

Working Paper

2022/13/DSC/TOM

(Revised version of 2019/44/DSC/TOM)

The When and How of Delegated Search

Sasa Zorc University of Virginia, <u>zorcs@darden.virginia.edu</u>

> Ilia Tsetlin INSEAD, <u>ilia.tsetlin@insead.edu</u>

Sameer Hasija INSEAD, <u>Sameer.hasija@insead.edu</u>

Stephen E. Chick INSEAD, <u>stephen.chick@insead.edu</u>

March 16, 2022

Firms often outsource search processes such as the acquisition of real estate, new technologies, or talent. To ensure the efficacy of such delegated search, firms need to carefully design incentive contracts to attenuate the ill effects of agency issues. We model this problem using a dynamic principal-agent framework, embedding the standard sequential search model. The optimal contract pays the agent a fixed per-period fee, plus a bonus for finding a suitable alternative. The bonus size is defined a priori and decreases over time, while the range of values deemed suitable expands over time. If the principal is unable to contract on the value of the delivered alternatives, the optimal contract consists of two parts: early in the search process, the agent is granted a small bonus for every alternative brought to the principal, irrespective of whether the principal accepts it; late in the search process, the agent is awarded a comparatively larger bonus, which is decreasing in time, but only if the principal accepts the alternative. We also consider situations where the principal chooses between searching in-house and outsourcing. This decision is shown to hinge on the principal's tradeoff between speed and quality. The age-old aphorism "if you want it done right, do it yourself" holds, as in-house search is optimal for a principal who prioritizes quality. Yet, in the context of our model, we also establish an addendum: "if you want it done fast, hire someone else to do it."

Keywords: Search; Contract Theory; Dynamic Programming.

Electronic copy available at: <u>https://ssrn.com/abstract=3463626</u>

Working Paper is the author's intellectual property. It is intended as a means to promote research to interested readers. Its content should not be copied or hosted on any server without written permission from publications.fb@insead.edu

Find more INSEAD papers at https://www.insead.edu/faculty-research/research

Copyright © 2022 INSEAD

The When and How of Delegated Search

Saša Zorc

Darden School of Business, University of Virginia, 100 Darden Blvd, Charlottesville, VA 22903 zorcs@darden.virginia.edu

> Ilia Tsetlin, Sameer Hasija INSEAD, 1 Ayer Rajah Avenue, Singapore 138676, Singapore ilia.tsetlin@insead.edu, sameer.hasija@insead.edu

Stephen E. Chick INSEAD, Boulevard de Constance, Fontainebleau 77305, France stephen.chick@insead.edu

March 16, 2022

Firms often outsource search processes such as the acquisition of real estate, new technologies, or talent. To ensure the efficacy of such delegated search, firms need to carefully design incentive contracts to attenuate the ill effects of agency issues. We model this problem using a dynamic principal-agent framework, embedding the standard sequential search model. The optimal contract pays the agent a fixed per-period fee, plus a bonus for finding a suitable alternative. The bonus size is defined a priori and decreases over time, while the range of values deemed suitable expands over time. If the principal is unable to contract on the value of the delivered alternatives, the optimal contract consists of two parts: early in the search process, the agent is granted a small bonus for every alternative brought to the principal, irrespective of whether the principal accepts it; late in the search process, the agent is awarded a comparatively larger bonus, which is decreasing in time, but only if the principal accepts the alternative. We also consider situations where the principal chooses between searching in-house and outsourcing. This decision is shown to hinge on the principal's tradeoff between speed and quality. The age-old aphorism "if you want it done right, do it yourself" holds, as in-house search is optimal for a principal who prioritizes quality. Yet, in the context of our model, we also establish an addendum: "if you want it done fast, hire someone else to do it."

Key words: search; contract theory; dynamic programming

1. Introduction

Consider a search for a new house to buy. How many houses should be surveyed before a decision is made? How can one balance time pressure with the desire to find something better than what has already been observed? These questions arise not only in real estate but also in job search, the search for profitable investment opportunities, and the process of discovering and adopting new technologies. A ubiquitous practice in business is to hire someone else to conduct the search on your behalf. Looking for a new employee? A headhunter will find the person for you. Need to sell a house? Hire a real estate agent. Looking to acquire new technology? – technology scout.

Suppose, for example, that you are looking for a house. Would it be better to search for a property yourself or rather to hire a real estate agent for that task? This is a nuanced decision that depends on several factors. In the first place, there is a cost to hiring an agent, yet it will save you time by delegating activities you would otherwise have to do yourself. Moreover, there are other possible benefits to hiring an agent. The agent, who conducts searches for a living, can probably search more rapidly and efficiently than you can, might be aware of unlisted but available properties that give you more options, and may be more proficient at bargaining – giving you a better price than if you had negotiated yourself.

However, the agent may have other incentives besides your best interest. For example, a standard real estate contract pays the agent a percentage of the final purchase price for the house. In that case, the agent has no incentive to seek a lower price because doing so would reduce his fee. So if the agent is aware of two similar houses with different prices, he earns more by bringing only the higher-priced house to your attention (a classic adverse selection problem). Even in situations where the price is not an issue, such as when you offer the agent a flat fee, there are still agency frictions. In particular, the agent will then be incentivized to find you a property of minimum acceptable quality and as rapidly as possible – instead of patiently looking for a property you will like the most. Similarly, if the agent is paid for his time, the agent might not even search only to report that nothing suitable was found (a classic moral hazard problem).

Another factor to bear in mind is that once the search has been delegated to an agent who maximizes his own utility, you must rely on the information that this agent presents to you. Finally, in addition to the moral hazard and adverse selection issues that create frictions, there are many settings in which the agent (e.g., an independent expert or a niche consulting firm) is risk-averse, and so requires compensation for any uncertainty. All these agency issues make it difficult to determine whether you are best served by hiring an agent and, if so, to devise incentives that will yield the best possible outcome from your perspective.

Our paper aims to resolve this conundrum by answering two questions. (i) What conditions call for, respectively, hiring an agent and conducting the search oneself? (ii) When search is delegated, what is the principal's optimal contract with the agent?

Related Literature. The problem of a single decision maker undertaking a search process has been the topic of extensive research in search theory. For foundational work on this topic (including reviews), the reader is referred to McCall (1970), Gronau (1971), Lippman and McCall (1976), Weitzman (1979), Lancaster and Chesher (1983), and Mortensen (1986). Variants of the sequential search model have been used to describe a variety of situations, including job search (Mortensen 1986), adoption of new technologies (McCardle 1985), the real estate market (Baucells and Lippman 2004), and innovation (Rahmani and Ramachandran 2021). Some of the more recent contributions to this literature are Tang et al. (2009), Ulu and Smith (2009), Lippman and Mamer (2012), Massala and Tsetlin (2015), Mai and Pekeč (2018), Hu and Tang (2021), Alizamir et al. (2019), and Zorc and Tsetlin (2020). We draw from this stream of work by building on the canonical model of sequential search without recall (Gronau 1971, Lippman and McCall 1976).

