

WORKING PAPERS

N° TSE-475

March 2014

“Optimal Delegated Search with Adverse Selection and Moral Hazard”

Robert Ulbricht

Optimal Delegated Search with Adverse Selection and Moral Hazard*

Robert Ulbricht

Toulouse School of Economics

March 11, 2014.

Abstract

The paper studies a model of delegated search. The distribution of search revenues is unknown to the principal and has to be elicited from the agent in order to design the optimal search policy. At the same time, the search process is unobservable, requiring search to be self-enforcing. The two information asymmetries are mutually enforcing each other; if one is relaxed, delegated search is efficient. With both asymmetries prevailing simultaneously, search is almost surely inefficient (it is stopped too early). Second-best remuneration is shown to optimally utilize a menu of simple bonus contracts. In contrast to standard adverse selection problems, indirect nonlinear tariffs are strictly dominated.

Keywords: adverse selection, bonus contracts, delegated search, moral hazard, optimal stopping.

JEL Classification: D82, D83, D86, C72.

*I am extremely grateful to Klaus M. Schmidt for extensive feedback and guidance throughout this project. I also thank Guy Arie, Dirk Hofmann, Niko Matouschek, and Caspar Siegert for useful comments and suggestions. Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 is gratefully acknowledged. *Email:* robert.ulbricht@tse-fr.eu.

1 Introduction

Searching is an important aspect of many agency relationships. To name a few, recruiting agencies are hired to search for job candidates. Real estate agents are contracted to search for prospective tenants or housing. And insurance brokers are employed to attract new clients. More generally, many forms of problem-oriented thinking require searching for ideas or solutions. This includes research centers searching for new product ideas, advocates thinking about strategies to defend their clients, and business consultancies (or managers) looking for profitable business strategies.

This paper analyzes optimal searching when it is delegated to an agent. The model is based on the standard (single-agent) search model, in which a problem-solver sequentially samples “solutions” from a time-invariant distribution until she is satisfied (McCall, 1970; Mortensen, 1970). Departing from the standard search model, I study equilibrium search when the revenues are not collected by the problem-solver but by a distinct principal.

I consider two information asymmetries governing the relationship between the problem-solving agent and the principal. First, motivated by the aforementioned examples, I model the agent as an expert who has an *ex ante* informational advantage over the principal in assessing the prospects of searching. For instance, real estate agents are likely better in assessing the likelihood that a house sells at a certain price compared to house owners; recruiting agencies are likely to be better informed about the chances of finding qualified candidates than their clients; *et cetera*. In an effort to capture this notion of asymmetry, I assume that payoffs x are sampled from a time-invariant but state-dependent distribution $F(x|\theta)$, where θ is privately known by the agent. Second, I assume that the search process itself cannot be observed (or verified) by the principal. This second asymmetry reflects that many search routines are either intrinsically unobservable (e.g., thinking for ideas), or are hard to be verified due to their soft and easily manipulatable nature (e.g., sampling a *genuine* buyer).

In this search environment the precise configuration of information frictions is crucial to the delegated search process. If either of the two asymmetries occurs in isolation, then the efficient benchmark can be sustained as a delegation equilibrium. This holds true independent of cash constraints or the risk attitude of the agent. If, however, both asymmetries prevail simultaneously, then each acts as a catalyst to the

other one, and search is almost surely inefficient (it is stopped too early).¹

A natural question then is: how should one design the contractual arrangements to achieve second-best optimality?

Confronted with both asymmetries, the challenge is to bring the agent to reveal the optimal search policy (which depends on θ) and, at the same time, to induce her to also search according to the revealed policy. It turns out that the second-best optimum can be implemented via a menu of simple bonus contracts. Each contract pays a fixed bonus when a previously specified target is reached, and nothing otherwise. Other information about the realized outcome is optimally ignored.

Underlying this result is an (endogenous) misalignment between the perceived returns to searching between the agent and the principal. Contracts that are more sensitive to the outcome of the search process than bonus contracts are shown to increase this misalignment in states where the agent is least inclined to truthfully reveal the optimal search policy. Accordingly, more complex contracts make it (weakly) more costly to learn the optimal search policy. Simpler contracts, on the other hand, are precluded by the requirement to incentivize the agent to search according to the revealed policy. The same logic also rules out indirect tariffs that only condition on the realized search revenues.

This result provides a novel angle to the common practice of using bonus schemes rather than fully state-contingent schedules to set incentives (e.g., Moynahan, 1980; Churchill, Ford and Walker, 1993). It thereby complements a small literature that seeks to explain why real world compensation schemes are often simpler than standard theories would suggest.² In particular, Herweg, Muller and Weinschenk (2010) have recently demonstrated that bonus schemes are optimal if agents are averse to losses relative to an expectation-based reference point.³

More generally, the paper relates to a large literature focusing on the delegation of

¹Put into context of the principal-agent literature, the combination of asymmetries works similarly to how cash constraints and risk-aversion unleash the moral hazard in traditional settings (Holmstrom, 1979; Innes, 1990). A difference is, however, that the unleashing of asymmetries in the search environment works both ways, a point that becomes important for the second-best contracting scheme (see below).

²Seminal examples include Holmstrom and Milgrom (1987), Townsend (1979) and Innes (1990) rationalizing linear compensations schemes and standard debt contracts.

³See also Park (1995), Kim (1997), Demougin and Fluet (1998) and Oyer (2000) showing that bonus schemes are “knife-edge” optimal under limited liability if agents are exactly risk-neutral, while they are generally suboptimal if agents are risk-averse to only the slightest degree (Jewitt, Kadan and Swinkels, 2008).

certain tasks subject to contracting constraints. The delegation of search has recently been explored by Lewis and Ottaviani (2008) and Lewis (2012).⁴ Lewis and Ottaviani, however, study search over a long-term horizon, using techniques from the dynamic moral hazard literature (Toxvaerd, 2006; Sannikov, 2008). In these environments search revenues are decreasing in time and the central challenge is to induce the agent to search at the right *speed*.

In this paper, in contrast, search is taken to take place during a comparatively short time span and the main challenge is to learn (and induce) the preferred stopping rule. From the view of the principal, the difficulty lies thereby in disentangling an *ex ante* poor distribution of search revenues from a poorly chosen search policy. At a more technical level, this aspect of our paper is closely related to the theory of regulation and procurement where similar combinations of asymmetries are often taken into account (e.g., Laffont and Tirole, 1986, 1993; McAfee and McMillan, 1987).⁵

On the empirical side, the efficiency of search agencies has been studied in the context of the real estate industry. In line with the findings in this paper, the literature documents that search spells are generally inefficiently short (Rutherford, Springer and Yavas, 2005; Levitt and Syverson, 2008).

The remainder of this paper is organized as follows. Section 2 presents the model. Section 3 provides the first-best benchmark and shows how it can be implemented as an equilibrium outcome if only one of the two asymmetries is active. Section 4 analyzes the solution to the model with both information asymmetries, and Section 5 concludes. All proofs are confined to the appendix.

2 A simple model of delegated search

There are two parties, a principal and an agent. Both parties are risk-neutral and have unlimited access to cash.⁶ The principal hires the agent to operate a search technology

⁴For an overview of the (non-delegated) search literature, see Mortensen (1986) and Rogerson, Shimer and Wright (2005). Multi-agent variations include Albrecht, Anderson and Vroman (2010) and Compte and Jehiel (2010) looking at non-homogeneous search committees that jointly decide over the continuation of a search process.

⁵Several other literatures are based on similar combinations of asymmetries, perhaps most notably the seminal Mirrlees (1971) framework. See Bolton and Dewatripont (2005, ch. 6) for an overview.

⁶Both assumptions can be relaxed without affecting the results. Maintaining them helps simplifying notation, avoids dealing with (irrelevant) corner solutions, and highlights that any inefficiency emerging in this search environment is unrelated to limited liability and risk-sharing.

that yields a monetary outcome $x \in X = [0, B]$. The agent samples outcomes at constant (non-monetary) costs $c > 0$ from a twice differentiable cumulative distribution function $F(x|\theta)$, where θ is an exogenous state with support $\Theta = [\underline{\theta}, \bar{\theta}]$. The prior cumulative distribution function of θ is common knowledge, is denoted by P , and has a differentiable density p such that $p(\theta) > 0$ for all $\theta \in \Theta$. Each time the agent samples an outcome, she can either stop search and select any previously sampled outcome, or continue searching. If she selects an outcome, the principal collects its monetary value, the agent receives her remuneration, and the game ends. The outside option from not selecting any outcome and from not contracting is normalized to zero for both parties. I restrict attention to the case where, in the absence of information asymmetries, solving the problem is profitable in all states ($\mathbb{E}\{x|\theta\} \geq c$ for all $\theta \in \Theta$).

To conclude the model, let me specify an information structure. I consider two information asymmetries.

Assumption A1 (Adverse Selection). *The state θ is privately revealed to the agent before she contracts with the principal. The principal knows the set of potential states Θ and their distribution $P(\theta)$.*

Assumption A2 (Moral Hazard). *Search by the agent and the sampled selection of outcomes cannot be observed by the principal. The value of the selected outcome is observable and verifiable.*

3 Benchmark cases

For reference, I first describe the full information (first-best) benchmark and examine the special cases where only one of the two information asymmetries prevails.

3.1 Full information

Under full information the principal reaps the (joint) surplus and implements the search policy that maximizes it. This is merely the standard search model. I skip the derivation and simply state the result. For details, see e.g. McCall (1970).

Proposition 1. *In the first best, the agent searches as long as for all previously sampled outcomes it holds that $x \leq \bar{x}^{FB}(\theta)$. Otherwise she stops search and selects the last-sampled outcome. The first-best stopping rule is defined by a mapping $\bar{x}^{FB} : \Theta \rightarrow X$,*

such that \bar{x}^{FB} satisfies:

$$c = \int_{\bar{x}^{FB}(\theta)}^B (x' - \bar{x}^{FB}(\theta)) dF(x'|\theta). \quad (1)$$

Under full information, the problem is separated across states. Conditional on θ , search continues until the agent samples a solution of at least value $\bar{x}^{FB}(\theta)$. The optimal “stopping rule” $\bar{x}^{FB}(\theta)$ is hereby chosen to equate the marginal expected benefit of finding a better outcome than $\bar{x}^{FB}(\theta)$ (the right-hand side of equation (1)) with the marginal (social) cost of continuing search c .

