frac{\partial e_T}{\partial \tau_F} d\tau_F + B(1 - p_S(e))\frac{de_F}{d\tau_F} d\tau_F = d\tau_T \frac{\partial e_T}{\partial \tau_T}\frac{1}{1 - p_S(e)} > 0$$

increasing effort, multiplied by the present value change in effort between today and the future, which is strictly positive. Surplus increases by shifting rewards to the present because the dynamic cost of agency, identified in (2), is reduced. That is, by shifting compensation away from the future, the cost of inducing the agent to perform today is reduced. As a consequence the agent works harder today, knowing her expected future profit is lower. This causes a reallocation of effort to the present, thus accelerating the arrival of surplus, which benefits the owner.

While tilting rewards to the present improves current performance, it causes the future rate of discovery to decrease if the present search fails. A possible fix for this is to switch to a new agent when today's search fails. This would eliminate incentives for the agent to procrastinate, as he has just one chance at making a discovery. However, replacing agents after each failed attempt at discovery is costly as each new agent must be instructed what to search for. If the costs, $c_S > 0$, of preparing each agent are small enough, though, employing multiple short term agents is preferable. To see why, recall the expected reward for inducing effort e for one period from a long lived agent is,

$$p_S(e_T)\,\tau(e_T) = \frac{ce_T}{\gamma} + B\pi(\tau)$$

By comparison, the cost of inducing the same effort from a sequence of multiple agents is

$$p_S(e_T)\,\tau(e_T) = \frac{ce_T}{\gamma} + Bc_S$$

Hence switching agents is preferable to employing one long lived agent if the cost of preparation sufficiently are small.

In summary this example demonstrates several properties of delegation that hold under more general conditions as I show in the sections to follow. Namely there are two agency costs of delegating search There is the usual static agency cost which is the rent the agent commands from his hidden action, and there is a dynamic agency cost which is the opportunity cost of discovery for the long lived agent. Agents wish to procrastinate when hired over the long term to search. The example demonstrates three contractual responses to reduce these agency costs. These responses are what cause delegated search to differ from the text book decision theoretic search in three respects. The first is that delegation reduces the incentives to search due to the added costs of agency. The second is tilting the incentives to search to the present, dissuades the agent from procrastinating. This results in higher discovery rates at first but lower rates in the future as the search continues. The third difference lies in the possible advantage of employing multiple agents. Hiring multiple agents increases switching costs but reduces agency rent over the long run. These findings are summarized in:

Proposition 0: *Delegated search differs from efficient decision theoretic search in three respects. (a) The rate of search is slower, (b) the rate of discovery is decreasing with the duration of the search and (c) search costs may decline when multiple agents are employed for short periods.*

The causes and consequences of these contractual responses to delegated agency search are explored in detail in the sections to follow.

3 Model of Continuous Time Search with Agents

Building on the example above consider again the owner who contracts with a single risk neutral agent to solve her production problem. I consider multi agent search in section 4. The search for a solution is modeled as a continuous time sampling process. The continuous time formulation provides an easier and more transparent derivation of the key results. Under the process, the agent exerts effort at rate $e \geq 0$ and incurs costs ce to draw ideas at the rate of $\phi(e) = e^{\gamma}$ for $\gamma \in (0,1)$[6]. Given $e > 0$, an idea I is discovered at the rate $\phi(e)\, p_I$ where $p_I > 0$ is the known conditional probability of drawing idea $I = S, D$ with $p_S + p_D = 1$.[7]. An idea is either D a *diagnosis* of the problem, or S, a *solution* for the problem which completes the search. The default state is $I = O$ if no ideas have yet been discovered. Once the problem is solved the owner derives a final benefit of $V_S > 0$. If the problem is first diagnosed the owner's benefit is V_D. Once a solution is discovered the owner receives the residual benefit of $V_S - V_D$.

3.1 The Optimal Agency Search Agreement

Throughout the search, the agent's rate of search effort $e(t)$ and his actual discoveries $\tilde{I}(t)$ are not directly observed. Hence there are both moral hazard and hidden information

[6] The adoption of this functional form streamlines the analysis without limiting the generality of the results.

[7] I abstract from issues of learning about the distribution, that while important for some applications are not critical for this analysis.

problems for the owner to contend with. To induce the agent to search, the owner must promise to reward the agent for the discoveries he discloses, $I(t)$. The agent may disclose all or part of his discoveries so that $I(t) \subseteq \tilde{I}(t)$.[8] Without loss of generality we require full disclosure to be incentive compatible so that $I\left(\tilde{I}(t)\right) = \tilde{I}(t)$ for all t.[9] The owner proposes a long term agreement that promises payments and a share of profits as compensation for the agent's disclosed discoveries. In principle the terms of the agreement could depend on the entire public history of discovery and compensation at each instant. Such agreements would be difficult to analyze, however, because of the large variation in the history of play that can arise. Fortunately, though, it can be shown that all the information contained in any history of previous play is effectively summarized by the agent's continuation profit, $\pi_I(t)$, that the contract promises.[10] To analyze the optimal agreement, we may therefore restrict attention to contracts that use the agent's profit as the state variable, on which the terms of employment and compensation are based.

At the beginning of the project, when no ideas have yet been discovered, the owner proposes the following long term agreement, $\tilde{A} = \{E, \tau_J(\pi_I(t)), e_I(\pi_I(t)), \pi_J(\pi_I(t))\}$, to the risk neutral agent with an known initial wealth endowment of $E \geq 0$. The agreement is a complete plan for the search arrangement, stipulating that:

$$\begin{aligned} E &\geq 0 = \text{the agent's required investment} \\ \tau_J(\pi_I(t)) &\geq 0 = \text{the agent's payment for discovery of } J = S, D \text{ starting from stage } I = O, S \\ e_I(\pi_I(t)) &\geq 0 = \text{the agent's recommended effort rate} \\ \pi_J(\pi_I(t)) &\geq 0 = \text{the agent's continuation profit at stage } J \end{aligned}$$

If the agent accepts the contract, the agreement becomes binding for both parties. The agreement is *feasible* provided it implements the desired plan. This requires the agent to be *obedient*, (O), in supplying the recommended effort; that the agreement is (IC), *incentive compatible,* so that agent fully discloses his discovery; and that it is *individually*

[8] That is the agent may disclose $\{O, D, S\}$ when he discovers S, and he may disclose $\{O, D\}$ when he discovers D.

[9] We appeal to the revelation principle in restricting attention to incentive compatible agreements. The revelation principle holds for this model by the nesting property of Green and Laffont (1986) despite the fact that the agent can only under report his discoveries.

[10] This important insight that continuation surplus encapsulates the relevant information from the history of play is due to Green (1987) and Spear and Srivastava (1987) and Phelan and Townsend (1991). They demonstrate that under mild boundedness conditons, which also hold for our problem, recursive dynamic programing techniques may be employed to solve dynamic moral hazard problems.

rational (IR) for the agent to accept the contract. The agent's profit under a feasible contract is defined recursively by the following:[11]

$$\begin{aligned}
r\pi_I(t) &= \underbrace{-ce(\pi_I(t))}_{\text{flow costs}} + \underbrace{\phi(e(\pi_I(t)))\,\Delta\pi_I(t)}_{\text{discovery appreciation}} + \underbrace{\dot{\pi}_i(t)}_{\text{time appreciation}} && (A)\\
\text{where} &: \ \Delta\pi_I(t) = p_S(\tau_S(\pi_I) - \pi_I) + p_D(\tau_D(\pi_I) + \pi_D(\pi_I) - \pi_I)\cdot 1_{I=0} \\
\text{subject to} &: \\
\phi'(e(\pi_I))\,\Delta\pi_I(t) &= c && (O)\\
I\left(\tilde{I}(t)\right) &= \tilde{I}(t) \ \text{ for all } t && (IC)\\
\pi_{I_0}(0) &\geq E && (IR)
\end{aligned}$$

Condition (A) is an accounting identity that specifies the rate of appreciation in agent profits must be equal to flow costs of effort, plus the appreciation from discovery, given $\Delta\pi_I(t)$, plus the time appreciation in profits, $\dot{\pi}_I(t)$. The agent's profit is defined subject to the conditions of the agreement requiring that the agent supply $e(\pi_I)$ which is guaranteed by (O)[12], that the agent truthfully disclose his discoveries, which requires (IC) and that it be individually rational for the agent to participate which is guaranteed by (IR)[13].