Several papers on delegated search address the outsourcing of a search process. Erat and Krishnan (2012) study the case of a principal who seeks to solve a design problem by setting up a contest among several agents. Their paper differs from ours in terms of how the process is outsourced (contract vs. contest) and also in the nature of the process being outsourced: Erat and Krishnan study a single-period problem in which the key decision is *where* to search, as in the classic case of a treasure buried in one of several possible locations; in contrast, we use the canonical model of *sequential* search. Rahmani and Ramachandran (2021) examine search delegation in the context of innovation, focusing on whether it is beneficial for the principal to commit to a stopping rule for the search process. They do not identify the optimal contract, instead viewing the contract between the principal and the agent as exogenous, with its form restricted to a tractable parametric family.

Lewis (2012) explores delegation of a sequential search process, but his model differs from ours with regard to the nature of the process being outsourced and to the source of agency frictions. Lewis presents a search model in which alternatives represent technologies, and the principal receives a continuous income stream from the best technology identified so far. Thus there is no decision to be made about which alternatives to accept and which to reject; the only decision involved in forming a search strategy concerns the amount of effort to put into it. Hence there is also no adverse selection problem with the agent; however, moral hazard arises because he may exert less effort than the principal would prefer (our model features both moral hazard and adverse selection). The model we develop also differs vis-à-vis contracting: whereas Lewis considers a riskneutral agent with limited liability, the agent in our model is risk-averse and has unlimited liability. Ulbricht (2016) develops a model where search with recall is outsourced to a risk-neutral agent with no capital constraints. The source of agency frictions in that model is information asymmetry, as the principal does not know the distribution of the alternatives. Also related to this stream is the study by Poblete and Spulber (2017), who examine R&D partnerships in which one party is responsible for the invention and the other for innovation (where innovation is modeled as a sequential search); these authors find that the optimal contract amounts to call options. The model of Gerardi and Maestri (2012) studies delegation of information acquisition, which closely resembles a delegated version of the technology adoption model of McCardle (1985); they find the optimal contracts are performance-based, consisting of a single payment upon conclusion of the search. Kleinberg and Kleinberg (2018) focus on using "prophet inequalities" (Samuel-Cahn 1984) to establish theoretical bounds on the efficiency of delegated search.

Our work contributes to this literature by being the first to examine several aspects of delegated search. More specifically, we assess the effect of an agent's risk aversion and describe how not only the search dynamics but also the optimal contract evolve in the presence of a search deadline. We believe that this study is also the first to explore the topic – which is highly relevant from a practical standpoint – of when a search process should be outsourced instead of performed in-house.

From the perspective of methodology, the delegation of search is a problem in dynamic contract theory. The literature on contract theory and incentive design is vast, with a long history of applications and impact across many fields. Good reviews of the foundations of this literature are provided by Hart and Holmström (1987), Bolton and Dewatripont (2005), and Laffont and Martimort (2009). The study of dynamic contracting problems began with Rubinstein and Yaari (1983), Radner (1985), and Malcomson and Spinnewyn (1988); these papers consider multiperiod versions of classic contracting problems. One issue that has burdened dynamic contract theory models is the tendency for optimal contracts to be complex, history-dependent, and seldom expressible in analytic form. Spear and Srivastava (1987) show that – under the fairly mild assumptions of a risk-neutral agent, limited liability, and a common discount rate – history dependence can be made tractable because all of a history's relevant information can be reduced to a single variable: the agent's promised future utility. Spear's solution technique sparked a multitude of papers, primarily in economics and finance. The delegated search paper of Lewis (2012) relies on this approach. In the operations literature, an alternative technique is proposed by Plambeck and Zenios (2000); these authors develop a *dynamic principal-agent framework* for solving the delegation of Markov decision processes (MDPs). They employ the methodology of Fudenberg et al. (1990) and Smith (1998) and show that, given a risk-averse agent who maximizes utility from consumption and can access efficient banking, the contracting problem's temporal complexity can be decomposed into a sequence of (tractable) single-period problems. This solution technique, which has been adopted by several papers in operations management, is the one that we shall use.

The application of dynamic contracting is of interest also to the operations research (OR) and management science (MS) communities. One stream of literature builds on the dynamic principalagent approach of Plambeck and Zenios (2000). For instance, Plambeck and Zenios (2003) use it to examine the outsourcing of a make-to-stock single-server queueing system, and Fuloria and Zenios (2001) study the design of incentives for healthcare services delivery. Zhang and Zenios (2008) and Zhang (2012a,b) propose a related framework for the delegation of control over MDPs; in these papers, it is hidden information (not hidden actions) that creates agency frictions. Our paper contributes to this literature stream by extending the framework to systems in which the principal makes active monitoring decisions. Most other OR/MS treatises on dynamic contracting problems adopt the approach of "promised future utility": Li et al. (2013) model multiperiod contracting in a supply chain setting where the suppliers compete with each other; Belloni et al. (2016) develop computational methods for a dynamic procurement setting in which the optimal contract lacks closed form; Balseiro et al. (2019) examine a dynamic resource allocation problem that excludes the possibility of monetary transfers; Sun and Tian (2018) consider a continuous-time setting wherein the principal induces the agent to exert effort, which in turn increases the arrival rate of favorable events; and Chen et al. (2020) focus on determining an optimal monitoring schedule in a setting that is similar but in which agent effort *reduces* the arrival rate of *unfavorable* events.

Contributions and Key Results. To the best of our knowledge, this paper is the first to address the two central aspects of delegated search described next.

1. Finite horizon. A plethora of search decisions are made under the time pressure of an impending horizon. Many job markets, from the one for freshly minted MBAs to the one for star football players, have a clearly defined time horizon over which recruiting occurs. It has long been established that the presence of a horizon has a critical effect on optimal search strategies: the searcher becomes increasingly less selective as the end of the horizon approaches (Lippman and McCall 1976, Van den Berg 1990). However, such horizons have yet to be studied in the context of delegated search.¹ We shall demonstrate that a finite horizon affects delegated search even more than it does own search: Proximity of the horizon is a critical factor in decisions about whether or not to delegate search. The horizon also affects the type and structure of the optimal contract.