3.2 Only adverse selection

Consider now the case where the principal is able to observe (and verify) the sampled selection of outcomes, and only faces uncertainty from not knowing the state θ (Assumption A1 holds but not A2). In this case, the first-best search policies can be implemented by exactly compensating the agent for her search costs. Because this makes her payoffs effectively independent of the pursued search policy, she is indifferent and finds it (weakly) optimal to adopt the first-best policy. I state the precise result in the following.

Proposition 2. *Suppose Assumption A1 holds, but the principal is able to verify the sampled selection of outcomes. Then the first-best search policies can be implemented by specifying a transfer $T(N)$ from the principal to the agent, where $T(N) = Nc$, and N is the number of outcomes in the final sample.*

Proof Sketch. Consider the first-best equilibrium where the agent accepts the contract, pursues first-best search policies, and selects the outcome with the highest value. Under the proposed contract T , the agent breaks even independently of her search behavior, preventing any profitable deviation. Moreover, since returns to the principal define the joint first-best surplus, $x - Nc$, they are maximal (subject to the agent’s participation constraint), precluding any profitable deviation of the principal. *Q.E.D.*

With only adverse selection, the principal is able to construct a contract, in which the agent’s private knowledge about the state θ is not payoff-relevant to her. The agent is therefore willing to reveal the state without any explicit incentives. Critical to this contract is that the principal is able to verify the sampled selection of outcomes,

allowing him to assess the actual search costs of the agent. This is precisely what is prevented by Assumption A2. Under moral hazard, the principal can only form an expectation about how often the agent sampled before selecting an outcome, preventing him to differentiate a poor distribution of outcomes (caused by θ) from an too early termination of search by the agent. In this sense, Assumption A2 “unleashes” Assumption A1 by rendering the agent’s private knowledge of θ necessarily payoff-relevant (for any non-trivial contracting).

3.3 Only moral hazard

A similar conclusion holds regarding the flipside scenario where Assumption A2 holds but not Assumption A1. Again the first-best search policies can be implemented via a simple contractual arrangement. Perhaps the most obvious approach is to make the agent residual claimant, as it is then in her own interest to maximize the joint surplus.⁷

Proposition 3. *Suppose Assumption A2 holds, but the principal learns θ prior to contracting the agent. Then first-best search policies can be implemented by specifying a transfer $T(x)$ from the principal to the agent, where $T(x) = -\bar{x}^{FB}(\theta) + x$.*

Proof Sketch. Under the proposed contract, returns to the agent, $-\bar{x}^{FB}(\theta) + x - Nc$ define the joint surplus up to a constant. Search is hence efficient. There remains the question whether both parties agree to the specified “price” $\bar{x}^{FB}(\theta)$. To verify this, note that $\bar{x}^{FB}(\theta)$ defines the first-best expected surplus (conditional on θ). Hence the principal reaps all the surplus, and therefore happily proposes this contract, which the agent accepts in equilibrium since she breaks even. *Q.E.D.*

Again, this arrangement is not feasible if both asymmetries co-exist. The reason is that with θ unknown the price $\bar{x}^{FB}(\theta)$ will be subject to adverse selection in the original sense (Akerlof, 1970). Accordingly, adverse selection triggers the moral hazard problem in the search environment similar to how risk aversion and limited liability unleash moral hazard in traditional principal-agent settings.

⁷This requires that the agent has sufficient cash reserves and is not too risk-averse. However, there exists an alternative scheme similar to the optimal scheme in the next section that here implements the first best without violating limited liability and without imposing any monetary risk on the agent.

4 Adverse selection and moral hazard

I now turn to the case where both asymmetries co-exist. The challenge for the principal then becomes to design an incentive scheme that brings the agent to reveal her knowledge of the state θ and, at the same time, induces her to search according to the search policies that the principal finds optimal given θ .

Let a contract be a (possibly state-contingent) mapping $T_\theta : X \rightarrow \mathbb{R}$, which specifies, for every outcome $x \in X$, a transfer from the principal to the agent. Under Assumption A2 it is clear that all incentives to search have to be self-enforcing given T_θ . Taking into account the mapping from T_θ to search policies, the principal's objective is to maximize expected search revenues net of transfers. By the revelation principle, a solution to this problem may be obtained via a direct revelation mechanism in which the agent truthfully reports the state θ , and for each θ is assigned a contract T_θ . The principal's problem is then to find the optimal set of contracts $\{T_\theta\}_{\theta \in \Theta}$.

I approach this problem as follows. Since any contract T_θ effectively designs a search problem from the perspective of the agent, I first characterize the agent's optimal search policy for an arbitrary contract. With this implementability constraints at hand, I then examine the optimization problem of the principal and obtain some defining properties of the optimal menu. In particular, I establish the optimality of bonus contracts. After simplifying the problem accordingly, I lastly solve for the optimal menu $\{T_\theta\}_{\theta \in \Theta}$ and derive the equilibrium search policies.

4.1 Implementability constraints

Once the agent has chosen a contract $T_{\bar{\theta}}$ from the menu offered to her, sequential rationality requires that she pursues the search policy which is then optimal for her. Since the agent is effectively facing a search problem over the transfers $T_{\bar{\theta}}(x)$ specified by the chosen contract, equilibrium search is characterized by the solution to this search problem. The following lemma states the solution.

Lemma 1. *An agent with distribution θ and contract $T_{\bar{\theta}}$ searches as long as for all previously sampled outcomes it holds that $T_{\bar{\theta}}(x) \leq \bar{T}_{\bar{\theta}}(\theta)$. Otherwise she stops search and selects the last-sampled outcome. Let $\psi_{\bar{\theta}} : X \times \mathbb{R} \rightarrow \{0, 1\}$ be an indicator function,*

such that

$$\psi_{\bar{\theta}}(x, \bar{T}_{\bar{\theta}}(\theta)) = \begin{cases} 0 & \text{if } T_{\bar{\theta}}(x) \leq \bar{T}_{\bar{\theta}}(\theta), \text{ and} \\ 1 & \text{if } T_{\bar{\theta}}(x) > \bar{T}_{\bar{\theta}}(\theta). \end{cases}$$

Then the stopping rule is defined by a mapping $\bar{T}_{\bar{\theta}} : \Theta \rightarrow X$, such that $\bar{T}_{\bar{\theta}}(\theta)$ satisfies

$$c = \int \left(\{T_{\bar{\theta}}(x') - \bar{T}_{\bar{\theta}}(\theta)\} \cdot \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta)) \right) dF(x'|\theta), \quad (2)$$

whenever $\int T_{\bar{\theta}}(x') dF(x'|\theta) \geq c$. Otherwise the agent does not search at all.

Similar to the first-best case, the optimal stopping rule $\bar{T}_{\bar{\theta}}$ equates the marginal cost of searching c with the marginal expected benefits from finding a better outcome. However, in contrast to the first best, the value of searching from the perspective of the agent is now defined by $T(x)$ rather than x . For what is coming next, it will be crucial to formulate the solution to the agent's problem in terms of outcomes $x \in X$.⁸ To ensure that $T_{\bar{\theta}}(x)$ maps back into a unique solution in X , I therefore impose the following restriction.

Assumption A3. *Contracts are monotonically increasing, i.e. $T_{\theta}(x') \leq T_{\theta}(x'')$ for all $(x', x'', \theta) \in \{X^2 \times \Theta \mid x' \leq x''\}$.*

It is well known that this assumption can be rationalized by the possibility of free disposal; i.e., the ability of the agent to freely downscale any realized outcome x .⁹ Under Assumption A3, inverting $\bar{T}_{\bar{\theta}}$ then immediately defines a stopping rule in X , given by,

$$\bar{x}(T_{\bar{\theta}}, \theta) = \max_x \{x : T_{\bar{\theta}}(x) \leq \bar{T}_{\bar{\theta}}(\theta)\}. \quad (3)$$

The following proposition formulates the resulting implementability constraints by defining $\bar{x}(T_{\bar{\theta}}, \theta)$ directly as a function of $T_{\bar{\theta}}$ (eliminating the intermediate dependence on $\bar{T}_{\bar{\theta}}$).

⁸This guarantees that the Spence-Mirrlees property holds with respect to any change in the payment scheme dT . See Footnote 11 for details.

⁹To see this note that with free disposal the agent can guarantee herself a payoff of $T_{\theta}^*(x) \equiv \max_{x' \in [0, x]} \{T_{\theta}(x')\}$. Hence, w.l.o.g., one could replace T_{θ} by \hat{T}_{θ} , which for all x , pays $\hat{T}_{\theta}(x) = T_{\theta}^*(x)$, whereas it can be easily verified that \hat{T}_{θ} is indeed increasing in x .