The following Lemma records three properties that must hold for any feasible contract, that is one satisfying (O) and (IC).

Lemma 1: *An agreement $\tilde{A}$ for which $\tau_S(t)$ and $\tau_D(t)$ are non decreasing, is (O) and (IC) iff*

$$\begin{array}{ll}
(a) & p_S\tau_S(\pi_I) + p_D\tau_D(\pi_I)\cdot 1_{I=0} = p_S(\pi_I) + p_D(\pi_I - \pi_D(\pi_I))\cdot 1_{I=0} + \frac{ce(\pi_I)}{\gamma\phi(e)}\\
(b) & \dot{\pi}_I = r\pi_I - \frac{1-\gamma}{\gamma}ce(\pi_I)\\
(c) & \tau_S(\pi_0) = \tau_D(\pi_0) + \tau_S(\pi_D(\pi_0))
\end{array}$$

Condition (a) of Lemma 1 requires that the expected payment be equal to the agent's promised profits plus a rent for supplying effort. Condition (b) specifies the time rate of change in agency profits. Requirements for the agreement to be incentive compatible are given in (c) under the condition that payments τ_S and τ_D are non increasing, which is

[11] In what follows I drop the dependence of e_I on π_I where no confusion exists. The variable $1_{I=J}$ is an indicator varible which equals 1 when $I = J$ and zero otherwise.

[12] Condition (O) is necessary and sufficient for the supply of $e(\pi_I)$ given the concavity of $\phi(e)$.

[13] If $\pi_0 \geq E$, the agent accepts the contract initially. Under the optimal agency contract the agent's continuation profit turns out to be positive so that his continued participation is individually rational.

satisfied on the equilibrium path. Payments that decrease with time induce the agent to disclose discoveries immediately. (c) implies that the payment for finding S starting from O is the same if it is disclosed at once or if it is revealed by a disclosure of D followed by an immediate disclosure of S.

I can now state the optimal contracting problem. The principal's ex ante surplus under a feasible agreement $\tilde{A}$ starting at stage, I_0, with profit, $\pi_{I_0}(0)$, is recursively defined by:

$$rW_{I_0}(\pi_{I_0}) = \underbrace{\phi(e_{I_0})(p_S V_S + p_D V_D \cdot 1_{I=D} - \Delta\pi_{I_o})}_{\text{discovery appreciation}} \underbrace{+\dot{W}_{I_0} + rE}_{\text{time appreciation}} \qquad (P)$$

The owner's surplus satisfies the accounting identity (P) requiring that the rate of surplus appreciation equals the discovery appreciation, plus the time rate of appreciation in surplus. Using Lemma 1 to substitute for payments and profits, the optimal contract can be written as the solution to the following problem

$$\max_{\pi_{I_o} \geq E} rW_{I_o}(\pi_{I_o}) \qquad (PP)$$

where

$$\begin{aligned} rW_I(\pi_I) &= \max_{e_I} \phi(e_I)(p_S(V_S - V_D \cdot 1_{I=D} - \pi_I - W_I(\pi_I)) \\ &\quad + \phi(e_I) p_D (V_D + W_D(\pi_D) + \pi_D - \pi_I - W_I(\pi_I)) \cdot 1_{I=0} \\ &\quad + \frac{dW_I}{d\pi_I}\left(r\pi_I - \frac{(1-\gamma)}{\gamma} ce_I\right) - \frac{ce_I}{\gamma} + rE \end{aligned}$$

The principal's problem, $[PP]$, is to choose the optimal initial profit, π_{I_o} and the effort rate, $e_I(t)$ to maximize the surplus flow at each instant. Surplus consists of the current value flow of income minus costs, the combined rate of appreciation from discovery, and the surplus change due to appreciation in the agent's profits. The optimal contract that solves the principal's problem is described in Section 4.

4 Optimal Delegated Search

4.1 Simple Search

This section explores the properties of delegated search and the contracts employed to implement the search. As noted by the example in section 2, simple search is the easiest assignment to delegate. In contrast to more complicated searches simple search just requires the agent to find the solution to a well defined problem that has already been diagnosed. With the stage of the search already known, there is no problem of inducing the agent to reveal his progress. Moreover the choice of compensation is straightforward, the agent is paid a fee once he discovers a solution. Hence, managing the search is essentially a repeated moral hazard problem where the goal of the principal is to pay the agent just enough to perform without giving away too much surplus.

This pay for performance trade-off is evident from Figure 1 which depicts the attainable set of surpluses $(W_D(\pi_D), \pi_D)$ generated from a single search starting at stage D. The domain of feasible profits (gross of E) is $\pi_D \in [0, \pi_D^*]$. The lower bound on agent profit is 0 which equals his outside income. The upperbound on profit, π_D^*, equals the maximum surplus that the search can generate. This is what the agent would receive if he purchased the project with an initial investment of $E = \pi_D^*$. For smaller investments the agent is allocated enough profit, $\pi_D \geq E$ to insure his participation in the search. It follows from routine dynamic programing arguments that $W_D(\pi_D)$ is concave, that it is increasing for small π_D and decreasing for π_D sufficiently large.[14] The $W_D(\pi_D)$ boundary lies inside of the Pareto frontier $W_D^*(\pi_D)$ except at the point $(W_D^*(\pi_D^*)\, \pi_D^*)$ where the agent essentially owns the project. As the agent's rent falls from π_D^* total surplus decreases as the agent searches less. However the principal's rent $W_D(\pi_D)$ increases because she retains a larger share of total surplus. The principal's surplus continues to increase with reductions in agent rent until $\pi_D = \hat{\pi}_D$ where $W_D(\pi_D)$ is maximized. For profit less than $\hat{\pi}_D$ the principal's surplus decreases towards zero.

The allocation of surplus indicated by Figure 1 reflects several of the important properties of delegated search. The most apparent characteristic is how the efficiency of the search is constrained by the amount of the agent's investment E. Wealth constrained agents will not search as efficiently when their surplus share is rationed by the size of their investment. Agents with no wealth will nonetheless be promised a positive profit of $\hat{\pi}_D$ since the owner's expected surplus is maximized by allocating that profit to the agent. Investment greater than $\hat{\pi}_D$ will induce the principal to offer the agent a greater share of the surplus which will cause the agent to search more efficiently. This inevitable rationing of surplus implies another property of the search arrangement; it must be governed by a long term contract that may not be renegotiated. This follows because the surplus frontier is

[14] See Spear and Srivastava (1987) for example.

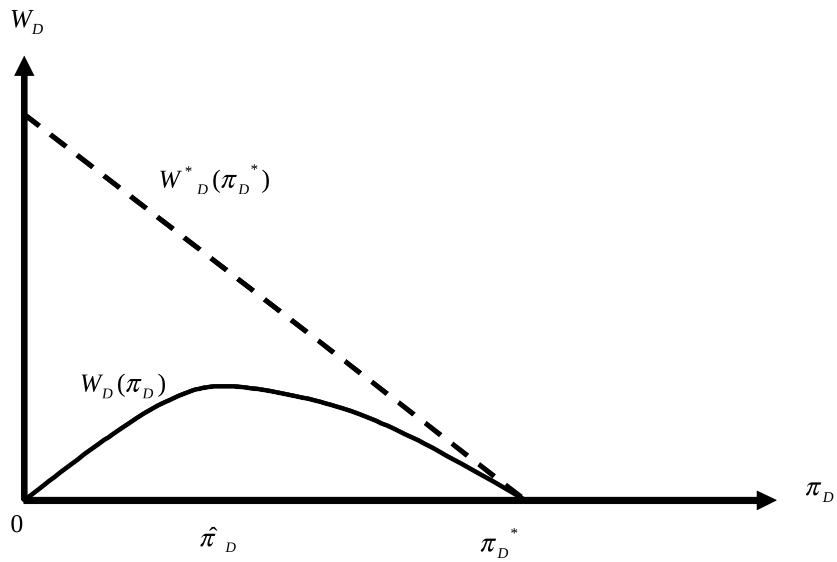

Figure 1: Surplus Allocations

not everywhere downward sloping, implying by the arguments of Fudenberg, Holmstrom, and Milgrom(1989) that the set of agent incentives can not be offered within the set of incentive compatible efficient contracts.[15] Moreover the contracts for implementing delegated search will appear quite different than the renegotiation-proof arrangements that would otherwise govern the search as described in Lewis and Ottaviani(08).