2. The outsourcing decision. The literature on delegated search is chiefly concerned with the question of how best to outsource a search process. However, in real-world search settings, an omnipresent conundrum is whether the search should be delegated at all. Although companies that seek to fill a position can delegate the search process to an agent (a recruiter or a headhunter), they have the option of conducting the same process in-house. Under what conditions is one or the other of these options optimal? Our model is designed to tackle this question. We show that the ages-old aphorism "if you want it done *right*, do it yourself" holds, as the principal's sensitivity to quality leads to the optimality of in-house search. Yet, we also find that, instead, delegation is optimal when the principal is more sensitive to speed than to quality. Hence we dare add to the aphorism: "if you want it done *fast*, hire someone else to do it."

 $^{^{1}}$ In the delegated search literature, only the model of Rahmani and Ramachandran (2021) has a finite horizon. However, that paper does not solve the principal's contracting problem.

In addition, we derive optimal contracts between principal and agent in the context of delegated search – in other words, contracts that enable the principal to attain her second-best outcome when outsourcing. The optimal contracts must be identified if we are to compare the performance of outsourcing and in-house search. Interestingly, we find that the specificities being studied in this paper lead to optimal contracts that are different from ones in the prior literature.

3. Optimal contract type. If the values of alternatives found in a search process are verifiable, then the contract types that emerge as optimal in our model consist of a per-time fee and a bonus for delivering an alternative of sufficiently high value. It turns out that the per-time component of the optimal contract vanishes if there is no opportunity cost for accepting a contract, a result that mirrors observations from practice. However, our findings differ from current business customs with respect to how the bonus amount is determined and when the bonus is awarded. In practice, bonuses are either fixed or a function of the search result's value – as when a real estate agent receives some fraction of a house's sales price, or a headhunter receives some percentage of the recruited employee's first-year salary. Yet, in our model, the optimal bonus size is determined ex-ante (when the contract is signed) and decreases with time, while the bonus is sometimes awarded even for alternatives that the principal is not willing to accept. Our results suggest that existing contracts could be improved if the principal is endowed with sufficient information about the parameters of the search process to be able to construct such a contract optimally.

If the values of the alternatives brought to the principal are *not* verifiable, then two contract types emerge as optimal: a *per-delivery* contract, which pays the agent for every alternative delivered (regardless of its quality); and a *purchase-right* contract, under which the principal has the right to purchase any of alternatives found by the agent. In this latter arrangement, the agent receives a pre-specified bonus (and thus is rewarded only for the alternatives that the principal accepts). An interesting result is that if the search horizon is short enough, then there is no loss of efficiency from the inability to contract on the quality of the alternatives.

2. Model

We begin by introducing the canonical sequential search model (Gronau 1971, Lippman and McCall 1976). This is the setting faced by a principal who is conducting an in-house search process – that is, without hiring an agent. It is also an essential building block of our delegated search model.

2.1. Canonical Model: Sequential Search without Recall

The search process lasts up to N time periods, during each of which the following sequence of events occurs. The searcher first decides on whether or not to pay the search cost c. If he does, then with probability p, he finds an alternative (of value u) via a random draw from a distribution

with cumulative distribution function (cdf) $H(\cdot)$, such that supp $H \subseteq [\underline{u}, \overline{u}]$, where $0 \leq \underline{u} < \overline{u}$. The searcher then decides whether to keep this drawn alternative, receiving its value as a payoff and ending the search process, or to continue searching. Thus, the searcher can only accept one alternative. The searcher's goal is to devise a policy that will maximize his net present value (NPV) for a given per-period discount factor $\alpha \in (0,1)$. Denoting by V_n the searcher's value function, it follows from Lippman and McCall (1976) that a decreasing threshold policy $\{\overline{\xi}_n\} = \{\alpha V_{n+1}\}$ is optimal. That is, in period n, it is optimal to accept any alternative with a value of at least αV_{n+1} and to reject all others. Denoting $\overline{H}(x) := P(X \ge x)$, where X is a random variable distributed according to cdf H, the value function is then given recursively by

$$V_n = \left(p \int_{\alpha V_{n+1}}^{\bar{u}} \bar{H}(x) \, dx - c \right)^+ + \alpha V_{n+1} \tag{1}$$

with boundary condition $V_{N+1} = 0$. This condition represents the fallback option – that is, the searcher's payoff if the process is completed (N periods have passed) and no alternative has been accepted. The fallback option's value can be normalized to zero without loss of generality.

Several properties of the searcher's optimal policy are relevant for our analysis. If the (bracketed) term on the right-hand side (RHS) of (1) is positive, then it is optimal to pay the search cost; otherwise, it is optimal not to pay it. If there is any period during which paying the search cost is optimal, then paying it is optimal in every period – that is, both before and after.² In short: either it is optimal never to search (a degenerate solution), or it is optimal to search in every period.

2.2. Delegated Search Process

The principal can delegate control of the search process to an agent. (To facilitate the exposition, we continue using feminine and masculine pronouns for the principal and the agent, respectively.) She can therefore offer the agent a contract s that specifies how he will be paid for this task. Informally, the contract is a sequence of functions $s = \{s_n(\cdot)\}$, each of which describes the payment to the agent in period n, which depends on the history of publicly observable events until time of payment; we formalize it once all the relevant notation has been introduced. The decision of whether to search in period n is denoted by $\omega_n \in \{0, 1\}$ (1 standing for the decision to search). If the search finds an alternative, the agent observes its value, which is drawn from a distribution with cdf $F(\cdot)$. We assume that the distribution specified by $F(\cdot)$ is discrete with finite support supp $F \subset [u, \bar{u}]$ (where $0 \leq u < \bar{u}$) and a constant probability mass function $\mu(\cdot)$ (so $\mu(x) = 1/|\operatorname{supp} F|$ for all $x \in \operatorname{supp} F$). Such $F(\cdot)$ can be interpreted as a discrete approximation of a continuous distribution and $G(\cdot)$ with support $[\underline{u}, \bar{u}]$; two common discretization techniques which satisfy our assumptions and

 $^{^{2}}$ This property is established in the proof of Lemma A1 (see Appendix B).

can be made arbitrarily precise are discussed in Appendix C.1. (A variant where the distribution of alternatives is itself continuous is considered in Appendix C.2.)