Proposition 4. *Suppose Assumptions A2 and A3 hold. Let \mathcal{M} be the space of monotonically increasing functions $X \rightarrow \mathbb{R}$. Then search is determined by a function $\bar{x} : \mathcal{M} \times \Theta \rightarrow X$, which specifies, for a contract $T_{\bar{\theta}} \in \mathcal{M}$ and a state $\theta \in \Theta$, a number $\bar{x}(T_{\bar{\theta}}, \theta)$, such that the agent searches as long as for all previously sampled outcomes it holds that $x \leq \bar{x}(T_{\bar{\theta}}, \theta)$. Otherwise she stops search and selects the last-sampled outcome. The stopping rule \bar{x} is defined by the following inequalities.*

$$c \leq \int_{\hat{x}}^B (T_{\bar{\theta}}(x') - T_{\bar{\theta}}(\hat{x})) dF(x'|\theta) \quad \text{for all } \hat{x} \leq \bar{x}(T_{\bar{\theta}}, \theta) \quad (4a)$$

$$c > \int_{\hat{x}}^B (T_{\bar{\theta}}(x') - T_{\bar{\theta}}(\hat{x})) dF(x'|\theta) \quad \text{for all } \hat{x} > \bar{x}(T_{\bar{\theta}}, \theta). \quad (4b)$$

4.2 Contract properties

I am now ready to characterize the problem from the perspective of the principal. The optimal menu of contracts $\{T_{\theta}\}_{\theta \in \Theta}$ is given by the solution to the following maximization problem:

$$\max_{\{T_{\theta}\}_{\theta \in \Theta}} \left\{ \int_{\theta \in \Theta} \int_{\bar{x}(T_{\theta}, \theta)}^B \left(\frac{x' - T_{\theta}(x')}{\bar{F}(\bar{x}(T_{\theta}, \theta)|\theta)} \right) dF(x'|\theta) dP(\theta) \right\}$$

subject to the constraints,

$$\frac{1}{\bar{F}(\bar{x}(T_{\theta}, \theta)|\theta)} \left[\int_{\bar{x}(T_{\theta}, \theta)}^B T_{\theta}(x') dF(x'|\theta) - c \right] \geq 0 \quad (IR_{\theta})$$

$$\begin{aligned} \frac{1}{\bar{F}(\bar{x}(T_{\theta}, \theta)|\theta)} \left[\int_{\bar{x}(T_{\theta}, \theta)}^B T_{\theta}(x') dF(x'|\theta) - c \right] \\ \geq \frac{1}{\bar{F}(\bar{x}(T_{\bar{\theta}}, \theta)|\theta)} \left[\int_{\bar{x}(T_{\bar{\theta}}, \theta)}^B T_{\bar{\theta}}(x') dF(x'|\theta) - c \right] \quad (IC_{\theta, \bar{\theta}}) \end{aligned}$$

for all $(\theta, \bar{\theta}) \in \Theta^2$, where $\bar{x}(T_{\bar{\theta}}, \theta)$ is characterized by

$$c \leq \int_{\hat{x}}^B (T_{\bar{\theta}}(x') - T_{\bar{\theta}}(\hat{x})) dF(x'|\theta) \quad \text{for all } \hat{x} \leq \bar{x}(T_{\bar{\theta}}, \theta) \quad (SP_{\theta, \bar{\theta}}^-)$$

$$c > \int_{\hat{x}}^B (T_{\bar{\theta}}(x') - T_{\bar{\theta}}(\hat{x})) dF(x'|\theta) \quad \text{for all } \hat{x} > \bar{x}(T_{\bar{\theta}}, \theta). \quad (SP_{\theta, \bar{\theta}}^+)$$

The objective of the principal here is to maximize his expected payoff subject to three kinds of constraints. First, constraints (IR_θ) require that it must be individually rational for the agent in state θ to accept contract T_θ , rather than choosing her outside option. Second, constraints ($IC_{\theta,\tilde{\theta}}$) require that it must be optimal for the agent in state θ to truthfully reveal the state to the principal by choosing T_θ from the menu of all contracts $\{T_{\tilde{\theta}}\}_{\tilde{\theta}\in\Theta}$. These constraints stem from the principal not knowing the state θ . The third set of constraints reflect the requirement to also incentivize search by the agent. As follows from Proposition 4, ($SP_{\theta,\tilde{\theta}}^-$) and ($SP_{\theta,\tilde{\theta}}^+$) pin down the stopping rule $\bar{x}(T_{\tilde{\theta}},\theta)$ implemented in state θ under contract $T_{\tilde{\theta}}$.

Inspecting ($SP_{\theta,\tilde{\theta}}^-$) and ($SP_{\theta,\tilde{\theta}}^+$) it can be seen that for a general contract T_θ , the implemented search policy depends on the state $\tilde{\theta}$. In particular, off-equilibrium search policies $\bar{x}(T_\theta,\tilde{\theta})$ do not necessarily equal $\bar{x}(T_\theta,\theta)$. The key question guiding the following analysis will be: what search policy should a contract designed for state θ implement in state $\tilde{\theta}$? An intuitive answer here is that the principal wants to design a contract T_θ that implements a search policy in all states $\tilde{\theta}$ other than θ that makes it as unattractive as possible in the relevant subset of these states to choose contract T_θ . In the following I will show that bonus contracts achieve this goal.

Before proving this claim, let me impose some structure on the distribution of outcomes $F(x|\theta)$ and states $P(\theta)$. Let $\bar{F} \equiv 1 - F$, and let $H \equiv \partial\bar{F}^{-1}/\partial x$. Then:

Assumption A4. $\partial H/\partial\theta \leq 0$, and $\partial^2 H/\partial\theta^2 \geq 0$.

Assumption A5. $\partial H/\partial x \leq 0$, and $\partial^2 H/\partial x\partial\theta \leq 0$.

The first part of Assumption A4 introduces a stochastic ordering over distributions in θ . A sufficient condition for H to be decreasing is the commonly used monotone likelihood ratio condition.¹⁰ Intuitively, the imposed ordering in H requires that at any point of search, continuing search will yield higher outcomes—in the sense of first-order stochastic dominance—in state θ'' than in state $\theta' < \theta''$. At a technical level, this guarantees that the Spence-Mirrlees property holds in a stochastic sense.¹¹ The second part of Assumption A4 strengthens the ordering, such that H is convexly

¹⁰Another condition which is less strict than the monotone likelihood ratio condition, and which also implies H decreasing, is sometimes referred to as monotone hazard ratio condition.

¹¹More precisely, Assumption A4 implies that the agent's indifference curves between *expected* transfers and different stopping rules cross only once over different states. To see this, let $T_\theta^e \equiv \mathbb{E}\{T(x) | x \geq \bar{x}, \theta\}$ denote the expected transfers to the agent in state θ with a given contract T , and let $u_\theta(T_\theta^e, \bar{x}) \equiv T_\theta^e - c/\bar{F}(\bar{x}|\theta)$ denote the expected utility of the agent when pursuing stopping rule \bar{x} . Then the single crossing property holds, if for any $(\theta, \theta') \in \{\Theta^2 | \theta > \theta'\}$ it holds that

increasing in θ . Intuitively, this requires that the benefit of being in a better state than θ is decreasing in θ . Assumption A5 is of more technical nature, ensuring that the objective function of the principal is concave.¹²

To keep the results clear, I impose two further assumptions that ensure the existence of an interior optimum.

Assumption A6. $\frac{d}{d\theta} \left(\frac{p(\theta)}{1-P(\theta)} \right) \geq 0$.

Assumption A7. $\frac{\partial^2}{\partial \bar{x} \partial \theta} \mathbb{E}\{x|x \geq \bar{x}, \theta\} \geq 0$.

Assumption A6 is standard in many mechanism design applications, meaning that the likelihood to be in a better state than θ is decreasing in θ . A sufficient condition for this to hold is that the likelihood $p(\theta)$ is weakly decreasing in θ (e.g., θ being uniform). Finally, Assumption A7 requires the marginal benefit of search to be increasing in θ . Together with Assumptions A4 and A6 this ensures the optimum to be interior.¹³

Parametric distributions for F consistent with these assumptions exist, for instance, within the Beta family and the (generalized) family of Pareto distributions (see the end of the next subsection for a particular simple example).¹⁴

I am now ready to show that the principal optimally designs a menu of contracts which is exclusively comprised of bonus contracts as defined below.

Definition. Let τ be a nonrandom constant. Then a contract T is called a *bonus contract* when it is of the following form:

$$T(x) = \begin{cases} 0 & \text{if } x < \bar{x} \\ \tau & \text{if } x \geq \bar{x}. \end{cases}$$

The argument establishes that a menu of bonus contracts implements any set of equilibrium search policies in a weakly optimal way and, therefore, any equilibrium

$-\frac{\partial u_\theta / \partial \bar{x}}{\partial u_\theta / \partial T_\theta^e} \leq -\frac{\partial u_{\theta'} / \partial \bar{x}}{\partial u_{\theta'} / \partial T_{\theta'}^e}$, which simplifies to $H(x|\theta) \leq H(x|\theta')$. In conjunction with Assumption A3 this then ensures that indifference curves are indeed single-crossing for any differential dT , since for any monotonic contract it holds that $dT_\theta^e \geq dT_{\theta'}^e$.

¹²This assumption can be relaxed to $HH_{x\theta} - H_x H_\theta \leq 0$, with subscripts denoting partial derivatives.

¹³In many cases Assumption A6 is not necessary. Specifically, when $p(\theta)/(1-P(\theta))$ is increasing at a sufficiently high rate, or when H is sufficiently convex ($HH_{\theta\theta} - H_\theta^2 \geq 0$, with subscripts denoting partial derivatives), it can be simply dropped. For details see the proof of Proposition 6.

¹⁴Given certain regularity conditions that ensure that a first-best solution exists, the analysis in this paper also seamlessly extends to the case where $B \rightarrow \infty$, permitting distributions for F with half-bounded supports (e.g., the exponential distribution with $\bar{F}(x|\theta) = e^{-x/\theta}$ for $\theta > 0$).

can be implemented by a menu of such contracts. First, I investigate the implications of a menu that consists exclusively of bonus contracts. The following lemma asserts that in this case the optimal menu resembles many of the characteristics of a standard adverse selection problem. In particular, I have that (IR_θ) , $(IC_{\theta,\bar{\theta}})$, $(SP_{\theta,\bar{\theta}}^-)$ and $(SP_{\theta,\bar{\theta}}^+)$ in the principal's maximization problem can be replaced by (a), (b) and (c) below.

Lemma 2. *Suppose Assumption A1–A4 hold. Let $\{T_\theta\}_{\theta \in \Theta}$ be a menu of bonus contracts, let $\bar{x}(\theta)$ be the search targets implemented by contract T_θ , and let $\tau(\theta)$ be the transfer paid to the agent when $x > \bar{x}(\theta)$. Then if $\{T_\theta\}_{\theta \in \Theta}$ is a solution to the principal's optimization program, then it is also a solution to a program where the principal maximizes her payoff subject to the following constraints:*

(a) $\bar{x}(\theta)$ is nondecreasing in θ ,

(b) $U(\theta) = 0$, and

(c) $\frac{dU}{d\theta} = -\frac{\partial}{\partial \theta} \left(\frac{c}{F(\bar{x}(\theta)|\theta)} \right)$,

where $U(\theta) \equiv u(\bar{x}(\theta), \tau(\theta), \theta)$ denotes the agent's indirect utility in state θ (given contract $T_\theta = (\bar{x}(\theta), \tau(\theta))$).

Equipped with Lemma 2, I show the following result.