A more detailed analysis of delegated simple search yields the following:

Proposition 1: *Suppose $E = \pi_D^*$. Then (i) The agent is paid the full surplus for discovery of a solution with $\tau_S(t) = V_S$, and (ii) The agent works efficiently with $e_D(t) = e_D^*$ to complete the search.*

Proposition 2: *Suppose $E < \pi_D^*$. Then (i) The agent is paid a share of the search*

[15]Hence inefficent allocations lying on the upward sloping porton of the surplus frontier can not be improved without increasing the agent's surplus. The parties will commit not to increase the agent's surplus in order to force him to perform.

surplus, $\tau_S(t) < V_S$, which decreases with the time of completion, so that $\dot{\tau}_S(t) < 0$, (ii) The agent's search rate is less than efficient, such that $e_D(t) < e_D^*$, (iii) The agent's search rate decreases with time $\dot{e}_D(t) < 0$ and (iv) The agent's compensation and induced rate of search are smaller the less he invests.

Combining the findings of Proposition 1 with 2 illustrate the important effects of delegation on the process of search. Figure 2 depicts the difference in compensation arising in delegated search. When the agent posts a large enough investment to buy the project, he is compensated with the full surplus from search such that $\tau_S(t) = V_S$ and therefore he allocates effort efficiently throughout time until the solution is discovered. In contrast when the agent's investment is constrained by his initial wealth he is offered just a share of the surplus, such that $\tau_S(t) < V_S$. As a result he searches too slowly and the search takes too long to complete. These are the usual types of static inefficiencies resulting from delegation, wherein the agent expends too little effort because he bears all of the costs but receives only a share of the surplus he creates. Predictions like these on the effects of delegation on search are confirmed for example, by Levitt and Syverson(forthcoming) who find that realtors who market their own homes typically obtain higher sales prices in less time than comparable client-owned houses that they market.

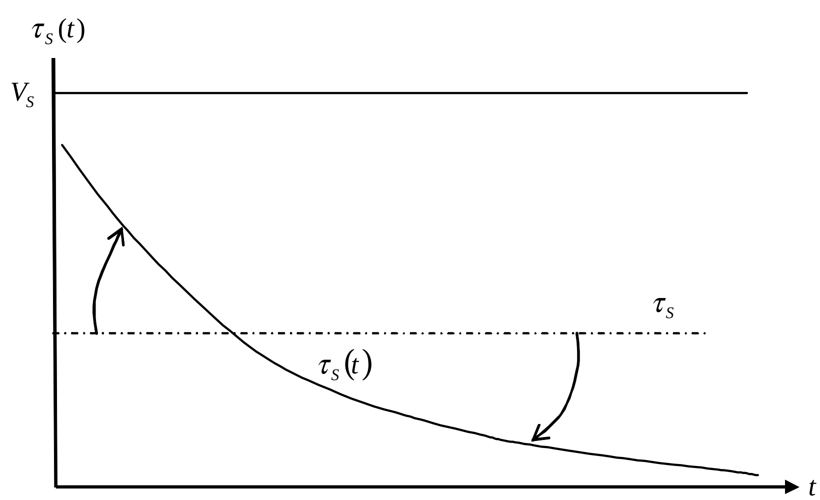

Figure 2: Simple Search Compensation

At the same time that search incentives are lower, they are also tilted towards the present away from the future under delegation. Figure 2 shows how a constant reward of τ_S is transformed into a time of completion payment $\tau_S(t)$ for rewarding the agent for early discovery and penalizing him for delay. As a result, the agent responds by increasing his search in the near term and decreasing it in future periods. Unlike the classical decision theoretic search that calls for steady work until completion, work is accelerated under delegation with the result that the likelihood of completion decreases the longer the search proceeds.

The intuition for this result stems from the observation that agents have a tendency to procrastinate in long term agreements. Consider the agent's incentives to search when offered a small stationary rewards. Since it is increasingly costly for the agent to search at a faster rate, he delays searching today when he can search in the future and receive the same reward.[16] By tilting payments to the present, the owner induces the agent to work harder in the present period, or otherwise receive smaller rewards in the future. While there is a cost to accelerating the search, in so far as effort is allocated less efficiently and future discovery is slower, overall the principal benefits by forcing an earlier completion of the search.

This is not the first analysis to point out the benefits of performance forcing in agency relationships. Work by Holmstrom and Milgrom (1991) and Baker (1992) on multi tasking focuses on the benefits of forcing agents to work on specified jobs to reduce their risk and the cost of inducing performance. By analogy, time of completion contracts work by focusing the agent's attention on searching today instead of the future which reduces the cost of agent procrastination in delegated search. In a different context to this Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997) show that performance forcing may arise in social benefit programs. Penalizing workers for long spells of unemployment by decreasing their benefits may reduce insurance costs and decrease the duration of unemployment. The present analysis demonstrates the benefits of basing compensation on the time of completion in order to align the intertemporal incentives of the agent and principal.

[16] The agent must expend increasingly more effort to increase her rate of success, since the discovery rate is concave in effort.

4.2 Multiple Stage Search

While many tasks delegated to agents, like finding a home buyer, are simple search activities, other jobs that agents are asked to perform, like solving a complicated communications or management problem are clearly not. For these more complex assignments, the agent searches in stages, first to discover a diagnosis for the problem and second (and sometimes simultaneously) to find a solution. Whereas it is fairly easy for the principal to delegate a simple search, managing a multistage agency search is likely to be more difficult. With multiple phases the owner wants the search to sequentially adjust with each new discovery. However only the agent knows what he has found. Consequently the owner must choose how to compensate the agent to induce him to disclose his progress as well as to modify his effort as the search transitions to the next stage.

There are two types of compensation contracts for multi stage search. Let t_I be the time $I = D, S$ is discovered and denote by $t_d = \min[t_D, t_S]$ the time of the first discovery.

Definition: *Time of Completion contract.* A time of completion contract pays the agent a fee $\tau_S(t_S)$ only at the time a solution is discovered.

Definition: *Progress Payment Contract:* A progress payment contract pays the agent, $\tau_D(t_D)$ and $\tau_S(t_S)$ at the times of discovery for a diagnosis and for a solution.

Payment under the time of completion contract depends only on the time when the search is completed. As a consequence, the agent's search rate can not be adjusted sequentially in response to a new discovery. For instance such adjustments are desirable after the discovery of a diagnosis, to account for the reduction in residual benefits from further search. In contrast, progress payment contracts are more refined, in that they allow for payments to be tailored to the current stage of search. Adjusting compensation in response to a discovery induces the agent to change his rate of search at different stages.

The choice of compensation between time of completion and progress payments is a matter of how much the agent is expected to profit from the search.

Proposition 3: *In the solution to* $[PP]$

(i) If $\pi_0(t_d) > \pi_D^*$ *then the optimal agreement is a progress payment contract with:*

$$\begin{aligned} \pi_D(\pi_0(t_d)) &= \pi_D^* \\ \tau_D(\pi_0(t_d)) &> 0 \\ \tau_S(\pi_0(t_d)) &= \tau_D(\pi_0(t_d)) + \tau_S(\pi_D^*) \end{aligned}$$

(ii) If $\pi_0(t_d) \leq \pi_D^*$ *then the optimal agreement is a time of completion contract with:*

$$\tau_S(t_S) > 0, \tau_D(\pi_0(t_d)) = 0$$

Proposition 3 reveals that there is a "pecking order" for compensating the agent. The form of reward is determined by the agent's equity in the current search as measured by his profit, $\pi_0(t_d)$ at his first discovery. By part (a) if the first discovery is a diagnosis, D, the preferred compensation for the agent is a full equity share π_D^* in the next stage of the search. In addition the agent receives a progress payment of τ_D financed from the left over funds, $\pi_o(t_d) - \pi_D^*$. However, when there is insufficient equity to fully fund the next stage search, the agreement reverts to a time of completion contract, by part (b). In this case the agent receives only a terminal payment of $\tau_S(t_S)$ at the time of completion, even though this is not the owner's preferred form of compensation.