After observing the value, the agent must decide whether to report the alternative to the principal. We denote this decision by $\Xi_n \in \mathcal{F}$, the set of all alternatives that will be reported to the principal if found, where \mathcal{F} is the set of all subsets of supp F. If the alternative is reported to the principal, the principal will incur an evaluation cost e, observe the value of the alternative u, and decide whether to accept it and end the search process (receiving u as the payoff), or reject it and continue. We denote this decision of the principal by $\Psi_n \in \mathcal{F}$, the set of all alternatives that she will accept. Thus, the search will conclude at a stochastic stopping time T, which is either the time at which the principal accepts an alternative, or N, if an alternative is never accepted. The evaluation cost e is key to our model because it motivates the principal would prefer to evaluate all alternatives herself, and so the problem would be reduced to just dynamic moral hazard.

We refer to the sequence of agent's decisions on whether to search as his search policy $\{\omega_n\}$, which alternatives he should report to the principal as his reporting policy $\{\Xi_n\}$, and to the principal's decision about which alternatives to accept as her acceptance policy $\{\Psi_n\}$. All of these policies are dynamic, so can be contingent on history of the process until the time at which the decision needs to be made; we formalize the policies in Section 2.3. This process is illustrated in Figure 1.



This model features complete but imperfect information; that is, both players are aware exante of the game's parameters (e.g., the search cost c, the probability p of finding an alternative, the distribution $F(\cdot)$), yet information asymmetries arise over time because the principal cannot observe all of the agent's actions. More precisely, the principal is not able to observe whether the agent actually searched in any particular period (i.e., he paid the search cost); neither does she know whether the agent found an alternative – that is, unless he reports it to her for evaluation.

We express the common knowledge about period n of the search process through

 $X_n = \begin{cases} \emptyset & | \text{ If } no \text{ alternative was reported to the principal in period } n; \\ (u, A) & | \text{ If an alt. of value } u \text{ was reported to and accepted by the principal in period } n; \\ (u, R) & | \text{ If an alt. of value } u \text{ was reported to but rejected by the principal in period } n. \end{cases}$

Using X_n , we can represent the delegated search process as a competitive Markov decision process with state variable X_n , starting in $X_0 = \emptyset$. Denote the set of all possible states in period n by $\mathcal{X}_n :=$ $\{\emptyset\} \cup \{(u, R) | u \in \operatorname{supp} F\} \cup \{(u, A) | u \in \operatorname{supp} F\}$ and denote the set of possible states where no alternative was accepted by $\bar{\mathcal{X}}_n := \{\emptyset\} \cup \{(u, R) | u \in \operatorname{supp} F\}$. We also define $X^n := (X_0, X_1, X_2, ..., X_n)$, which is the state history of the Markov process up to the end of period n. Denote by P(S) the probability that a random variable with cdf $F(\cdot)$ has a realization in set $S \in \mathcal{F}$. If the current state is $X_{n-1} \in \overline{\mathcal{X}}_{n-1}$, where $n \in \{1, ..., N\}$, then the search is still ongoing, and the state transitions of the Markov process are governed by the distribution $F(\cdot)$ and decisions ω_n, Ξ_n, Ψ_n as follows. With probability $1 - \omega_n p P(\Xi_n)$, the agent does not find an alternative that he reports to the principal, and the state transitions to $X_n = \emptyset$. For every $u \in \Xi_n$, the agent finds an alternative of value u with probability $\omega_n p\mu(u)$, which he reports to the principal. If $u \in \Psi_n$, the principal accepts it, resulting in state transition to the terminal state $X_n = (u, A)$, otherwise she rejects it and the state transitions to $X_n = (u, R)$. Defining all X_N states to be terminal as well (when the horizon N is reached), the stopping time T (the time at which the search process is completed) is the time at which a terminal state is reached. With this is mind, the set of all possible state histories up to period n is $\mathcal{X}^n := \{ (X_0, ..., X_n) | (X_0 = \emptyset) \land (\forall k \in \mathbb{N}, k < n : X_k \in \overline{\mathcal{X}}_k) \land (X_n \in \mathcal{X}_n) \}, \text{ and the set of possible histo-}$ ries where no alternative was accepted is $\bar{\mathcal{X}}^n := \{(X_0, ..., X_n) | (X_0 = \emptyset) \land (\forall k \in \mathbb{N}, k \le n : X_k \in \bar{\mathcal{X}}_k) \}.$

An important property for the tractability of the model is that the private knowledge acquired by the agent in one period (e.g., the value of the alternative when it was not reported to the principal) has no bearing on any probabilities in the future periods, due to the search process consisting of i.i.d. draws from the same known distribution.³

The history variable X^n also serves as the sole contractable variable in our main model. We can formalize the contract as $s = \{s_n(X^n)\}$, where for every $n, s_n(\cdot) : \mathcal{X}^n \to \mathbb{R}$. In other words, s is a sequence of payment rules, one for each period, with each payment rule being a function of the state history. We assume that, in any given period, contractual payments are made at the start of

³ This property would not apply, and the methodology of Fudenberg et al. (1990) and Plambeck and Zenios (2000) we rely on would not work, if, e.g., the distribution $F(\cdot)$ was not known and had to be inferred from observations. In this case, every time the agent observes a draw which the principal does not, that event would endow him with private knowledge about the distribution, which will remain relevant throughout the search horizon.

the period – i.e., before any actions in that period are taken. Once the search process terminates (at time T), there is one final payment to agent based on the terminal state, after which the contractual relationship dissolves. This formulation allows for a large variety of contracts, e.g., paying per unit of time, paying a percentage of the final result's value, and more complex arrangements such as the "tenure evaluation" contract of Sun and Tian (2018).

Two problems central to contract theory are moral hazard and adverse selection (Bolton and Dewatripont 2005). The information asymmetries in our model give rise to both of these problems. First, *moral hazard* is present because the principal cannot distinguish between the agent searching unsuccessfully and exerting no search effort. Hence the agent is incentivized to forgo paying the search cost and to report that he did not find a suitable alternative. Second, *adverse selection* is also an issue, because the agent might find an alternative that the principal would accept, yet not present it to her (or he might present her with alternatives of unacceptable value). For example, this can occur within a contract that pays the agent only for his time (i.e., for the duration of a search process), which naturally creates an incentive to prolong the search. While the adverse selection dimension of the problem comes down to a binary decision of whether to report or not, in section 4.3 we consider a variant where the agent can also misrepresent the value of the alternative, which is more in line with the more general adverse selection settings in the literature.