Proposition 5. *Suppose Assumptions A1–A7 hold. Then to any equilibrium in the game defined in Section 2, there corresponds an associated menu of bonus contracts $\{T_\theta\}_{\theta \in \Theta}$, which implements the same actions and is (weakly) preferred by the principal.*

The proof establishes a weak optimality of bonus contracts by showing that for any set of search policies a menu of bonus contracts implementing these policies defines a lower bound on the utility of the agent. To develop an intuition, suppose the principal wants to implement the stopping rule $\bar{x}(\theta)$ in state θ . In order to incentivize the agent to continue search for all $x < \bar{x}(\theta)$, he needs to provide her with a certain expected benefit of finding $x \geq \bar{x}(\theta)$. Let $\hat{\tau}(\theta)$ denote the expected payment necessary to implement this benefit.¹⁵ Then it must hold that $\int_{\bar{x}(\theta)}^B T_\theta(x') dF(x'|\theta) \geq \hat{\tau}(\theta)$. The precise shape of T_θ on $[\bar{x}(\theta), B]$ is, however, irrelevant for the purpose of incentivizing the agent in state θ and can be freely used to reduce the agent's temptation of misreporting

¹⁵Generally $\hat{\tau}(\theta)$ depends on $\sup\{T_\theta(x) : x \in [0, \bar{x}(\theta)]\}$ in conjunction with $(SP_{\theta,\bar{\theta}}^\pm)$. For small values of the former term it is pinned down by (IR_θ) instead.

the state θ . As is typical for adverse selection problems, the relevant temptation in this context is to underreport the state, giving the agent in all states better than θ a stochastic advantage in finding high outcomes relative to state θ . Because of this stochastic advantage, any schedule T_θ that is strictly increasing on $[\bar{x}(\theta), B]$ yields an expected return that is strictly higher than $\hat{\tau}(\theta)$ in all states better than θ . By paying a fixed remuneration, bonus contracts eliminate this premium associated with misreporting.¹⁶

Regarding the question which search policy should be implemented off-equilibrium, it immediately follows:

Corollary. *Let $\{T_\theta\}_{\theta \in \Theta}$ be a solution to the principal's optimization program. Then without loss of generality one may assume that $\bar{x}(T_\theta, \tilde{\theta}) = \bar{x}(T_\theta, \theta) \equiv \bar{x}(\theta)$ for all $(\theta, \tilde{\theta}) \in \Theta^2$.*

Intuitively, implementing any off-equilibrium search policy that stops later than the one designed for state θ would necessarily generate a premium from misreporting the state. Any off-equilibrium search policy that stops earlier than in state θ does not need to be deterred.¹⁷ Accordingly, I henceforth suppress the second argument of $\bar{x}(T_\theta, \theta')$ and use $\bar{x}(\theta)$ to denote the search policy implemented by contract T_θ . With this (and Lemma 2), we can then write $\{T_\theta\}_{\theta \in \Theta}$ as a function of $\{\bar{x}(\theta)\}_{\theta \in \Theta}$:

Corollary. *The second-best optimum can be achieved by a menu of bonus contracts of the following form. For all $\theta \in \Theta$,*

$$T_\theta(x) = \begin{cases} 0 & \text{if } x < \bar{x}(\theta) \\ [\bar{F}(\bar{x}(\theta)|\theta)]^{-1}c + U(\theta) & \text{if } x \geq \bar{x}(\theta), \end{cases}$$

where

$$U(\theta) = \int_{\theta}^{\theta} -\frac{\partial}{\partial \tilde{\theta}} \left(\frac{c}{\bar{F}(\bar{x}(\tilde{\theta})|\tilde{\theta})} \right) d\tilde{\theta}$$

is the minimal rent necessary to deter misreporting of θ .

¹⁶There still remains a second type of rent associated with underreporting the state that accrues from having smaller expected costs of pursuing a particular search policy. Because of this second rent, delegated search is still inefficient in equilibrium.

¹⁷While it is possible that $\bar{x}(T_\theta, \tilde{\theta}) = 0$ for some $\tilde{\theta} < \theta$, it is never necessary in these cases to deter the agent from choosing T_θ . One may thus assume without loss of generality that $\bar{x}(T_\theta, \tilde{\theta}) = \bar{x}(T_\theta, \theta)$.

4.3 Optimal search policies

I now solve for the second-best search policies. Substituting for T_θ using the last corollary and again using Lemma 2, the challenge of incentivizing search is effectively solved and the maximization problem of the principal reduces to a standard adverse selection problem. The following proposition states the solution.

Proposition 6. *Suppose Assumptions A1–A7 hold. Then for some (nonempty) $\Phi \subseteq \Theta$, search is “directed”, with $\{\bar{x}(\theta)\}_{\theta \in \Phi}$ satisfying:*

$$c + D_\theta(\bar{x}(\theta)) = \int_{\bar{x}(\theta)}^B (x' - \bar{x}(\theta)) dF(x'|\theta), \quad (5)$$

for some $\bar{x}(\theta) > 0$, and where $D_\theta : X \rightarrow \mathbb{R}_+$ is defined by

$$D_\theta(x) = \begin{cases} 0 & \text{if } x = 0 \\ -\frac{1 - P(\theta)}{p(\theta)} \frac{\partial H(x|\theta)}{\partial \theta} \frac{c}{H(x|\theta)} & \text{if } x > 0. \end{cases} \quad (6)$$

For all $\theta \notin \Phi$, search is “undirected”, satisfying $\bar{x}(\theta) = 0$.

Comparing equation (5) to its first-best counterpart (1), the marginal cost of delegated search is inflated by an agency term D_θ . D_θ reflects the cost of learning the optimal search policy: Under delegation, increasing $\bar{x}(\theta)$ not only increases the expected search costs, but also makes it more tempting for the agent to misreport the search policy in states $\theta' \in \{\theta' \in \Theta : \theta' > \theta\}$. To offset for this additional temptation, the principal needs to pay the agent an (higher) rent $U(\theta')$ in all states θ' , making it (in expectations) more expensive to search in state θ .

Since the benefits of search are the same under delegation (the right-hand sides of (5) and (1)), we have:

Corollary. *Suppose Assumptions A1–A7 hold. Then delegated search is almost surely inefficient: $\bar{x}^{SB}(\theta) < \bar{x}^{FB}(\theta)$ for all $\theta \in [0, B)$.*

Proposition 6 distinguishes two cases. First, for all $\theta \in \Phi$, second-best search policies continue to be directed towards some target $\bar{x}(\theta)$, but the target is generally set too low (search stops too early). Second, for $\theta \notin \Phi$, sequential (delegated) search is not profitable at all and the principal simply asks the agent to sample a single outcome and to unconditionally select it as final.

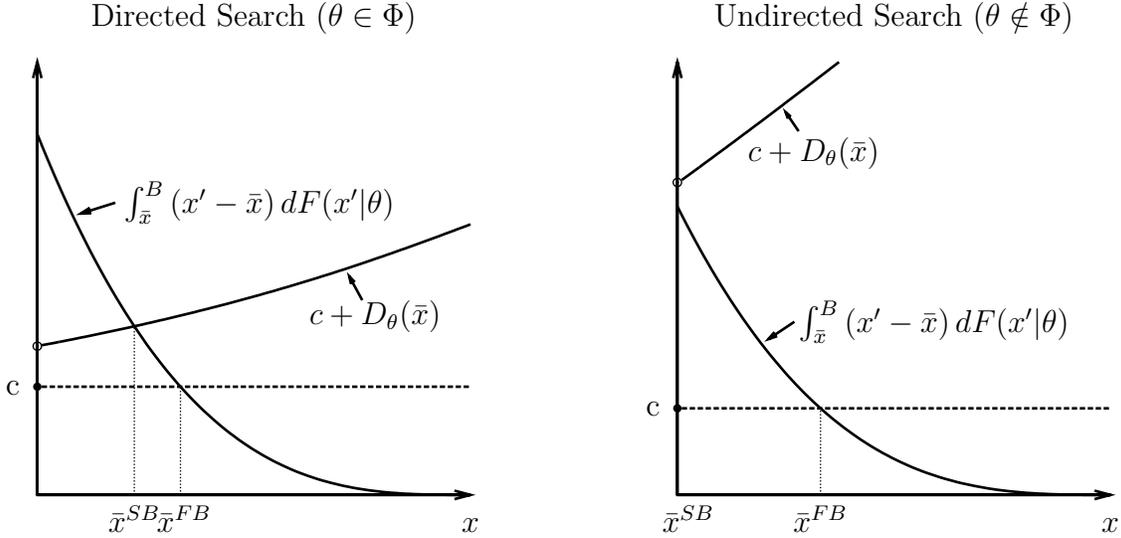


Figure 1: Second-best search.

Figure 1 illustrates the two cases. Delegated search is undirected ($\theta \notin \Phi$) if D_θ increases the left-hand side of (5) beyond the right-hand side for all $\bar{x} > 0$, either because it is unprofitable to search from an *ex post* perspective (taking into account the agency rents paid in θ), or because it is sufficiently unlikely to be in state θ , such that it is not worth the increase in rents in more likely states from an *ex ante* perspective. Because the agent is equally good at securing $x > 0$ in all states, there are no rents from deviating to $\bar{x} = 0$ that have to be compensated. Together with monotonicity of \bar{x} this implies $D_\theta(0) = 0$ for all θ . Undirected search is therefore preferred over no search whenever implementing $\bar{x} > 0$ is too costly.

A sufficient condition for $\theta \in \Phi$ is $c + D_\theta(\bar{x}) < \mathbb{E}(x|\theta)$ for a marginal \bar{x} :

$$c + \lim_{x \searrow 0} D_\theta(x) < \int x dF(x|\theta). \quad (7)$$

Using that agency rents $U(\theta')$ are increasing in $x(\theta)$ for all $\theta' > \theta$, the condition can be shown to be also necessary.

Proposition 7. *Suppose Assumptions A1–A7 hold. Then $\theta \in \Phi$ if and only if θ fulfills condition (7).*

In particular, since $\bar{x}(\theta)$ is increasing, it follows that Φ has the following “monotonicity” property.