The important role of the agent's wealth in delegating search is manifested in the preference for compensation that Proposition 3 reveals. When there are sufficient funds, the owner prefers to pay with equity. This is in contrast to relational contracts where it is matter of indifference how agents are compensated, (see Levin (2003)). Here, however, paying the agent with equity insures his performance in the next stage search since he owns the remaining phase of the development. In effect, when the search begins the agent is "working to own". His primary compensation is "sweat equity", a payment in kind for discovering a diagnosis. The owner defers as much of the agent's compensation as possible to bind him to perform efficiently in the latter stages of search. When there are insufficient funds to fully fund the next stage of search, the owner's next best option is to defer all payment until the end. This is provided for in a time of completion contract.

The foregoing analysis on the importance of agency wealth and timing in delegation, implies that an agent's search performance is largely a matter of luck. That is the agent's level and form of compensation as well as his success at search are determined by whether his first discovery occurs before or after a given deadline $\hat{t}_E$. This intuition is confirmed in the two Propositions that follow.

Proposition 4: *Early Discovery Search: Suppose* $t_d < \hat{t}_E$. *Then*

(i) *For* $t < t_d$ *continuation profits and incentives to search are decreasing with* $\dot{\pi}_0(t) < 0$ *and* $\dot{e}_0(t) < 0$

(ii) *At* t_d *search is governed by a progress payment contract such that*

$$\begin{aligned} \pi_D(\pi_0(t_d)) &= \pi_D^* \\ \tau_D(\pi_0(t_d)) &> 0 \\ \tau_S(\pi_0(t_d)) &= \tau_D(\pi_0(t_d)) + \tau_S(\pi_D^*) \end{aligned}$$

(iii) *The deadline* $\hat{t}_E$ *is increasing in the agent's initial investment* E.

The lucky agent defined in Proposition 4, is one whose first discovery occurs before the deadline $\hat{t}_E$. That agent receives greater payment provided by the progress payment contract and the earlier the discovery the larger is the agent's compensation. If the agent's first discovery is a solution, she receives the same compensation as if she first discovered a diagnosis followed in the next instant by a discovery of the solution. Otherwise if the diagnosis is first discovered, the agent's reward is a progress payment plus ownership of the remaining stage of search. Notice, the early discovery search is similar to a contract with job tenure; the agent is on probation until his first discovery, whereafter he is permitted to work without a deadline or pressure to perform. Moreover the agent who makes a greater initial investment in the search is afforded a longer period to prove himself

In contrast the unlucky agent who misses the discovery deadline is governed by the following agreement.

Proposition 5: *Late Discovery Search: Suppose* $t_d \geq \hat{t}_E$. *Then search is governed by a time of completion contract. The payment for completion and the incentives to search are decreasing with* t, *as* $\dot{\tau}_S(t) < 0$ *and* $\dot{e}(t) < 0$.

The late performing agent is destined to operate under a low powered-time of completion contract that offers less incentive to perform the longer the search proceeds. It is not only the agent but also the principal who is unlucky in this case. This is because the continuation search following the deadline is both slow and coarse as it is not possible to influence the agent's rate of search after a discovery with just a single time of completion payment. The rationale for employing this second best contract is that it serves as a threat to the agent who procrastinates and therefore runs the risk of missing the deadline.

4.3 Multiple Agent Search

Viewed in its entirety, the multistage search arrangement described in Propositions 4 and 5 is a performance forcing measure with rewards based on the time of discovery. The longer the agent takes for his first discovery the smaller is his progress payment and the less his continuation value in the ongoing search. The difficulty of employing a forcing contract to search with a single agent is that his incentives to search are diminished the longer the search lasts. The principal is betting the agent will complete the search early. But if the agent is unlucky in his initial search, the likelihood he will succeed in future periods grows smaller with each failed attempt.

One way to avoid this unfortunate outcome is to employ a sequence of agents to search, rather than just one. Suppose for instance there is a supply of identical wealth constrained agents each with the same ability to search. The owner employs the agents one at a time in sequence until one of the agents succeeds in completing the search. At the start of each search, the agent is promised a continuation profit of $\pi_0 \geq E$ that is large enough to cover the agent's initial investment. The agent is given a limited period of time, $\hat{t}_{C_s}$ to make his first discovery, otherwise he is replaced by another agent to resume the search. There is a switching cost of $c_S > 0$ the principal incurs in informing and training a new agent to search. Assuming this cost is not too large, the process continues until one agent succeeds in meeting the deadline $\hat{t}_{C_s}$ whereupon he proceeds to finish the search[17]. The properties of this multi agent search are characterized in the following:

Proposition 6: *For c_S sufficiently small and $E < \pi_0^*$ the optimal multi-agent contract is a sequence of short term progress payment contracts exhibiting these features: (i) Each agent begins his search with profit endowment $\pi_0(0) \geq E$, (ii) Until the first discovery, search rewards and incentives are decreasing with $\dot{\pi}_0 < 0, \dot{\tau}_S(t) < 0, \dot{\tau}_D(t) <$*

[17] If switching costs are too large, it may be impractical to use multiple agents so that only single agent search is employed.

0 *and* $\dot{e}_0 < 0$, *(iii) Agents who fail to discover by time* $\hat{t}_{C_s}$ *are replaced,* (iv) *At the time of replacement* $W_0\left(\pi_0\left(\hat{t}_{C_s}\right)\right) = W_0\left(\pi_0\left(0\right)\right) - c_S$..

Proposition 6 demonstrates the benefits and costs of employing multiple agents in a setting where efficient search requires only one agent to work. Part (i) indicates that with multi agent search, each agent starts out under a forcing contract in which his rewards and incentives to search decline with time. The longer he takes to make a discovery, the smaller are his payments for progress and he therefore searches less the longer the search proceeds. This is similar to the single agent search in that the agent is under pressure to perform at least until he make his first discovery. However, unlike single agent search, agents who are slow to progress are replaced with new agents under a mult agent search. The rationale for replacing one agent with another one of equal ability stems from the tendency for agents to procrastinate in dynamic settings. That is a long-lived agent, who is hired to search regards the effort he will expend in the future as a substitute for the effort he exerts today. The agent exerts less effort in the present knowing he can also search in the future. In effect the agent's incentives to work now versus later are in conflict.[18] Hence, one benefit of employing multiple agents to search, each for a limited time, is to avoid this intertemporal conflict. Each agent becomes more impatient to succeed and thus expends more effort, when his time to search is limited. In essence, the *efficiency wage* that agents command to perform (e.g. Shapiro and Stiglitz(1984)) is reduced by limiting their tenure.

A second benefit from using multiple agents stems from the "fresh start" each new agent receives when he is hired. Property (iii) indicates the incumbent agent is replaced once his performance deteriorates to the point where a new agent with a fresh start can perform better. By replacing a failing agent before his incentives to search fall too far, the owner avoids the slow down in search that otherwise occurs with single agent search. However, there is a cost of employing multiple agents. Each replacement agent must be trained and informed about the search at a cost of c_S to the principle. This cost is balanced against the benefits from replacement to determine the time $\hat{t}_{C_s}$ at which a new agent is brought in. Agents are replaced less frequently, the larger are the replacement costs, so that multi agent search evolves into a single agent search once costs grow sufficiently large.

[18] See for example Holmstrom and Milgrom (1991) on multi tasking, and Laux (2001) on adverse selection with multiple projects.

5 Conclusion

"If you want the job done right....do it yourself!" This is one message of agency theory, but it's certainly not the most important one. The theory predicts that the agent's performance is constrained by the principal's ability to monitor his effort and progress and by the contracts that the parties are able to enforce. The delegation of search to agents is no exception in this regard. Our analysis predicts that, delegated search will be slower and take longer to complete; that the rate of discovery will be greatest at the start of search and will decrease the longer the search lasts; and that search may be inefficiently organized and utilize more agents than is necessary. Given the extent of delegated search, these findings would appear to have important implications for the efficiency of exchange in matching markets, a prime topic for future research.