2.3. The Principal-Agent Problem of Delegated Search

We first describe the problem faced by an agent under a contract $\{s_n(X^n)\}$. In each period, the agent decides whether to pay the search cost (i.e., he determines ω_n) and also decides which found alternatives to bring to the principal (i.e., he determines Ξ_n). We assume that the agent is risk-averse with constant absolute risk aversion (CARA), where r is his coefficient of risk aversion, and that he seeks to maximize his utility from consumption. The agent starts with a certain wealth W_0 , and in each period, he determines γ_n – that is, how much wealth to consume (where the wealth not consumed is saved at accruing interest rate α^{-1}). So, in addition to making decisions about the search process, the agent faces a classic consumption–saving problem (Merton 1990). We assume frictionless access to banking: the agent's wealth can be negative, in which case he is charged the same interest rate α^{-1} on any loan required to maintain solvency. Suppose, by way of illustration, that the agent starts period n with wealth W_n and decides to pay the search cost c and to consume γ_n ; then he will receive utility $-\exp(-r\gamma_n)$, and his wealth in the next period will be $W_{n+1} = (W_n - \gamma_n - c)/\alpha + s_n(X^n)$. An agent who decides not to sign a contract with the principal can earn wage w in each period, and he will be able to earn w after any signed contract ends.

This combination of assumptions (CARA utility from consumption, frictionless banking, same discount rate) follows the approach of Fudenberg et al. (1990) and Plambeck and Zenios (2000) and

is essential for tractability of the principal-agent component of the model. What these assumptions jointly accomplish is that they remove the timing of the payment as an incentive device for the agent. In other words, they guarantee that the agent's expected utility depends on the payments only through how the payments affect his terminal wealth. Thus, the agent will be indifferent between all sequences of payments that have the same NPV.

With this in mind, we can formalize the set of admissible policies. The search policy $\{\omega_n\}$ is a sequence of functions $\{\omega_1(\cdot), ..., \omega_N(\cdot)\}$, each of which maps a state history in a period to a decision on whether to search or no in that situation, so for every $n \in \{1, ..., N\}$, $\omega_n(\cdot) : \bar{\mathcal{X}}^{n-1} \to \{0, 1\}$.⁴ Analogously, the reporting policy $\{\Xi_n\}$ is a sequence of functions that map the state history to a set of alternatives to be reported if found $(\Xi_n(\cdot) : \bar{\mathcal{X}}^{n-1} \to \mathcal{F})$, the acceptance policy $\{\Psi_n\}$ is a sequence of functions that map the state history to a set of alternatives to accept if reported $(\Psi_n(\cdot) : \bar{\mathcal{X}}^{n-1} \to \mathcal{F})$, and the consumption policy $\{\gamma_n\}$ is a sequence of functions that map state history and current wealth to the consumption decision for that period $(\gamma_n(\cdot) : \bar{\mathcal{X}}^{n-1} \times \mathbb{R} \to \mathbb{R})$.⁵ An important property here – which holds for all of the policies just described and follows from their domain being $\bar{\mathcal{X}}_n$ – is that they only need to specify what to do in the situation where the principal has not accepted an alternative yet; once a principal accepts an alternative, a terminal state is reached and there are no more decisions to be made.

The agent seeks to maximize expected utility from consumption, which we concisely express as

$$\mathbb{E}\left[-\left(\sum_{n=1}^{T}\alpha^{n-1}\exp(-r\gamma_n(X^{n-1},W_{n-1}))\right) - \frac{\alpha^T}{1-\alpha}\exp[-r(w+(1-\alpha)W_T)]\Big|W_0\right].$$
 (2)

Here, the sum is the discounted utility from consumption, whereas the $\alpha^T/(1-\alpha)\exp[-r(w+(1-\alpha)W_T)]$ term is the terminal utility after the dissolution of the contract, and reflects that the optimal consumption policy after contract completion is to consume the earned wage w as well as interest on terminal wealth in perpetuity.

The principal is risk-neutral and maximizes her (discounted) expected wealth. A principal's wealth increases when the search process is completed and she receives the search's result u as a payoff; her wealth is decremented by payments made to the agent and by her evaluation cost e. Thus, the principal aims to maximize

$$\mathbb{E}\left[-\left(\sum_{n=1}^{T}\alpha^{n}s_{n}(X^{n-1})+\mathbb{1}_{(X_{n\neq\emptyset})}\alpha^{n-1}e\right)+\alpha^{T}u_{T}\right],$$
(3)

⁴ Note the slight abuse of notation here where $\omega_n(\cdot)$ refers to a function that maps histories to decisions, but we also use ω_n (without arguments) to refer to the decision itself. The same convention is used for other policies. However, it should be clear from the context whether we refer to the decision or the policy. This issue also vanishes immediately following Theorem 2, which establishes the optimality of history-independent policies.

⁵ We could also allow the agent to make the policies $\{\omega_n\}$ and $\{\Xi_n\}$ contingent on his current wealth W_n . However, this option will never be used due to the *absence-of-wealth-effects* property established in Fudenberg et al. (1990), where the agent's CARA utility in combination with efficient banking implies that the optimal policies $\{\omega_n\}$ and $\{\Xi_n\}$ chosen by the agent in response to any contract do not depend on his wealth.

where u_T is the value of the alternative that the principal accepted or 0 if no alternative was accepted. While (2)-(3) clearly reflect the sources of both the principal and the agent's goals, they suppress complexity inside the parameters (e.g., the parameter W_T , which is critical for the agent's utility, will depend on all of the agent's and the principal's decisions as well as randomness).

To expose the full complexity of this dependence, consider the principal's value function in period n, denoted by $V_n^P(\{s_n(X^n)\}, \{\omega_n\}, \{\Xi_n\}, \{\Psi_n\}, X^{n-1})$. This term corresponds to her expected wealth, starting from period n, with history X^{n-1} , if she adopts acceptance policy $\{\Psi_n\}$ while the agent holds a contract $\{s_n(X^n)\}$, uses the search policy $\{\omega_n\}$, and the reporting policy $\{\Xi_n\}$. Using \oplus to denote concatenation (so e.g., $(X_0, ..., X_n) \oplus X_{n+1} = (X_0, ..., X_{n+1})$), for every non-terminal history $X^{n-1} \in \overline{X}^{n-1}$, $n \leq N$, we can express V_n^P recursively by:

$$\begin{aligned} V_{n}^{P}(\{s_{n}(X^{n})\},\{\omega_{n}\},\{\Xi_{n}\},\{\Psi_{n}\},X^{n-1}) &= \\ & \omega_{n}p \int_{\Xi_{n}\cap\Psi_{n}} \left(u-e-\alpha s_{n}(X^{n-1}\oplus(u,A))\right) dF(u) \\ & + \omega_{n}p \int_{\Xi_{n}\setminus\Psi_{n}} \left(-e-\alpha s_{n}(X^{n-1}\oplus(u,R)) + \alpha V_{n+1}^{P}(\{s_{n}(X^{n})\},\{\omega_{n}\},\{\Xi_{n}\},\{\Psi_{n}\},X^{n-1}\oplus(u,R))\right) dF(u) \\ & + \left(1-\omega_{n}p \int_{\Xi_{n}} dF(u)\right) \left(-\alpha s_{n}(X^{n-1}\oplus\emptyset) + \alpha V_{n+1}^{P}(\{s_{n}(X^{n})\},\{\omega_{n}\},\{\Xi_{n}\},\{\Psi_{n}\},X^{n-1}\oplus\emptyset)\right), \end{aligned}$$