Corollary. *Let $\theta'' > \theta'$. Then it holds that (i) if $\theta' \in \Phi$, then $\theta'' \in \Phi$; and (ii) if*

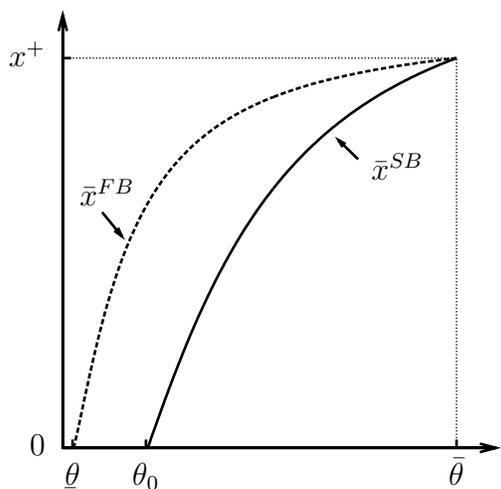


Figure 2: Optimal search policies.

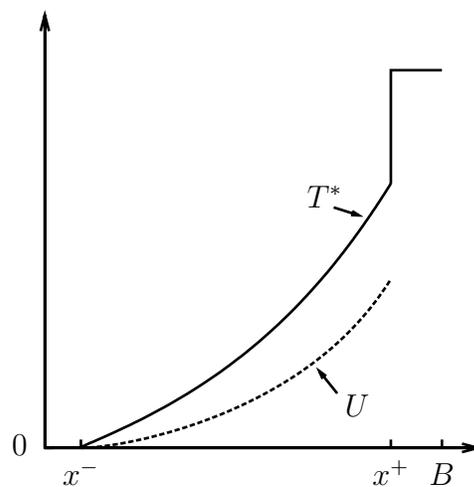


Figure 3: Indirect tariff.

$\theta'' \notin \Phi$, then $\theta' \notin \Phi$.

For an example, consider the case where θ is uniform on $[\frac{1}{10}, 4]$ and $\bar{F}(x|\theta) = (1-x)^{1/\theta}$.¹⁸ Figure 2 displays the optimal search policies as a function of θ . The example is chosen so that $\underline{\theta}$ is the lowest value of θ for which search is (first-best) profitable ($\bar{x}^{FB}(\underline{\theta}) \approx 0$). In state $\bar{\theta}$, delegated search is efficient (indicated by $x^+ \approx 0.81$). For all $\theta < \bar{\theta}$, second-best search stops too early compared to the first best, and is undirected on $[\underline{\theta}, \theta_0]$.

4.4 Indirect tariffs

The second-best bonus scheme is arguably a particular simple scheme among the class of direct mechanisms. An interesting question is, however, whether there also exists a simple indirect mechanism that implements the second best. In particular, does there exist a nonlinear (possibly discontinuous) tariff $T^* : X \rightarrow \mathbb{R}$ that implements the second best?

The answer is no. In contrast to pure adverse selection problems, any tariff T^* that implements the optimal search policies $\{\bar{x}(\theta)\}$ is strictly more costly than the direct bonus scheme.

¹⁸While this choice for F fails to satisfy the strict version of Assumptions A5 and A7, it is consistent with the relaxed versions given in Footnotes 12 and 13, so that all results apply.

Proposition 8. *Suppose Assumptions A1–A7 hold. Let $\{\bar{x}(\theta)\}$ define the second best search policies and let T^* define the tariff that implements $\{\bar{x}(\theta)\}$ at lowest expected costs. Then T^* exists, and expected transfers from the principal to the agent are strictly higher than $\tau(\theta)$ for all $\theta > \inf \Phi$.*

Underlying this inefficiency result is that the implementability conditions $(SP_{\theta, \bar{\theta}}^-)$ and $(SP_{\theta, \bar{\theta}}^+)$ deplete most degrees of freedom in designing T^* . Specifically, T^* must satisfy the following integral equation:

$$\bar{F}(x|q(x))T^*(x) + c = \int_x^B T^*(y) dF(y|q(x)) \quad \text{for all } x \in (x^-, x^+), \quad (8)$$

where $x^- \equiv \bar{x}(\inf \Phi)$, $x^+ \equiv \bar{x}(\bar{\theta})$, and $q(x) \equiv \bar{x}^{-1}(x)$ (for details, see the proof in the appendix). This leaves only the shape of T^* on $[x^+, B]$ as a means to replicate the second-best compensation scheme. It turns out that this degrees of freedom do not suffice to fulfill $(IC_{\theta, \bar{\theta}})$ and (IR_θ) at the (expected) second-best costs.

To build an intuition, note that $T(\bar{x}(\theta))$ defines the indirect utility of the agent in state θ , since for any $x = \bar{x}(\theta)$ she must be indifferent whether or not to continue searching. The key insight is that any solution to (8) is necessarily steeper than $U(q(x))$ (see Figure 3 for an illustration). Hence, satisfying individual rationality in state $\inf \Phi$ necessarily increases the rents in all other states beyond their second-best level. The reason is closely related to the optimality of bonus contracts. Because T^* must be strictly increasing on $[x^-, x^+]$ in order to implement the (continuous) mapping $\theta \mapsto \bar{x}$, the agent can generate the deviation premium that bonus contracts had eliminated by choosing $\bar{x}(\theta')$ in state $\theta > \theta'$. In order to nevertheless implement $\{\bar{x}(\theta)\}$, benefits of continued search have to compensate this premium, reflected in the steeper slope of T^* (defining the agent's utility under T^*) compared to U .

5 Summary

I have studied a model of delegated search under varying assumptions about what can be observed by the principal. If the principal can observe either the search process or shares the same information as the agent regarding its prospects, then delegated search is demonstrated to be efficient. If, however, the relation between the principal and the agent is governed by both imperfect monitoring of search and *ex ante* uncertainty

about its prospects, these sources of uncertainty exacerbate each other, and search is found to be stopped too early. In the light of this inefficiency, utilizing a menu of bonus contracts is shown to be second-best. The scheme strictly dominates any nonlinear (indirect) tariff.

Beyond the specific context of search agencies, the analysis may also illuminate the delegation of certain non-routine problems. Solving non-routine problems often requires investigating potential solutions that in the process may turn out unsatisfactory and require further thinking until a sufficiently promising solution strategy is conceived. Delegating such tasks resembles many aspects of the environment considered in this paper. For instance, managers are expected to come up with good business plans, consultants are hired to search for solutions to pending problems, and advocates need to find good strategies to defend their clients.

The optimal utilization of bonus contracts provides theoretical support for their widespread usage. The necessity to adopt these remuneration schemes depends, however, on the simultaneous relevance of the two asymmetries. If learning the optimal search policies is the main concern, a simple piece rate scheme that compensates the agent for each sampled solution suffices. If implementing search is the main concern, then linear payment schemes that make compensation highly sensitive to the value of the adopted solution are found to be optimal.

Of course, the precise configuration of asymmetries is often hard to determine for a particular application. With regards to the adopted notion of moral hazard, I conjecture that searching for solutions in non-routine tasks is often intrinsically unobservable, in particular when the search is of cognitive nature. Adverse selection regarding the optimal search policy, in contrast, is expected to increase in relevance with the expertise of the agent. I therefore suspect the delegation of search to be particularly challenging when tasks are both non-routine and require specialized skills. For less specialized tasks the experience of tenured agents may serve as a substitute.

A Mathematical Appendix

A.1 Proof of Lemma 1

Let $V(x|\theta)$ denote the indirect utility of the agent in state θ after sampling x . Then:

$$V(x|\theta) = \max \left\{ T_{\bar{\theta}}(x), -c + \int V(x'|\theta) dF(x'|\theta) \right\}, \quad (9)$$

whereas outcome x is selected as final whenever the associated transfer $T_{\bar{\theta}}(x)$ exceeds the expected utility from continuing search. Since this expected utility is independent of x , it holds that the agent selects outcome x whenever $T_{\bar{\theta}}(x) > \bar{T}_{\bar{\theta}}(\theta)$, where $\bar{T}_{\bar{\theta}}(\theta) = -c + \int V(x'|\theta) dF(x'|\theta)$. Let $\psi_{\bar{\theta}} : X \times \mathbb{R} \rightarrow \{0, 1\}$ be an indicator, such that

$$\psi_{\bar{\theta}}(x, \bar{T}_{\bar{\theta}}(\theta)) = \begin{cases} 0 & \text{if } T_{\bar{\theta}}(x) \leq \bar{T}_{\bar{\theta}}(\theta), \text{ and} \\ 1 & \text{if } T_{\bar{\theta}}(x) > \bar{T}_{\bar{\theta}}(\theta). \end{cases} \quad (10)$$

Then, using (9), I can rewrite $\bar{T}_{\bar{\theta}}(\theta)$ as

$$\bar{T}_{\bar{\theta}}(\theta) = -c + \int \left((1 - \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta))) \bar{T}_{\bar{\theta}}(\theta) + \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta)) T_{\bar{\theta}}(x') \right) dF(x'|\theta), \quad (11)$$

or

$$\begin{aligned} \bar{T}_{\bar{\theta}}(\theta) \left(1 - \int (1 - \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta))) dF(x'|\theta) \right) \\ = -c + \int \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta)) T_{\bar{\theta}}(x') dF(x'|\theta), \end{aligned} \quad (12)$$

or

$$c = \int \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta)) T_{\bar{\theta}}(x') dF(x'|\theta) - \bar{T}_{\bar{\theta}}(\theta) \int \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta)) dF(x'|\theta) \quad (13)$$

$$= \int \psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta)) \left(T_{\bar{\theta}}(x') - \bar{T}_{\bar{\theta}}(\theta) \right) dF(x'|\theta). \quad (14)$$

Because any increase in $\bar{T}_{\bar{\theta}}(\theta)$ weakly decreases $\psi_{\bar{\theta}}(x', \bar{T}_{\bar{\theta}}(\theta))$, the RHS of (14) is strictly decreasing in $\bar{T}_{\bar{\theta}}(\theta)$. Hence, if there exists a solution to (14), it is unique. Moreover, $\max_x T_{\bar{\theta}}(x)$ exists since X is compact, whereas the RHS of (14) is zero at $\max_x T_{\bar{\theta}}(x)$,

implying that (14) uniquely characterizes $\bar{T}_{\tilde{\theta}}(\theta)$ whenever $\int T(x) dF(x|\theta) \geq c$. Otherwise, marginal costs of searching always exceed the marginal benefits, and the agent trivially abstains from search.