Several factors shape the agent's behavior and performance in delegated search. Recognizing what these factors are and assessing how search changes in different settings provides some promising directions for future research. For instance, it's clear that the ability of parties to commit to contracts is important in determining search behavior. The analysis of this paper with Lewis and Ottaviani(2008) consider the extreme cases where commitment is either complete or nonexistent. An interesting avenue for future research is to explore delegation in intermediate settings where the parties may partially commit, either for limited periods of time or for some sets of contract provisions, as is the case in public sector procurement for example. The "type" of search that is delegated is another factor likely to affect performance. For instance, searching for "novelty" in the context of innovative activity (see for instance Manso(2006)) or searching for the "truth" in legal processes (see for instance Lester, Persico and Visschers (2008)) is likely to impact the delegation process differently. Finally the "sort" of agents that are assigned to search, whether it be one individual, a committee or a body of voters, will invariably affect how the search is conducted, the information is collected, and the decisions that are made. The recent work on search by committee by Albrecht, Anderson and Vroman(2008) and Compte and Jehiel(2008) is a promising step in that direction.

6 Appendix

PROOF OF LEMMA 1

Part (a) and (b): Substituting for $\phi(e) = e^{\gamma}$ into condition (O) and rearranging terms I obtain,

$$\Delta\pi_I = p_S\left(\tau_S(\pi_I) - \pi_I\right) + p_D\left(\pi_D\left(\pi_I\right) + \tau_D(\pi_I) - \pi_I\right) \cdot 1_{I=0} = \frac{ce_I}{\gamma\phi\left(e_I\right)} \tag{A1}$$

which implies

$$p_S\tau_S(\pi_I) + p_D\tau_D(\pi_I) \cdot 1_{I=0} = p_S\pi_I + p_D\left(\pi_I - \pi_D\left(\pi_I\right)\right) \cdot 1_{I=0} + \frac{ce_I}{\gamma\phi\left(e_I\right)}$$

Using (O) and $(A1)$ in the expression for $\dot{\pi}_I$ we obtain,

$$\begin{aligned} \dot{\pi}_I &= r\pi_I - \max_{e_I}\phi\left(e_I\right)\Delta\pi_I - ce_I \\ &= r\pi_I - \frac{(1-\gamma)}{\gamma}ce_I \end{aligned}$$

Part (c): Let $\pi_I(J,t)$ denote the agent's expected profit who has previously disclosed I and is currently at stage J at time t. If $J \neq I$ then (IC) requires that the agent's best response is to disclose J so that:

$$\begin{aligned} \pi_D(S,t) &= \max_{J\in\{D,S\}}\pi_J(S,t) = \tau_S\left(\pi_D\left(t\right)\right) && \text{(A2)} \\ \pi_0(S,t) &= \max_{J\in(0,D,S)}\pi_J(S,t) = \tau_S\left(\pi_O\left(t\right)\right) && \text{(A3)} \\ \pi_0(D,t) &= \max_{J\in(0,D)}\pi_J(S,t) = \tau_D\left(\pi_0(t)\right) + \pi_D\left(\pi_0(t)\right) && \text{(A4)} \end{aligned}$$

$(A2)$ is automatically satisfied when $\tau_S(t)$ is decreasing with time. Consider $(A3)$. A necessary and sufficient condition for the agent to report S directly instead of first reporting D and second reporting S is that

$$\tau_S\left(\pi_0\left(t\right)\right) = \tau_D\left(\pi_0\left(t\right)\right) + \tau_S\left(\pi_D\left(\pi_0\left(t\right)\right)\right) \tag{A5}$$

provided τ_S and τ_D are both decreasing. Finally consider $(A4)$. A sufficient condition for full disclosure is $(A5)$ which implies the agent can not benefit by delaying her disclosure of D while she discovers a solution, S.■

PROOF OF PROPOSITIONS 1 AND 2

For the simple search case search starts in stage D where I assume $V_D = 0$. In this instance the principal's problem $[PP]$ becomes

$$\max_{\pi_D(0)\geq E} rW_D(\pi_D(0)) \tag{B1}$$

where

$$\begin{aligned} rW_D(\pi_D) &= \max_{e_D} \phi(e_D)(p_S(V_S - \pi_D - W_D(\pi_D)) \\ &\quad + \lambda_D \left(r\pi_D - \frac{(1-\gamma)}{\gamma} ce_D \right) - \frac{ce_D}{\gamma} + rE \end{aligned} \tag{B2}$$

The problem is formally treated as a piecewise continuous stochastic control problem that is solved using control theoretic techniques (e.g. Dockner et al (2000), Chpt.8). The control variable for this problem is $e_D(t)$; the state variable is $\pi_D(t)$; the co state variable is $\lambda_D = dW_D/d\pi_D$ and the equation of motion for the state variable is $\dot{\pi}_D = r\pi_D - \frac{(1-\gamma)}{\gamma}ce_D$. The current valued Hamiltonian for this problem is $H(e_D, \pi_D, t) = rW(\pi_D)$ and the necessary Pontryagin conditions for the solution to $(B1) - (B2)$ are:

$$H_{e_D} = \phi'(e_D)\, p_S(V_S - (\pi_D + W_D(\pi_D))) - \frac{c}{\gamma} - \frac{c(1-\gamma)\lambda_D}{\gamma} = 0 \tag{B3}$$

$$\dot{\lambda}_0 = r\lambda_0 - H_{\pi_0} = \phi(e_D)p_S(1+\lambda_D) \tag{B4}$$

$$0 = \lim_{t\to\infty} e^{-rt}\lambda_D(t)\,\pi_D(t) \tag{B5}$$

and the optimal choice of $\pi_D(0)$ must satisfy:

$$\lambda_D(0)(\pi_D(0) - E) = 0 \tag{B6}$$

Differentiating $(B3)$ with respect to time and using $(B4)$ I obtain the following expression for $\dot{e}_D$,

$$\dot{e}_D = \frac{-\phi(e_D)(1-\gamma)\gamma^2}{c(1+(1-\gamma)\lambda_D)} \left\{ (1+\lambda_D)\left(r\pi_D + \frac{(1-\gamma)^2}{\gamma^2} ce_D \right) \right\} \tag{B7}$$

The expression for $\dot{\pi}_D$ is

$$\dot{\pi}_D = r\pi_D - \frac{(1-\gamma)}{\gamma} ce_D \tag{B8}$$

The system dynamics for this problem are described by $(B7)$ and $(B8)$. The relevant phase space for this simple search problem is the set $\{(e_0, \pi_0) \mid \pi_0 \in [0, \pi_D^*], e_0 \geq 0\}$ as indicated in Figure 3.

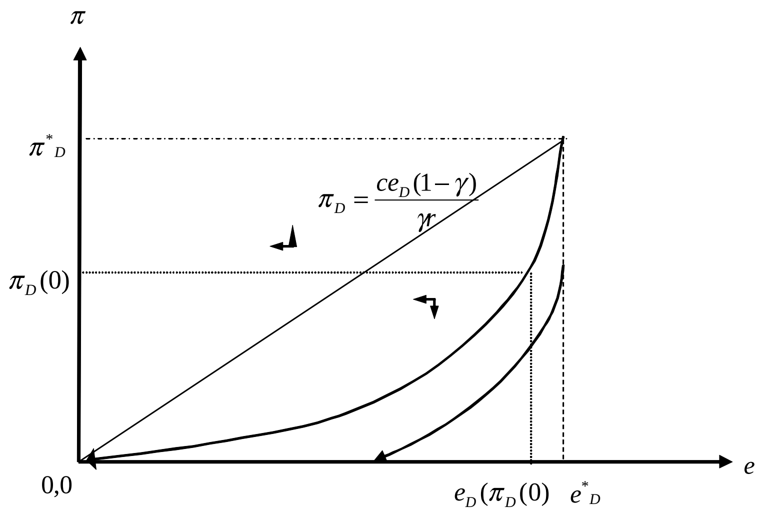

Figure 3: Simple Search Phase Diagram

There are two cases to consider which correspond respectively to Proposition 1 and to Proposition 2. First consider the setting of Proposition 1 where $E = \pi_D^*$. For that case $(B6)$ indicates that $\pi_D(0) = E = \pi_D^*$. Furthermore $(1 + \lambda_D) = 0$ when $\pi_D(t) = \pi_D^*$, so that $\dot{\lambda}_D = 0$ and $\dot{e}_D(t) = 0$ by $(B7)$ and $(B8)$. This implies that $e_D(t\ ) = e_D^*$ and $\pi_D(t) = \pi_D^*$ for all t. This completes the proof of Proposition 1.■