subject to boundary condition $V_{N+1}^P(\{s_n(X^n)\}, \{\omega_n\}, \{\Xi_n\}, \{\Psi_n\}, X^n) = 0.^6$ In (4), the second line is the scenario where the agent finds an alternative (which happens with probability $\omega_n p$), he reports that alternative to the principal (so its value is in Ξ_n) and she accepts the alternative (the value is also in Ψ_n), in which case the principal incurs the evaluation cost e, and pays the agent what the contract stipulates for that history $(s_n(X^{n-1} \oplus (u, A)))$, but receives the value of the accepted alternative u, and the search process ends (thus the absence of V_{n+1}^P in this scenario).

The third line is the scenario where the agent finds an alternative that he reports but the principal rejects, in which case the principal incurs the evaluation cost, the report and the rejection are added to the history variable X^{n-1} , the principal pays the agent what the contract stipulates for the just updated history, and the process proceeds to the next period.

The last line is the scenario in which the agent does not report any alternative to the principal, either due to not searching at all, or searching and not finding an alternative that is suitable to be reported ($u \in \Xi_n$). As in the previous case, the lack of report is added to the history variable X^{n-1} , the principals pays the agent what the contract stipulates for the just updated history, and the process proceeds to the next period.

⁶ As F(u) is discrete, integration in (4) could be replaced with summation. We will maintain integral notation though, as a) it will lead to more concise expressions in many places, b) it maintains notational consistency with both the canonical model and the variant of our model where $F(\cdot)$ is continuous, as given in Appendix C.

We can express the agent's value function in the same way:

$$\begin{aligned} V_{n}^{A}(\{s_{n}(X^{n})\},\{\omega_{n}\},\{\Xi_{n}\},\{\Psi_{n}\},\{\gamma_{n}\},W_{n},X^{n-1}) &= -\exp(-r\gamma_{n}) \\ &+ \omega_{n}p \int_{\Xi_{n} \cap \Psi_{n}} -\frac{\alpha}{1-\alpha} \exp\left[-r\left(w+(1-\alpha)\left(\frac{W_{n}-c-\gamma_{n}}{\alpha}+s_{n}(X^{n-1}\oplus(u,A))\right)\right)\right)\right] dF(u) \\ &+ \omega_{n}p \int_{\Xi_{n} \setminus \Psi_{n}} \alpha V_{n+1}^{A}(\{s_{n}(X^{n})\},\{\omega_{n}\},\{\Xi_{n}\},\{\psi_{n}\},\{\gamma_{n}\},\frac{W_{n}-\gamma_{n}-c}{\alpha}+s_{n}(X^{n-1}\oplus(u,R)),X^{n-1}\oplus(u,R)) dF(u) \quad (5) \\ &+ \left(1-\omega_{n}p \int_{\Xi_{n}} dF(u)\right) \alpha V_{n+1}^{A}(\{s_{n}(X^{n})\},\{\omega_{n}\},\{\Xi_{n}\},\{\psi_{n}\},\{\gamma_{n}\},\frac{W_{n}-\gamma_{n}-\omega_{n}c}{\alpha}+s_{n}(X^{n-1}\oplus\emptyset),X^{n-1}\oplus\emptyset),\end{aligned}$$

subject to boundary condition $V_{N+1}^A(\{s_n(X^n)\},\{\omega_n\},\{\Xi_n\},\{\psi_n\},\{\gamma_n\},W_{N+1},X^N) = -\frac{1}{1-\alpha}\exp[-r(w+(1-\alpha)W_{N+1})]$ (this is the agent's terminal utility, as in (2)).

In the expression above, there are a few additional details that are not present in the principal's value function. The right-hand side of the first line in the equation is the utility from consuming γ_n . The second line is the scenario where the agent finds and reports an alternative that the principal accepts, in which case the integrand is his terminal utility after his wealth has been adjusted for the search cost, consumption, and the final payment for successfully concluding the search process. The third line is the scenario where the agent finds an alternative that he reports but the principal rejects, in which case the process proceeds to the next period with both state variables (wealth and history) updated. The last line is the scenario where the agent finds and report did not report any alternative.

The value functions (4)-(5) are cumbersome to deal with directly (and we will avoid doing so until we derive results that enable their simplification: Lemma 1 and Theorems 1 and 2), but they allow us to formalize the dynamic principal agent problem as follows. The principal aims to solve what we refer to as the *main contracting problem*, which is the maximization in (6)

$$\max_{\{s_n(X^n)\},\{\omega_n\},\{\Xi_n\},\{\Psi_n\}} V_1^P(\{s_n(X^n)\},\{\omega_n\},\{\Xi_n\},\{\Psi_n\},X^0)$$
(6)

subject to the three constraints described next.

First, an incentive compatibility constraint dictates that the contract $\{s_n(X^n)\}$ implements the search policy $\{\omega_n\}$ and the reporting policy $\{\Xi_n\}$; that is, the pair $(\{\omega_n\}, \{\Xi_n\})$ needs to maximize the agent's expected utility when holding contract $\{s_n(X^n)\}$:

$$(\{\omega_n\},\{\Xi_n\},\{\gamma_n\}) \in \underset{\{\tilde{\omega}_n\},\{\tilde{\Xi}_n\},\{\tilde{\gamma}_n\}}{\arg\max} V_1^A(\{s_n(X^n)\},\{\tilde{\omega}_n\},\{\tilde{\Xi}_n\},\{\Psi_n\},\{\tilde{\gamma}_n\},W_0,X^0).$$
(7)

Second, there is an incentive compatibility for the principal as well. In particular, it must be in her best interest to use the acceptance policy $\{\Psi_n\}$ once the contract $\{s_n(X^n)\}$ has been signed:

$$\{\Psi_n\} \in \operatorname*{arg\,max}_{\{\tilde{\Psi}_n\}} V_1^P(\{s_n(X^n)\}, \{\omega_n\}, \{\Xi_n\}, \{\tilde{\Psi}_n\}, X^0).$$
(8)

Third, there is a participation (individual rationality) constraint: the agent will accept the contract only if it offers him expected utility that is no less than the disagreement payoff. Formally,

$$V_1^A(\{s_n(X^n)\},\{\omega_n\},\{\Xi_n\},\{\Psi_n\},\{\gamma_n\},W_0,X^0) \ge -\frac{1}{1-\alpha}\exp[-r(w+(1-\alpha)W_0)].$$
(9)

Equations (6)–(9) jointly form a dynamic contracting problem, which we now proceed to solve. We will focus on the case where the solution of this problem results in a strictly positive value function for the principal; otherwise, there would be no value to be generated from the delegation, and the principal should not offer any contract. The question of whether the principal should be delegating search in the first place is studied in greater detail in Section 3. Our notation is summarized in Table EC.1 in Appendix A (which includes the terms that have not yet been introduced).