A.2 Proof of Proposition 4

By construction of $\bar{x}(T_{\tilde{\theta}}, \theta)$, for some $x \in X$, $T_{\tilde{\theta}}(x) \geq \bar{T}_{\tilde{\theta}}(\theta)$ (see the proof to Lemma 1). Further, whenever search is directed, for some $x \in X$, $T_{\tilde{\theta}}(x) \geq \bar{T}_{\tilde{\theta}}(\theta)$. Thus for $T_{\tilde{\theta}}$ strictly increasing and continuous, $\bar{x}(T_{\tilde{\theta}}, \theta) \equiv T_{\tilde{\theta}}^{-1}(\bar{T}_{\tilde{\theta}}(\theta)) \in X$ obviously exists, and is given by (4a) and (4b). To verify the remaining cases, suppose that $\bar{T}_{\tilde{\theta}}(\theta)$ is not attained by $T_{\tilde{\theta}}(x)$ on X . Then from (4a) and (4b), $\bar{x}(T_{\tilde{\theta}}, \theta)$ is assigned to the point of discontinuity where $\lim_{x \nearrow \bar{x}(T_{\tilde{\theta}}, \theta)} T_{\tilde{\theta}}(x) < \bar{T}_{\tilde{\theta}}(\theta)$ and $\lim_{x \searrow \bar{x}(T_{\tilde{\theta}}, \theta)} T_{\tilde{\theta}}(x) > \bar{T}_{\tilde{\theta}}(\theta)$. So search given by $\bar{x}(T_{\tilde{\theta}}, \theta)$ is identical to search given by $\bar{T}_{\tilde{\theta}}(\theta)$. Finally, suppose that $\bar{T}_{\tilde{\theta}}(\theta)$ is attained on an interval $[\underline{x}, \bar{x}]$. Then from Lemma 1, the agent continues search for all $x \leq \bar{x}$ and stops search for $x > \bar{x}$. Thus $\bar{x}(T_{\tilde{\theta}}, \theta) = \bar{x}$, identical to the rule given by (4a) and (4b).

A.3 Proof of Lemma 2

To proof this lemma, I first characterize the set of stopping rules $\bar{x}(\theta)$ that are implementable via bonus contracts. From $(SP_{\theta, \tilde{\theta}}^-)$ and $(SP_{\theta, \tilde{\theta}}^+)$ it follows that an agent with bonus contract $T_{\tilde{\theta}} = (\bar{x}(\tilde{\theta}), \tau(\tilde{\theta}))$ chooses a stopping rule

$$\bar{x}(T_{\tilde{\theta}}, \theta) = \begin{cases} 0 & \text{if } \tau(\tilde{\theta}) < [\bar{F}(\bar{x}(\tilde{\theta})|\theta)]^{-1} c, \text{ and} \\ \bar{x}(\tilde{\theta}) & \text{if } \tau(\tilde{\theta}) \geq [\bar{F}(\bar{x}(\tilde{\theta})|\theta)]^{-1} c. \end{cases} \quad (15)$$

Accordingly, let $u(\theta, \tilde{\theta}) \equiv \max\{0, -[\bar{F}(\bar{x}(\tilde{\theta})|\theta)]^{-1} c + \tau(\tilde{\theta})\}$ denote the agent's indirect utility in state θ when she chooses bonus contract $T_{\tilde{\theta}} = (\bar{x}(\tilde{\theta}), \tau(\tilde{\theta}))$. Note that (IR_{θ}) implies that $\tau(\tilde{\theta}) \geq [\bar{F}(\bar{x}(\tilde{\theta})|\tilde{\theta})]^{-1} c$, and thus I can ignore $(SP_{\theta, \tilde{\theta}}^-)$ and $(SP_{\theta, \tilde{\theta}}^+)$ if $(IC_{\theta, \tilde{\theta}})$ holds.

Moreover, Assumption A4 implies that $u(\theta, \tilde{\theta}) \geq u(\tilde{\theta}, \tilde{\theta})$ for all $\theta \geq \tilde{\theta}$. Thus, $u(\theta, \tilde{\theta}) = -[\bar{F}(\bar{x}(\tilde{\theta})|\theta)]^{-1} c + \tau(\tilde{\theta})$ for all $\theta \geq \tilde{\theta}$. Therefore, for the agent in state θ to

not locally deviate, it must be that the first order condition

$$-H(\bar{x}(\tilde{\theta})|\theta) \frac{d\bar{x}(\tilde{\theta})}{d\tilde{\theta}} c + \frac{d\tau(\tilde{\theta})}{d\tilde{\theta}} = 0 \quad \text{for } \tilde{\theta} = \theta, \quad (16)$$

and the second order condition

$$\frac{d}{d\tilde{\theta}} \left(-H(\bar{x}(\tilde{\theta})|\theta) \frac{d\bar{x}(\tilde{\theta})}{d\tilde{\theta}} c + \frac{d\tau(\tilde{\theta})}{d\tilde{\theta}} \right) \leq 0 \quad \text{for } \tilde{\theta} = \theta \quad (17)$$

hold locally at $\tilde{\theta} = \theta$. Moreover, since (16) must hold for all $\theta \in \Theta$, it is an identity in θ , and thus

$$\frac{d}{d\tilde{\theta}} \left(-H(\bar{x}(\tilde{\theta})|\theta) \frac{d\bar{x}(\tilde{\theta})}{d\tilde{\theta}} c + \frac{d\tau(\tilde{\theta})}{d\tilde{\theta}} \right) - \frac{d}{d\theta} \left(H(\bar{x}(\tilde{\theta})|\theta) \frac{d\bar{x}(\tilde{\theta})}{d\tilde{\theta}} c \right) = 0 \quad \text{for } \tilde{\theta} = \theta. \quad (18)$$

Substituting (17) in (18) together with Assumption A4 thus yields $d\bar{x}/d\theta \geq 0$. This establishes that any solution to the principal's optimization program implies condition (a).

Next I argue that (16) is also sufficient to prevent the agent from deviating globally. First note that in the case where $\bar{x}(T_{\tilde{\theta}}, \theta) = 0$, $u(\theta, \tilde{\theta}) = 0 \leq u(\theta, \theta)$ and thus I can restrict attention to the case where the agent chooses $\bar{x}(T_{\tilde{\theta}}, \theta) = \bar{x}(\tilde{\theta})$ under contract $T_{\tilde{\theta}}$. Suppose to the contrary that the incentive constraint is violated in at least one state, i.e. $u(\theta, \tilde{\theta}) - u(\theta, \theta) > 0$ for some $(\theta, \tilde{\theta}) \in \Theta^2$, or by the fundamental theorem of calculus,

$$\int_{\theta}^{\tilde{\theta}} \left(-H(\bar{x}(\theta')|\theta) \frac{d\bar{x}(\theta')}{d\theta'} c + \frac{d\tau(\theta')}{d\theta'} \right) d\theta' > 0. \quad (19)$$

Suppose $\tilde{\theta} > \theta$. Then, Assumption A4 implies that $H(\bar{x}(\tilde{\theta})|\tilde{\theta}) \leq H(\bar{x}(\tilde{\theta})|\theta)$, and therefore (19) implies

$$\int_{\theta}^{\tilde{\theta}} \left(-H(\bar{x}(\theta')|\theta') \frac{d\bar{x}(\theta')}{d\theta'} c + \frac{d\tau(\theta')}{d\theta'} \right) d\theta' > 0 \quad (20)$$

since $d\bar{x}/d\theta \geq 0$. However, equation (16) implies that the integrand in (20) is equal to 0 for all θ' , contradicting that for any $\tilde{\theta} > \theta$, contract $T_{\tilde{\theta}}$ is preferred over T_{θ} . The

same logic establishes a contradiction for the case where $\tilde{\theta} < \theta$.

Now, let $U(\theta) \equiv u(\theta, \theta)$. Then

$$\frac{dU}{d\theta} = -H(\bar{x}(\theta)|\theta) \frac{d\bar{x}(\theta)}{d\theta} c + \frac{d\tau(\theta)}{d\theta} - \frac{\partial}{\partial \theta} \frac{c}{\bar{F}(\bar{x}(\theta)|\theta)}. \quad (21)$$

Thus (c) holds if and only if (16) holds. Therefore, (a) and (c) are both sufficient and necessary for $(IC_{\theta, \tilde{\theta}})$ to hold. Moreover, as shown above, $(IC_{\theta, \tilde{\theta}})$ implies $(SP_{\theta, \tilde{\theta}}^-)$ and $(SP_{\theta, \tilde{\theta}}^+)$. Thus I am left to show that (a), (b) and (c) are sufficient and (b) is necessary for (IR_{θ}) to hold. Consider sufficiency first. From (c) I have that $dU/d\theta$ is increasing in θ (by Assumption A4 implies $\partial \bar{F}(\bar{x}(\theta)|\theta) \geq 0$). Thus $U(\underline{\theta}) = 0$ implies (IR_{θ}) for all $\theta \in \Theta$. That $U(\underline{\theta}) = 0$ is also necessary for a solution to the principal's optimization program follows trivially from the above analysis, since shifting $U(\underline{\theta})$ doesn't affect any other constraints.