For the second setting of Proposition 2, suppose $E < \pi_D^*$. In that instance, $(B6)$ implies that $\pi_D(0) = \max[\hat{\pi}_D, E] < \pi_D^*$. From $(B4)$ it follows that $\dot{\lambda}_D(t) > 0$ for all t as long as $e_D(t) > 0$ and $1 + \lambda_D(t) > 0$. From $(B7)$ this implies that $\dot{e}_D(t) < 0$ as indicated by the arrows in Figure 3. According to $(B8)$

$$\dot{\pi}_D \left\{ \begin{array}{c} > \\ = \\ < \end{array} \right\} 0 \text{ as } \left\{ \begin{array}{c} \pi_D > \frac{1-\gamma}{\gamma r} c e_D \\ \pi_D = \frac{1-\gamma}{\gamma r} c e_D \\ \pi_D < \frac{1-\gamma}{\gamma r} c e_D \end{array} \right\} \tag{B9}$$

Thus π_D is increasing (decreasing) for points above (below) the $\pi_D = \frac{1-\gamma}{\gamma r} c e_D$ line in Figure 3.

The corresponding direction of motion at any point (e_D, π_D) in the phase diagram is indicated by the arrows in Figure 3. It's clear that any initial $\pi_D(0) < \pi_D^*$ there exists a unique trajectory starting at $(e_D(\pi_D(0)), \pi_D(0))$ that converges to $(0,0)$ which is the unique rest point that satisfies the transversality condition $(B5)$. Along this trajectory, $\dot{e}_D(t) < 0$ and $\dot{\pi}_D < 0$. Moreover, differentiating condition (a) of Lemma 1 with respect to time, I obtain the expression for $\dot{\tau}_S(t)$

$$\dot{\tau}_S(t) = \dot{\pi}_D(t) + \dot{e}_D(t)\frac{c(1-\gamma)}{\gamma^2 p_S} < 0$$

Consequently, along the optimal path the reward for discovery of a solution is declining with time. Finally notice the starting point moves further up the trajectory leading to $(0,0)$, the greater is the initial investment, for $E > \hat{\pi}_D$. This implies that the agents rewards and incentives to search are greater the more he invests initially. This completes the proof of Proposition 2.■

PROOF OF PROPOSITIONS 3-5

Propositions 3-5 are proved in the following sequence. I first solve the general multi-stage search problem $[PP]$ posed in section 3. Then I prove Proposition 3 which provides conditions for the existence of progress payments contracts and time of completion contracts. This is followed by a proof of Proposition 4 which characterizes the properties of progress payment contracts. The analysis concludes with a proof of Proposition 5.

SOLUTION TO MULTI STAGE SEARCH [PP]

For multi stage search the principal's problem $[PP]$ is,

$$\max_{\pi_{I_o}(0)\geq E} rW_{I_o}(\pi_{I_o}(0))$$

As in the simple search case this problem is formally treated as a piecewise continuous stochastic control problem that is solved using control theoretic techniques (e.g. Dockner et al (2000), Chpt.8). There are two stages O and D to this problem that I consider in sequence.

<u>**Stage** O :</u> For the stage O problem the controls are $e_0(t)$,$\tau_D(\pi_0), \tau_S(\pi_0)$ and $\pi_D(\pi_0)$, and the state is $\pi_0(t)$. The current valued Hamiltonian $H(e_0, \tau_s, \tau_D, \pi_D, t) = rW_0(\pi_0)$ for

the problem is:

$$\begin{aligned}H\left(e_0,\tau_s,\tau_D,\pi_D,t\right) &= rE-\frac{ce_0}{\gamma}\\ &+\phi\left(e_0\right)p_S(V_S-\pi_0-W_0\left(\pi_0\right))\\ &+\phi\left(e_0\right)p_D\left(V_D+W_D\left(\pi_D\right)+\pi_D-\pi_0-W_0\left(\pi_0\right)\right)\\ &+\lambda_0\left(r\pi_0-\frac{(1-\gamma)}{\gamma}ce_0\right)\\ &+\rho\left(\tau_S\left(\pi_0\right)-\tau_D\left(\pi_0\right)-\tau_S\left(\pi_D\left(\pi_0\right)\right)\right)\\ &+\eta\left(p_S\tau_S(\pi_0)+p_D(\tau_D\left(\pi_0\right)+\pi_D\left(\pi_0\right))-\pi_0-\Delta\pi_0\right)\\ &+\mu_{\tau_s}\tau_S\left(\pi_0\right)+\mu_{\tau_D}\tau_D\left(\pi_0\right)+\mu_{\pi_D}\pi_D\left(\pi_0\right)\end{aligned}$$

where λ_0 is the multiplier for $\dot{\pi}_0$; ρ is the multiplier for the (IC)) constraint; η is the multiplier to the $\Delta\pi_0$ definition and $\mu_{\tau_s},\mu_{\tau_D}$.and μ_{π_D} are the multipliers for the non negativity constraints for $\tau_S\left(\pi_0\right),\tau_D\left(\pi_0\right)$ and $\pi_D\left(\pi_0\right)$ respectively. The stage O solution is characterized by the following necessary Pontryagin conditions:

$$\begin{aligned}H_{e_0} &= \phi'\left(e_0\right)p_S\left(V_S-\left(\pi_0+W_0\left(\pi_0\right)\right)\right)\\ &+\phi'\left(e_0\right)p_D\left(V_D+W_D\left(\pi_D\right)+\pi_D-\left(\pi_0+W_0\left(\pi_0\right)\right)\right)\\ &-\frac{c}{\gamma}-\lambda_0\frac{1-\gamma}{\gamma}c=0\end{aligned} \tag{ C1}$$

$$H_{\pi_D} = \phi\left(e_0\right)p_D\left(1+\lambda_D\right)-\rho\frac{d\tau_s\left(\pi_D\right)}{d\pi_D}+\eta p_D=0 \tag{C2}$$

$$H_{\tau_s(\pi_0)} = \rho+\mu_{\tau_S(\pi_0)}+\eta p_S=0 \tag{ C3}$$

$$H_{\tau_D(\pi_0)} = -\rho+\mu_{\tau_D(\pi_0)}+\eta p_D=0 \tag{C4}$$

$$\dot{\lambda}_0 = r\lambda_0-H_{\pi_0}=\phi(e_0)p_S\left(1+\lambda_0\right)+\eta \tag{C5}$$

$$0 = \lim_{t\to\infty}e^{-rt}\lambda_0\left(t\right)\pi_0\left(t\right) \tag{C6}$$

and

$$0=\lambda_o\left(0\right)\left(\pi_0(0)-E\right) \tag{C7}$$

is the condition for the optimal initial profit allocation to the agent. These conditions $(C1)-(C4)$ for the stage O solution to $[PP]$ are used to establish Proposition 3.

PROOF OR PROPOSITION 3: Suppose $\tau_D(\pi_0(t_d)) > 0$. Then adding $(C3)$ and $(C4)$ I obtain $\eta = -\mu_{\tau_D} = 0$. By $(C2)$ this implies that $(1+\lambda_D) = 0$ or that $\lambda_D = -1$. Since $\lambda_D = dW_D/d\pi_D$ it follows that $\pi_D(\pi_0) = \pi_D^*$. Hence whenever $\tau_D > 0$ the continuation profit for stage D, $\pi_D(\pi_0) = \pi_D^*$ is the profit that maximizes total surplus. Moreover, it follows from (IC) that $\tau_S(\pi_0(t_d)) = \tau_D(\pi_0(t_d)) + \tau_S(\pi_D^*)$ This proves part (i).