2.4. Analysis

The solution of the main contracting problem (6)-(9) is derived in two steps. First, we identify which search and reporting policies for the agent are implementable and determine, for each such policy pair, the contract that implements them at the lowest possible cost to the principal. This step is the focus of Theorem 1. Second, we use that theorem's results to reduce the main contracting problem to a classic dynamic program in which the principal optimizes over two choices: the acceptance policy and the induced reporting policy. Then the solution to this problem is the one given in Theorem 3, which completely characterizes the optimal contract.

Theorem 1 (Implementable Policies) The principal can implement any policy pair $(\{\omega_n\}, \{\Xi_n\})$, provided it satisfies the implementability condition

$$pP(\Xi_n) > \omega_n - \frac{1}{\exp[rc(1-\alpha)\alpha^{-1}]},\tag{10}$$

for all $n \in \{1, ..., N\}$. This can be done by offering the agent a contract that pays a base pay of s_n^0 in every period, replaced by bonus $s_n^B(\Xi_n)$ if the agent delivers an alternative with a value in Ξ_n in a period in which the principal would like him to search ($\omega_n = 1$), where

$$s_n^B(\Xi_n) := \frac{w+c}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{pP(\Xi_n)}{1 - (1 - pP(\Xi_n)) \exp[rc(1-\alpha)\alpha^{-1}]} \quad and \quad s_n^0 := \frac{w}{\alpha}.$$
 (11)

This contract implements $(\{\omega_n\}, \{\Xi_n\})$ at the lowest possible cost to the principal (optimized over the space of all contracts), and the contract is renegotiation proof.⁷ Under any such contract, the agent's optimal consumption policy is given by $\gamma_n(X^{n-1}, W_n) = (1-\alpha)W_n + w$, for all $n \in \{1, ..., N\}$.

⁷ A contract $\{s_n(X^n)\}$ is said to be *renegotiation proof* if, for any possible (partial) history of the game, there exists no contract $\{\bar{s}_n(X^n)\}$ such that both the principal and the agent would prefer to substitute it for $\{s_n(X^n)\}$.

From the history X^n , only one event is relevant to the optimal contract: whether an alternative that the principal would like reported (one in Ξ_n) was delivered in the current period. Formally,

$$s_n(X^n) = \begin{cases} s_n^B(\Xi_n) & \text{if } X^n \in \{(u, A), (u, R) | u \in \Xi_n\} \text{ and } \omega_n = 1\\ s_n^0 & \text{otherwise.} \end{cases}$$

This still does not exclude history-dependent contracts, as the principal can make the contract history-dependent by choosing to induce $\{\omega_n\}$ and $\{\Xi_n\}$ which are themselves history-dependent. However, we will show in Theorem 2 that is not an option she would like to use.

We observe that the principal can induce the agent to adopt almost any policy. In any period, she can induce him not to search ($\omega_n = 0$) by paying him only s_n^0 irrespective of the search outcome, or she can induce him to search and report all alternatives with a value in Ξ_n , by also awarding an appropriately sized bonus $-s_n^B(\Xi_n)$, as given by (11) – when such alternatives are found. The sole exception is that she cannot induce him to search for an extremely rare alternative, because that would violate the implementability condition (10). This limitation stems from the agent's CARA utility (which is bounded from above); thus, in situations where the bonus is awarded very rarely, even infinitely large bonus size will be insufficient to incentivize the agent to pay the search cost. For these reasons, the number of policies that cannot be implemented increases with the agent's risk aversion and also with his search cost. These limitations vanish as $c \to 0$ or $r \to 0$, which we can infer from the implementability condition.

The contract that induces any implementable policy pair $(\{\omega_n\}, \{\Xi_n\})$ at the lowest possible cost to the principal consists of paying the agent a fixed fee s_n^0 in every period, and a variable bonus s_n^B in periods where the $\omega_n = 1$ and the agent delivers to the principal an alternative with a value in Ξ_n . The size of this bonus depends not on the value of the delivered alternative but rather on $pP(\Xi_n)$, which is the probability that the agent finds an alternative of acceptable value. Observe that s_n^B is increasing in $pP(\Xi_n)$, and so, if the principal wants the agent to be more selective in his search, she will have to offer a higher bonus.

To solve the principal's main contracting problem, it remains to determine which policies $(\{\omega_n\}, \{\Xi_n\})$ she would prefer to implement and which acceptance policy $\{\Psi_n\}$ she should use. As those are all decisions of the principal, Theorem 1 reduces the problem to an MDP in which her decision is the triple $(\{\omega_n\}, \{\Xi_n\}, \{\Psi_n\})$. From (4), the Bellman equation of the resulting program is

$$V_n^P = \max_{\omega_n \in \{0,1\}, \Xi_n, \Psi_n \in \mathcal{F}} \omega_n p \int_{\Xi_n \cap \Psi_n} (u - e - \alpha s_n^B(\Xi_n)) dF(u)$$

+ $\omega_n p \int_{\Xi_n \setminus \Psi_n} (\alpha V_{n+1}^P - e - \alpha s_n^B(\Xi_n)) dF(u) + (1 - \omega_n p P(\Xi_n)) (\alpha V_{n+1}^P - \alpha s_n^0)$
= $\max_{\omega_n \in \{0,1\}, \Xi_n, \Psi_n \in \mathcal{F}} \omega_n p P(\Xi_n) \left(\alpha s_n^0 - e - \alpha s_n^B(\Xi_n)\right) + \omega_n p \int_{\Xi_n \cap \Psi_n} (u - \alpha V_{n+1}^P) dF(u) - \alpha s_n^0 + \alpha V_{n+1}^P,$ (12)

with boundary condition $V_{N+1}^P = 0$. This is still not a straightforward problem, so we start by identifying simplifications that can be made without loss of optimality.