A.4 Proof of Proposition 5

Consider an arbitrary menu of contracts $\{T_{\theta}\}_{\theta \in \Theta}$, and let $u(T, \bar{x}, \theta)$ be the utility of the agent in state θ when she chooses contract T and search policy \bar{x} . Then $u(T_{\tilde{\theta}}, \bar{x}(T_{\tilde{\theta}}, \theta), \theta) \geq u(T_{\tilde{\theta}}, \bar{x}(T_{\tilde{\theta}}, \tilde{\theta}), \theta)$, and therefore a necessary condition for $(IC_{\theta, \tilde{\theta}})$ to hold is that

$$U(\theta) = u(T_{\theta}, \bar{x}(T_{\theta}, \theta), \theta) \geq u(T_{\tilde{\theta}}, \bar{x}(T_{\tilde{\theta}}, \tilde{\theta}), \theta) \quad \text{for all } \tilde{\theta} \in \Theta.$$

Hence $\tilde{\theta} = \theta$ maximizes the RHS of the inequality, with $U(\theta)$ also being the value function of $\max_{\tilde{\theta}} u(T_{\tilde{\theta}}, \bar{x}(T_{\tilde{\theta}}, \tilde{\theta}), \theta)$. The envelope theorem implies

$$\frac{dU}{d\theta} = \left\{ \frac{\partial}{\partial \theta} \left(\frac{\int_{\bar{x}(T_{\tilde{\theta}}, \tilde{\theta})}^B T(x') dF(x'|\theta)}{\bar{F}(\bar{x}(T_{\tilde{\theta}}, \tilde{\theta})|\theta)} \right) - \frac{\partial}{\partial \theta} \left(\frac{c}{\bar{F}(\bar{x}(T_{\tilde{\theta}}, \tilde{\theta})|\theta)} \right) \right\} \Bigg|_{\tilde{\theta}=\theta}, \quad (22)$$

and since by Assumptions A3 and A4 the first term in (22) is positive, I have that

$$\frac{dU}{d\theta} \geq -\frac{\partial}{\partial \theta} \left(\frac{c}{\bar{F}(\bar{x}(T_{\theta}, \theta)|\theta)} \right). \quad (23)$$

Moreover, (IR_{θ}) implies $U(\underline{\theta}) \geq 0$. Thus for any \bar{x} that is nondecreasing in θ , the agent's utility under a menu of bonus contracts as given by Lemma 2(b) and (c)

constitutes a lower bound on the agent's utility under any menu of contracts that implements \bar{x} . As long as the restriction on \bar{x} is not binding for bonus contracts, a menu of bonus contracts is therefore optimal. That the restriction that \bar{x} is increasing is indeed not binding is established in the proof of Proposition 6.

A.5 Proof of Proposition 6

Using $\{T_\theta\}_{\theta \in \Theta}$ as given by the second corollary to Proposition 5, the principal's objective function is

$$\int_{\underline{\theta}}^{\bar{\theta}} \left(\int_{\bar{x}(\theta)}^B \frac{x'}{\bar{F}(\bar{x}(\theta)|\theta)} dF(x'|\theta) - \frac{c}{\bar{F}(\bar{x}(\theta)|\theta)} + \int_{\theta}^{\bar{\theta}} \frac{\partial}{\partial \tilde{\theta}} \left(\frac{c}{\bar{F}(\bar{x}(\tilde{\theta})|\tilde{\theta})} \right) d\tilde{\theta} \right) dP(\theta), \quad (24)$$

or, after an integration by parts,

$$\int_{\underline{\theta}}^{\bar{\theta}} \left(\int_{\bar{x}(\theta)}^B \frac{x'}{\bar{F}(\bar{x}(\theta)|\theta)} dF(x'|\theta) - \frac{c}{\bar{F}(\bar{x}(\theta)|\theta)} + \frac{1 - P(\theta)}{p(\theta)} \frac{\partial}{\partial \theta} \left(\frac{c}{\bar{F}(\bar{x}(\theta)|\theta)} \right) \right) dP(\theta). \quad (25)$$

By Lemma 2 the only relevant constraint which I have to take care of is that $\bar{x}(\theta)$ is nondecreasing in θ . Ignoring this constraint for the moment, the maximizer of (25) is given by

$$c + D_\theta(\bar{x}(\theta)) = \int_{\bar{x}(\theta)}^B (x' - \bar{x}(\theta)) dF(x'|\theta), \quad (26)$$

where function $D_\theta : X \rightarrow \mathbb{R}_+$ is defined by,

$$D_\theta(x) = -\frac{1 - P(\theta)}{p(\theta)} \frac{\partial H(x|\theta)}{\partial \theta} \frac{c}{H(x|\theta)}. \quad (27)$$

(From Assumption A5, (25) is concave at any $\bar{x}(\theta)$ that satisfies (26), so (25) is globally quasi-concave, and the second-order condition is satisfied.)

I still have to show two things. First, that $\bar{x}(\theta)$ as given by (26) is indeed nonde-

creasing. Second, I have to extend the analysis to the possibility that for some $\theta \in \Theta$, a corner solution might be optimal.

From Assumption A4, the RHS of (26) is increasing in θ , so $\bar{x}(\theta)$ will trivially be nondecreasing whenever $D_\theta(\bar{x}(\theta))$ is nonincreasing in θ . This will be the case whenever $p/(1-P)$ is increasing at a sufficiently high rate, and/or H is sufficiently convex in θ . Otherwise, $\bar{x}(\theta)$ will still be increasing if the RHS of (26) is increasing at a sufficiently high rate. A sufficient condition for this to be the case is that

$$H(\bar{x}(\hat{\theta})|\theta) \int_{\bar{x}(\hat{\theta})}^B (x' - \bar{x}(\hat{\theta})) dF(x'|\theta) \quad (28)$$

is increasing in θ at $\hat{\theta} = \theta$. Assumption A7 ensures that this is always the case. Hence, any interior solution to the principal's program is characterized by equations (26) and (27).

Because for $\bar{x}(\theta) = B$ benefits of search (the RHS of (26)) are equal to 0, corner solutions may at most be given by $\bar{x}(\theta) = 0$. They occur whenever the left-hand side of (26) exceeds the right-hand side for all values of $\bar{x}(\theta)$, indicating that the "solution" to (26) is negative. Given the constraint $\bar{x} \geq 0$ and quasi-concavity of the problem, we hence have $\bar{x} = 0$ as long as the principal wishes to still implement search. To see that this indeed the case, note that the last term in (25) drops out for $\bar{x} = 0$ since $\bar{F}(0|\theta) = 1$ for all θ . Hence, for $\bar{x} = 0$ the problem collapses to the first-best problem where per assumption $\mathbb{E}(x|\theta) \geq c$ for all $\theta \in \Theta$.

A.6 Proof of Proposition 7

From Proposition 6 marginal costs of searching are given by $c + D_\theta(\bar{x})$. Differentiating with respect to \bar{x} yields

$$-\frac{1-P}{P} \left(\frac{\partial^2 H}{\partial x \partial \theta} \frac{1}{H} - \frac{\partial H}{\partial \theta} \frac{\partial H}{\partial x} \frac{1}{H^2} \right) c \geq 0, \quad (29)$$

by Assumptions A4 and A5. Moreover, marginal benefits are trivially decreasing in \bar{x} . Thus a necessary and sufficient condition for $\bar{x}(\theta) > 0$ is that for $\hat{x} \searrow 0$ directed search is beneficial:

$$\lim_{\hat{x} \searrow 0} \left\{ \int_{\hat{x}}^B (x' - \hat{x}) dF(x'|\theta) - c - D_\theta(\hat{x}) \right\} > 0, \quad (30)$$

or

$$c + \lim_{x \searrow 0} D_\theta(x) < \int x \, dF(x|\theta). \quad (31)$$

A.7 Proof of Proposition 8

Existence Consider existence of T^* first. By continuity of \bar{x} , a necessary and sufficient condition for T^* to implement $\{\bar{x}(\theta)\}$ is that the equivalent to $(SP_{\theta,\theta}^\pm)$ holds with equality for all θ in $(\theta_0, \bar{\theta})$, where $\theta_0 \equiv \inf \Phi$, or equivalently:

$$c = \int_x^B (T(y) - T(x)) \, dF(y|q(x)), \quad (32)$$

or

$$T(x) = -\frac{c}{\bar{F}(x|q(x))} + \frac{1}{\bar{F}(x|q(x))} \int_x^B T(y) \, dF(y|q(x)), \quad (33)$$

for all $x \in (x^-, x^+)$, $x^- \equiv \bar{x}(\theta_0)$, $x^+ \equiv \bar{x}(\bar{\theta})$, and where $q(x) \equiv \bar{x}^{-1}(x)$.¹⁹ Separate T^* into T^- defined on $[0, x^+]$ and T^+ defined on $[x^+, B]$, and fix some T^+ . Then T^- is given by the functional \mathcal{T} ,

$$(\mathcal{T}f)(x) = g(x) + \frac{1}{\bar{F}(x|q(x))} \int_x^{x^+} f(y) \, dF(y|q(x)), \quad (34)$$

where

$$g(x) = \frac{1}{\bar{F}(x|q(x))} \left(-c + \int_{x^+}^B T^+(y) \, dF(y|q(x)) \right). \quad (35)$$

Inspecting \mathcal{T} , it is clearly increasing in f . Moreover, for any constant k ,

$$\mathcal{T}(f+k)(x) = (\mathcal{T}f)(x) + \frac{F(x^+|q(x)) - F(x|q(x))}{1 - F(x|q(x))} k.$$

By (1), $x^+ \in (0, 1)$, such that the term multiplying k is in $(0, 1)$. Hence, \mathcal{T} satisfies Blackwell's sufficient conditions to be a contraction, establishing existence of a unique

¹⁹Here I assume without loss of generality that $\partial H/\partial \theta < 0$, so that \bar{x} is invertible. If it were not, I could simply define a new variable that treats all instances of θ where H is constant as a single state and carry out the following analysis with respect to that variable.

T^- for a given T^+ .

Having taken care of implementing $\{\bar{x}(\theta)\}$, the only other constraints to address are individual rationality. Since \mathcal{T} is increasing in T^+ , individual rationality can be guaranteed by setting T^+ accordingly. We conclude that there exists a T^* implementing the second best search policies and leave it to the reader to formally establish the shape of T^+ that defines the cost-minimizing tariff. (The answer is: $T^+(x) = \text{const}$ for all $x \in [x^+, B]$, where const is set such that $T^-(x^-) = 0$.)

Inefficiency Suppose there exists a tariff T that implements $\{\bar{x}(\theta)\}$ at the same costs as in the second best in all states $\theta \in \Theta$. Since both parties are risk neutral, this implies that $U(\theta)$ corresponds to the second-best rents for all θ . Contradicting the existence of such a T , I first show that T necessarily violates individual rationality for $\theta \rightarrow \theta_0$. Subsequently I then argue that restoring individual rationality for θ_0 requires increasing rents for all $\theta \in \Phi$ above their second-best level.