If $\pi_0(t_d) \leq \pi_D^*$ then $\tau_D(\pi_0(t_d)) = 0$, otherwise part (i) of the proposition implies that $\pi_0(t_d) > \pi_D^*$ which is a contradiction. Given $\tau_D(\pi_0(t_d)) = 0$ it follows that $\pi_0(t_d) = \pi_D(\pi_0(t_d))$. This completes the proof of Proposition 3.■

PROOF OF PROPOSITION 4: According to Proposition 3 the optimal agreement is a progress payment contract when $\pi_0(t_d) > \pi_D^*$. In that case payments $\tau_S(t)$ and $\tau_D(t)$ are offered for either the discovery of a solution or a diagnosis. The dynamics of $\tilde{A}$ in this case are described by the differential equations,

$$\dot{e}_0 = \frac{-\phi(e_0)(1-\gamma)\gamma^2}{c(1+(1-\gamma)\lambda_0)}\left\{(1+\lambda_0)\left(r\pi_0 + \frac{(1-\gamma)^2}{\gamma^2}ce_0\right)\right\} \tag{C8}$$

$$\dot{\pi}_0 = r\pi_0 - \frac{1-\gamma}{\gamma}ce_0 \tag{C9}$$

where $(C8)$ is obtained by differentiating $(C1)$ with respect to time. The relevant phase space for progress payment search is the set $\{(e_0, \pi_0) \mid \pi_0 \in (\pi_D^*, \pi_0^*], e_0 \geq 0\}$ as indicated in Figure 4. The direction of motion within that set are indicated by the arrows. Following the previous arguments in the proof of Proposition 2 it is straightforward to show that for any $\pi_0(0)$ there is a unique $e_0(\pi_0(0))$ such that the trajectory beginning at $e_0(\pi_0(0)), \pi_0(0)$ leads to the point (e_D^*, π_D^*) where the time of completion search begins. Along this path $\dot{\pi}_0 < 0$, $\dot{e}_0 < 0$, until the first discovery is made, where upon the search is either completed if S is discovered, or the continuation search is carried out at the efficient rate until a solution is discovered if D is discovered initially.

The deadline $\hat{t}_E$ for the progress payment search is the time it takes to reach (e_D^*, π_D^*) along the optimal trajectory starting at $(e_0(\pi_0(0)), \pi_0(0))$ Since $\pi_D(0)$ is increasing with E for $E > \hat{\pi}_D(0)$ as investment increases the starting point is placed further up on the trajectory.

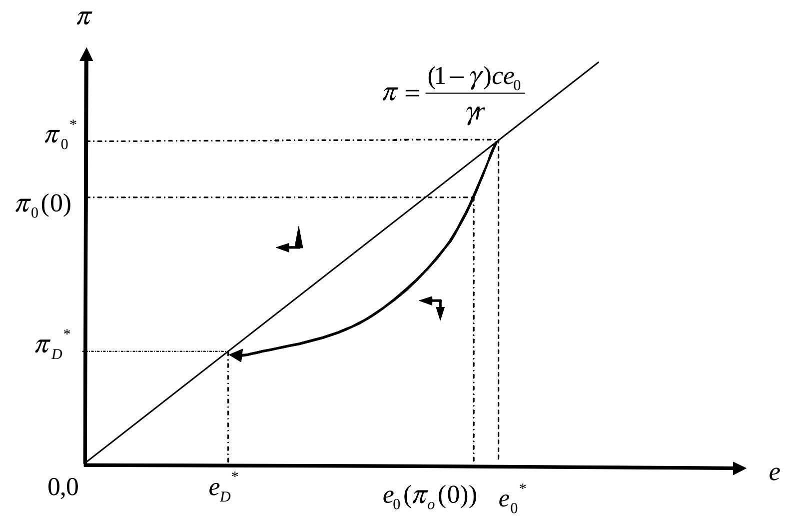

Figure 4: Progress Payment Phase Diagram

Hence the deadline in increasing in E (and strictly so) whenever$E > \hat{\pi}_D(0)$ ■

PROOF OF PROPOSITION 5: Proposition 3 shows that the search is governed by a time of completion contract for $\pi_0(t_d) \leq \pi_D^*$. Under time of completion search the agent receives a single payment, $\tau_S(t)$, once he completes the project. Consequently there is no distinction between stages O and D as the agent is searching only for a solution S. In this case $[PP]$ is conveniently solved in one stage as an open-loop piecewise deterministic control problem. (see Dockner et al (2000), Chpt. 8).

An open loop strategy is one which depends only on time. In particular it does not depend on the state of the agent's search. Hence in order to formulate $[PP]$ as an open loop control problem, I take $e(t)$ as the control variable and $\pi(t)$ as the state variable. In addition it is useful to define another state variable, $\sigma_D(t)$, which is the conditional probability that the search is in phase D, given a solution hasn't yet been found. The probability the search is in state O is $\sigma_0(t) = 1 - \sigma_D(t)$. The equation of motion for the state $\sigma_D(t)$ is determined by taking the time derivative of $\sigma_D(t)$ so that

$$\begin{aligned}
\frac{d}{dt}\sigma_D(t) &= \frac{d}{dt}\left(\frac{\sigma_D(t)}{\sigma_D(t)+\sigma_0(t)}\right) \\
&= \left(\frac{\dot{\sigma}_D(t)\left(\sigma_D(t)+\sigma_0(t)\right)-\left(\dot{\sigma}_0(t)+\dot{\sigma}_D(t)\right)\sigma_D(t)}{\left(\sigma_D(t)+\sigma_0(t)\right)^2}\right) \\
&= \dot{\sigma}_D(t)-\left(\dot{\sigma}_0(t)+\dot{\sigma}_D(t)\right)\sigma_D(t) \\
&= \phi(e)\left(p_D\left(1-\sigma_D(t)\right)\right) \qquad \text{(C10)}
\end{aligned}$$

The equation of motion for $\pi(t)$ is,

$$\dot{\pi}(t) = r\pi(t) - \frac{1-\gamma}{\gamma}ce(t) \qquad \text{(C11)}$$

The current valued Hamiltonian corresponding to this control problem is given by

$$\begin{aligned}
H &= rW(\pi,\sigma_D) \\
&= rE + r\sigma_D V_D - \frac{ce}{\gamma} \\
&\quad +\phi(e)\left(p_S(V_S-\pi-W(\pi,\sigma_D)\right) \\
&\quad +\lambda\left(r\pi-\frac{(1-\gamma)}{\gamma}ce\right) \\
&\quad +\lambda_{\sigma_D}\left(\phi(e)\left(1-\sigma_D\right)p_D\right)
\end{aligned}$$

where λ and λ_{σ_D} are the two co state variables attached to $\dot{\pi}$ and $\dot{\sigma}_D$.

The necessary Pontryagin conditions for the solution to $[PP]$ satisfy

$$\begin{aligned} 0 &= \phi'(e)\left(p_S\left(V_S - \left(\pi + W\left(\pi, \sigma_D\right)\right) + \lambda_{\sigma_D}\left(1-\sigma_D\right)p_D\right)\right) \\ &\quad -\frac{c}{\gamma} - \frac{\lambda c(1-\gamma}{\gamma} \end{aligned} \tag{C12}$$

$$\dot{\lambda} = r\lambda - H_\pi = \phi(e)p_S\left(1+\lambda\right) \tag{C13}$$

$$\dot{\lambda}_{\sigma_D} = r\lambda_{\sigma_D} - H_{\sigma_D} = r\left(\lambda_{\sigma_D} - V_D\right) + \phi\left(e\right)\left(p_S + p_D\right)\lambda_{\sigma_D} \tag{C14}$$

$$0 = \lim_{t\to\infty} e^{-rt}\lambda\pi\left(t\right) \tag{C15}$$

$$0 = \lim_{t\to\infty} e^{-rt}\lambda_{\sigma_D}\sigma_D \tag{C16}$$

The dynamics for this problem are given by the four differential equations $(C10), (C11)\,(C13)\,(C14)$. Using phase diagrammatic techniques to find the solution to this problem is cumbersome because there are two state variables and hence four differential equations to be analyzed. Nonetheless it is possible to characterize the solution by examining the evolution of the agent's effort and profit through time. Differentiating $((C12))$ with respect to time and using $(C13), (C14)$ we get an expression for the time rate of change in effort,

$$\dot{e} = \frac{-\phi(e)\left(1-\gamma\right)\gamma^2}{c\left(1+\left(1-\gamma\right)\lambda_0\right)}\left\{\left(1+\lambda\right)\left(r\pi + \frac{\left(1-\gamma\right)^2}{\gamma^2}ce\right) + p_D\left(1-\sigma_D\right)r\left(V_D - \lambda_{\sigma_D}\right)\right\} \tag{C17}$$

It follows that $\lambda_{\sigma_D}(t) \leq V_D$ since $\lambda_{\sigma_D}(t) = \frac{\partial W_{\sigma_D}}{\partial \sigma_D}$ is the shadow value of increasing the likelihood that a payment of V_D has been received. Therefore $(C17)$ implies $\dot{e} < 0$.