Let $U(\theta)$ denote the second-best rents as given by the second corollary to Proposition 5, and suppose that T implements $U(\theta)$ for all θ . Then from (33), $T(x) = U(q(x))$. Hence, implementability requires

$$\begin{aligned} \bar{F}(x|q(x))\tilde{U}(q(x)) = \\ -c + \int_{x^+}^B T(y) dF(y|q(x)) + \int_x^{x^+} \int_{\theta}^{q(x)} m(\tilde{\theta}) d\tilde{\theta} dF(y|q(x)), \end{aligned} \quad (36)$$

for all $x \in (x^-, x^+)$, where $\tilde{U}(\theta)$ is the actual utility implemented by T , and

$$m(\theta) = -\frac{\partial}{\partial \theta} \left(\frac{c}{\bar{F}(\bar{x}(\theta)|\theta)} \right). \quad (37)$$

After an integration by parts, a change in variables, and a collecting of terms, (36) becomes

$$\begin{aligned} \bar{F}(x|q(x))\tilde{U}(q(x)) = \\ -c + \int_{x^+}^B T(y) dF(y|q(x)) - F(x|q(x)) \int_{\theta}^{q(x)} m(\tilde{\theta}) d\tilde{\theta} \\ + F(x^+|q(x)) \int_{\theta}^{\bar{\theta}} m(\tilde{\theta}) d\tilde{\theta} - \int_{q(x)}^{\bar{\theta}} m(\tilde{\theta}) F(\bar{x}(\tilde{\theta})|q(x)) d\tilde{\theta}, \end{aligned} \quad (38)$$

or, after another change in variables, and substituting again for $U(\theta)$,

$$\begin{aligned} \bar{F}(\bar{x}(\theta)|\theta)\tilde{U}(\theta) &= -c + \int_{x^+}^B T(y) dF(y|\theta) - F(\bar{x}(\theta)|\theta)U(\theta) \\ &\quad + F(x^+|\theta)U(\bar{\theta}) - \int_{\theta}^{\bar{\theta}} m(\tilde{\theta})F(\bar{x}(\tilde{\theta})|\theta) d\tilde{\theta}. \end{aligned} \quad (39)$$

Consider the first two terms on the right-hand side. By Assumption A4, $\tilde{U}(\theta)$ is maximized for all $\theta < \bar{\theta}$, subject to $U(\bar{\theta})$ being fixed, by setting T^+ constant. Hence,

$$\begin{aligned} \int_{x^+}^B T(y) dF(y|\theta) - c &\leq \frac{\bar{F}(x^+|\theta)}{\bar{F}(x^+|\bar{\theta})} \int_{x^+}^B T(y) dF(y|\bar{\theta}) - c = \\ &\quad \bar{F}(x^+|\theta) \left(U(\bar{\theta}) + \frac{c}{\bar{F}(x^+|\bar{\theta})} \right) - c. \end{aligned} \quad (40)$$

Substituting in (39) yields

$$\begin{aligned} \bar{F}(\bar{x}(\theta)|\theta)\tilde{U}(\theta) &\leq \\ &\quad - F(\bar{x}(\theta)|\theta)U(\theta) + U(\bar{\theta}) - c \left(1 - \frac{\bar{F}(x^+|\theta)}{\bar{F}(x^+|\bar{\theta})} \right) - \int_{\theta}^{\bar{\theta}} m(\tilde{\theta})F(\bar{x}(\tilde{\theta})|\theta) d\tilde{\theta} = \\ &\quad \bar{F}(\bar{x}(\theta)|\theta)U(\theta) - c \left(1 - \frac{\bar{F}(x^+|\theta)}{\bar{F}(x^+|\bar{\theta})} \right) + \int_{\theta}^{\bar{\theta}} m(\tilde{\theta})\bar{F}(\bar{x}(\tilde{\theta})|\theta) d\tilde{\theta}. \end{aligned} \quad (41)$$

By Assumption A4, $\partial H(x|\tilde{\theta})/\partial \tilde{\theta} = \partial^2[\bar{F}(x|\tilde{\theta})]^{-1}/\partial x \partial \tilde{\theta} \leq 0$. Moreover, $\bar{F}(x|\theta)$ is also (strictly) decreasing in x . Hence, the last term on the right-hand side satisfies

$$-c \int_{\theta}^{\bar{\theta}} \frac{\partial}{\partial \tilde{\theta}} \left(\frac{\bar{F}(\bar{x}(\tilde{\theta})|\theta)}{\bar{F}(\bar{x}(\tilde{\theta})|\tilde{\theta})} \right) d\tilde{\theta} < -c \int_{\theta}^{\bar{\theta}} \frac{\partial}{\partial \tilde{\theta}} \left(\frac{\bar{F}(x^+|\theta)}{\bar{F}(x^+|\tilde{\theta})} \right) d\tilde{\theta} = -c \left[\frac{\bar{F}(x^+|\theta)}{\bar{F}(x^+|\tilde{\theta})} \right]_{\tilde{\theta}=\theta}^{\bar{\theta}}. \quad (42)$$

Substituting in (41), the last two terms on the right-hand side cancel out. Hence,

$$\tilde{U}(\theta) < U(\theta) \quad \text{for all } \theta \in (\theta_0, \bar{\theta}), \quad (43)$$

contradicting $\tilde{U}(\theta) = U(\theta)$. In particular, it holds that for $\theta \rightarrow \theta_0$, $\tilde{U}(\theta_0) < 0$ since $U(\theta_0) = 0$. From the existence proof above, it is clear that the only degree of freedom to restore individual rationality under the constraint of implementability consists in

raising T^+ . Under Assumption A4 any change dT^+ implies $dU(\theta) > dU(\theta_0)$ for all $\theta > \theta_0$. Moreover, since \bar{x} is increasing in θ , the difference between the left-hand side to the right-hand side in (42) is decreasing in θ . Hence, any T^+ that sets $\tilde{U}(\theta_0) = 0$ necessarily sets $\tilde{U}(\theta) > U(\theta)$ for all $\theta > \theta_0$.

References

- Akerlof, George A.** 1970. "The market for 'lemons': Quality uncertainty and the market mechanism." *The Quarterly Journal of Economics*, 84(3): 488–500.
- Albrecht, James, Axel Anderson, and Susan Vroman.** 2010. "Search by committee." *Journal of Economic Theory*, 145(4): 1386–1407.
- Bolton, Patrick, and Mathias Dewatripont.** 2005. *Contract theory*. MIT press.
- Churchill, Gilbert A, Neil M Ford, and Orville C Walker.** 1993. *Sales force management*. Irwin/McGraw-Hill Chicago.
- Compte, Olivier, and Philippe Jehiel.** 2010. "Bargaining and majority rules: A collective search perspective." *Journal of Political Economy*, 118(2): 189–221.
- Demougin, D., and C. Fluet.** 1998. "Mechanism sufficient statistic in the risk-neutral agency problem." *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, 622–639.
- Herweg, F., D. Muller, and P. Weinschenk.** 2010. "Binary payment schemes: Moral hazard and loss aversion." *The American Economic Review*, 100(5): 2451–2477.
- Holmstrom, B.** 1979. "Moral hazard and observability." *The Bell Journal of Economics*, 74–91.
- Holmstrom, B., and P. Milgrom.** 1987. "Aggregation and linearity in the provision of intertemporal incentives." *Econometrica: Journal of the Econometric Society*, 55(2): 303–328.
- Innes, R.D.** 1990. "Limited liability and incentive contracting with ex-ante action choices." *Journal of Economic Theory*, 52(1): 45–67.

- Jewitt, I., O. Kadan, and J.M. Swinkels.** 2008. "Moral hazard with bounded payments." *Journal of Economic Theory*, 143(1): 59–82.
- Kim, S.K.** 1997. "Limited liability and bonus contracts." *Journal of Economics & Management Strategy*, 6(4): 899–913.
- Laffont, Jean-Jacques, and Jean Tirole.** 1986. "Using cost observation to regulate firms." *The Journal of Political Economy*, 614–641.
- Laffont, Jean-Jacques, and Jean Tirole.** 1993. *A theory of incentives in procurement and regulation*. MIT Press, Cambridge, MA.
- Levitt, Steven D., and Chad Syverson.** 2008. "Market Distortions when Agents Are Better Informed: The Value of Information in Real Estate Transactions." *The Review of Economics and Statistics*, 90(4): 599–611.
- Lewis, Tracy R.** 2012. "A theory of delegated search for the best alternative." *The RAND Journal of Economics*, 43(3): 391–416.
- Lewis, Tracy R., and Marco Ottaviani.** 2008. "Search Agency." *Mimeo*.
- McAfee, R Preston, and John McMillan.** 1987. "Competition for agency contracts." *The Rand Journal of Economics*, 296–307.
- McCall, John J.** 1970. "Economics of Information and Job Search." *The Quarterly Journal of Economics*, 84(1): 113–126.
- Mirrlees, James A.** 1971. "An exploration in the theory of optimum income taxation." *The review of economic studies*, 175–208.
- Mortensen, D.T.** 1970. "Job search, the duration of unemployment, and the Phillips curve." *The American Economic Review*, 60: 847–862.
- Mortensen, D.T.** 1986. "Job search and labor market analysis." *Handbook of labor economics*, 2: 849–919.
- Moynahan, J.K.** 1980. *Designing an effective sales compensation program*. New York:Amacom.
- Oyer, Paul.** 2000. "A theory of sales quotas with limited liability and rent sharing." *Journal of labor Economics*, 18(3): 405–426.

- Park, E.S.** 1995. “Incentive contracting under limited liability.” *Journal of Economics & Management Strategy*, 4(3): 477–490.
- Rogerson, Richard, Robert Shimer, and Randall Wright.** 2005. “Search-theoretic models of the labor market: a survey.” *Journal of Economic Literature*, 43(4): 959–988.
- Rutherford, Ronald C., Thomas M. Springer, and Abdullah Yavas.** 2005. “Conflicts Between Principals and Agents: Evidence from Residential Brokerage.” *Journal of Financial Economics*, 76(3): 627–665.
- Sannikov, Yuliy.** 2008. “A continuous-time version of the principal-agent problem.” *The Review of Economic Studies*, 75(3): 957–984.
- Townsend, R.** 1979. “Optimal contracts and competitive markets with costly state verification.” *Journal of Economic theory*, 21(2): 265–293.
- Toxvaerd, Flavio.** 2006. “Time of the Essence.” *Journal of Economic Theory*, 129(1): 252–272.