Suppose there is some $t < \infty$ such that $\dot{\pi}(t) \geq 0$. By $(C11)$ this implies that

$$\pi\left(t\right) \geq \frac{1-\gamma}{\gamma r}ce\left(t\right) \tag{C18}$$

Since $\dot{e}(t) < 0$ and $\dot{\pi}(t) \geq 0$, $(C18)$ implies that for $t' > t$,

$$\pi\left(t'\right) > \frac{1-\gamma}{\gamma r}ce\left(t'\right)$$

and therefore that $\pi(t'') > \dot{\pi}(t') > 0$ for all $t'' > t'$. Consequently $\lim_{t''\to\infty} \pi(t'') = \infty$ which is impossible. Therefore it follows that $\dot{\pi}(t) < 0$ for all t. This implies that

$$\dot{\tau}_s(t) = \dot{\pi}(t) + \frac{c(1-\gamma)\dot{e}}{\gamma^2 p_s} < 0$$

This completes the proof of Proposition 5■

PROOF OF PROPOSITION 6: The multi agent search problem requires the principal to maximize the search surplus from the current agent plus the expected value of surplus from replacing the incumbent with a new agent at time $\hat{t}_{c_s}$. To perform the maximization, the principal select controls $\{e_0, \tau_S, \tau_D, \pi_D\}$, a terminal time $\hat{t}_{c_s}$ and an initial profit $\pi_0(0)$ in order to

$$max_{\{e_0,\tau_S,\tau_D,\pi_D\},\hat{t}_{c_s},\pi_0(0)} \left\{ \int_0^{\hat{t}_{c_s}} e^{-rt} H(e_0, \tau_S, \tau_D, \pi_D, t)\, dt + e^{-r\hat{t}_{c_s}} (W_0(\pi_0(0)) - c_S) \right\} \tag{D1}$$

where $H(e_0, \tau_s, \tau_D, \pi_D, t)$ is the current value Hamiltonian for $[PP]$ and $e^{-r\hat{t}_{c_s}}(W_0(\pi_0(0)) - c_S)$ is the continuation surplus of employing a new agent. The necessary conditions for the maximization of $(D1)$ includes $(C1)-(C4)$ plus the condition for the optimal replacement time, (see Seierstad and Sydsaeter, (Thm. 13, pgs. 390-91,1987)

$$H\left(e_0^*, \tau_s^*, \tau_D^*, \pi_D^*, \hat{t}_{c_s}\right) - r(W_0(\pi_0(0)) - c_S) = 0 \tag{D2}$$

where $H\left(e_0^*, \tau_s^*, \tau_D^*, \pi_D^*, \hat{t}_{c_s}\right)$ is the maximized Hamiltonian at time $t = \hat{t}_{c_s}$. Noting that $H\left(e_0^*, \tau_s^*, \tau_D^*, \pi_D^*, \hat{t}_{c_s}\right) = rW_0\left(\pi_0\left(\hat{t}_{c_s}\right)\right)$ and rearranging $(D2)$ we obtain,

$$W_0\left(\pi_0\left(\hat{t}_{c_s}\right)\right) = (W_0(\pi_0(0)) - c_S) = 0 \tag{D3}$$

This proves part (iv) of the Proposition.

The behavior of this system is characterized by the same equations of motion $(\dot{e}, \dot{\pi}_0)$ that govern the progress payment search of Proposition 4. Therefore it can be shown that the controls $e_0, \tau_s, \tau_D, \pi_D$ have the same properties. In particular, $\dot{\tau}_s < 0$, $\dot{\tau}_D < 0$, and $\dot{e}_0 < 0$ until the first discovery or until the incumbent agent is replaced. This completes the proof of Proposition 6.■

References

Albecht, James, Axel Anderson and Susan Vroman. 2008. "Search by Committee," mimeo

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

Atkeson, Andrew. 1991. "International Lending with Moral Hazard and Risk of Repudiation." *Econometrica*, 59(4): 1069–1089.

Baker, George P. 1992. "Incentive Contracts and Performance Measurement." *Journal of Political Economy*, 100(3): 598–614.

Clementi, Gian Luca, and Hugo A. Hopenhayn. 2006. "A Theory of Financing Constraints and Firm Dynamics." *Quarterly Journal of Economics*, 121(1): 229–265.

Compte, Olivier and Phillip Jehiel. 2008. Bargaining and Majority Rules: Collective Search Perspective, mimeo

DeMarzo, Peter M., and Michael J. Fishman. 2007. "Optimal Long-Term Financial Contracting." *Review of Financial Studies*, 20(5): 2079– 2128.

Dockner, Englebert, Steffen Jorgensen, Ngo Van Long, and Gerhard Sorger. 2000. *Differential Games in Economics and Management Science*, Cambridge University Press.

Fudenberg, Drew, Bengt Holmström, and Paul Milgrom. 1990. Short-term Contracts and Long-term Agency Relationships. *Journal of Economic Theory*. 51(1): 1–31.

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.

Green, Jerry R., and Jean-Jacques Laffont. 1986. "Partially Verifiable Information and Mechanism Design." *Review of Economic Studies*, 53(3): 447–456.

Holmström, Bengt, and Paul Milgrom. 1991. "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design." *Journal of Law, Economics, and Organization*, 7: 24–52.

Hopenhayn, Hugo A., and Jean Pablo Nicolini. 1997. “Optimal Unemployment Insurance.” *Journal of Political Economy*, 105(2): 412-438.

Laux, Christian. 2001. “Limited-Liability and Incentive Contracting with Multiple Projects.” *RAND Journal of Economics*, 32(3): 514–526.

Lester, Benjamin, Nicola Persico, and L. Visschers 2008 "Information Acquisition and the Exclusion of Evidence in Trails," mimeo

Levin, Jonathan. 2000. “Relational Incentive Contracts.” *American Economic Review*, 93(3): 835-847.

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, and Marco Ottaviani. 2008. “Search Agency.” Mimeo.

Mason, Robin and Juuso Valimaki. 2008 "Dynamic Moral Hazard and Project Completion" mimeo

McCall, John J. 1970. “Economics of Information and Job Search.” *Quarterly Journal of Economics*, 84(1): 113–126.

Manso, Gustavo 2007 "Motivating Innovation", Sloan School of Management, MIT mimeo

Mortenson, Dale 1970 "Job Search, the Duration of Unemployment and the Phillips Curve" *American Economic Review,* 60(3): 847-62

Phelan, Christopher, and Robert M. Townsend 1991. “Computing Multi-period, Information-Constrained Optima.” *Review of Economic Studies*. 58(5): 853–81.

Rogerson, Richard, Shimer and Wright (2005), "Search Theoretic Models of Labor," *Journal of Economic Literature*, Vol XLIII pp. 959-988

Rogerson, William. 1985. “Repeated Moral Hazard.” *Econometrica*, 53(1): 69–76.

Sannikov, Y. (forthcoming) "A Continuous Time Version of the Principal-Agent Problem," *Review of Economic Studies*

Seierstad, Atle and Knut Sydsaeter. 1987 *Optimal Control Theory and Economic Applications* North Holland

Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review*, 74(3), 433–444.

Shavell, Steven, and Laurence Weiss. 1979. "Optimal Payment of Unemployment-Insurance Benefits over Time." *Journal of Political Economy*, 87(6): 1347–1362.

Spear, Stephen E., and Sanjay Srivastava. 1987. "On Repeated Moral Hazard with Discounting." *Review of Economic Studies*, 54(4): 599–617.

Toxvaerd, Flavio. 2006 "Time of the Essence," *Journal of Economic Theory* 129 252-72