Theorem 2 (Optimality of Memoryless Threshold Policies) Without loss of optimality, the choice of $\{\omega_n\}$, $\{\Psi_n\}$, and $\{\Xi_n\}$ in (6)-(9) can be restricted to history-independent policies, i.e., for all $n \in \{1, ..., n\}$: $\omega_n(X^{n-1})$, $\Psi_n(X^{n-1})$, and $\Xi_n(X^{n-1})$ are constant functions. There exist sequences of thresholds $\{\xi_n\}$ and $\{\psi_n\}$, such that under an optimal contract, the agent uses reporting policy $\{\Xi_n\} = \{[\xi_n, \bar{u}] \cap \text{supp } F\}$ and the principal uses evaluation policy $\{\Psi_n\} = \{[\psi_n, \bar{u}] \cap \text{supp } F\}$.

The theorem shows that the optimality of history-independent threshold policies from the canonical model extends to our setting as well. This is a convenient property for our analysis, as it allows us to conduct the optimization in the main contracting problem (6)–(9) over sequences of thresholds $\{\xi_n\}$ and $\{\psi_n\}$ rather than policies $\{\Xi_n\}$ and $\{\Psi_n\}$ (real numbers rather than functions). Thus, henceforth, we will use $\{\xi_n\}$ (resp. $\{\psi_n\}$) in place of $\{\Xi_n\}$ (resp. $\{\Psi_n\}$), referring to it also as the reporting (resp. acceptance) policy. In this notation $\{\omega_n\}, \{\xi_n\}$, and $\{\psi_n\}$ are simply sequences of nnumbers. We also find it useful to introduce notation r x $\gamma := \inf\{y \in \text{supp } F | y \ge x\}$; this operation rounds up x to the closest number in the support of $F(\cdot)$.⁸

Next, we identify the policies that the principal would like the agent to use – that is, if there were no conflicts of interest between the principal and the agent and if the principal could simply dictate all of the agent's actions while reimbursing him for the costs incurred. We refer to these as the *first-best* (FB) policies, and we formally derive them in Lemma A2 (in Appendix B). The principal's FB value function is given by

$$V_{n}^{\rm FB} = \left(p \int_{\alpha V_{n+1}^{\rm FB} + e}^{\bar{u}} \bar{F}(u) du - c \right)^{+} - w + \alpha V_{n+1}^{\rm FB}, \tag{13}$$

with boundary condition $V_{N+1}^{\text{FB}} = 0$. The first-best reporting and acceptance policies are $\{\xi_n^{\text{FB}}\} = \{\forall \alpha V_{n+1}^{\text{FB}} + e \forall\}$. For all $n \in \{1, ..., N\}$, the first-best search policy is

$$\omega_n^{\rm FB} = \begin{cases} 1 & \text{if } p \int_{\alpha V_{n+1}^{\rm FB} + e}^{\bar{u}} \bar{F}(x) \, dx - c > 0; \\ 0 & \text{otherwise.} \end{cases}$$

Comparing (13) to the canonical problem in (1), we can see that the first-best search differs only through the addition of the costs w and e which are absent in the canonical model.⁹ To understand

⁸ Note that $\inf \emptyset = \infty$, so if x is grater than all elements of $\sup F$, then $\forall x \models \infty$.

⁹ Assuming parameter values are such that it is optimal to search in the canonical search problem, the first-best policy $\{\xi_n^{\text{FB}}\}$ has two intuitive interpretations that link it to the canonical policy. First, $\{\xi_n^{\text{FB}}\}$ is also the solution of the canonical problem in which the search cost is w higher, and the evaluation cost needs to be paid only once (i.e., when an alternative is accepted). Second, $\{\xi_n^{\text{FB}}\}$ is the solution of the canonical problem in which the search cost is w higher and the distribution is shifted to the left by e (so alternatives come from G(x) := F(x+e), instead of F(x)).

this resemblance, consider what happens in the very last period: if the principal is searching on her own, then she will accept any alternative with a positive value; if the search is delegated to an agent, then she remains willing to accept any (positive-value) alternative yet would prefer that the agent not bring her alternatives with a value lower than the evaluation cost (because that would then yield a loss). The same need – that is, for the alternative's value to be at least e higher than the value of the rest of the search – also prevails in all previous periods.

The following lemma combines all of the principal's costs into one well-behaved function, which has the additional advantage of simplifying her choice of which reporting policy to induce.

Lemma 1 (The Principal's Costs) The expected single-period costs incurred by the principal when implementing search policy $\{\omega_n\}$ and reporting policy $\{\xi_n\}$ are w if $\omega_n = 0$, and are otherwise given by the quasi-convex function

$$z(\xi_n) = w + pF(\xi_n)(\alpha s_n^B(\xi_n) + e - w).$$
(14)

The restriction $z(\xi_n)|_{\text{supp }F}$ is strictly increasing (strictly decreasing) on $\xi_n \leq \xi^{\text{LC}}$ ($\xi_n > \xi^{\text{LC}}$), where

$$\xi^{\rm LC} := \min_{\substack{\xi_n \in \operatorname{supp} F}} z(\xi_n).$$
(15)

The fundamental trade-off involved in choosing ξ_n is that, while making the agent *more* selective in his search is more expensive for the principal (as s_n^B increases with $p\bar{F}(\xi_n)$), an agent who is *less* selective entails higher expected evaluation costs for the principal (as the principal incurs evaluation cost with probability $p\bar{F}(\xi_n)$). Using Lemma 1, we identify the lowest-cost (LC) reporting policy ξ^{LC} , which resolves this tradeoff. Notice from the last part of the lemma that it is possible for the cost-minimizing policy not to be unique, as $\arg\min_{\xi_n \in \text{supp } F} z(\xi_n)$ can contain up to two (consecutive) elements. In order to break the tie, we assume that if the principal is ever indifferent between these two policies (due to them yielding the same expected utility), she will choose ξ^{LC} .

Thus, using Lemma 1 and Theorem 2, we can reformulate (12) as

$$V_n^P = \max_{\omega_n \in \{0,1\}, \xi_n, \psi_n \in \text{supp } F} \omega_n p \int_{\max\{\xi_n, \psi_n\}}^{\bar{u}} (u - \alpha V_{n+1}^P) dF(u) - (1 - \omega_n) w / \alpha - \omega_n z(\xi_n) + \alpha V_{n+1}^P, \quad (16)$$

with boundary condition $V_{N+1}^P = 0.^{10}$ We now use this result to characterize the optimal (secondbest) contract and its properties.

¹⁰ We could allow the choice of ξ_n and ψ_n to be over $[0, \bar{u}]$ instead of over supp F. Doing so would result in a well-defined problem and an identical value function; however it would introduce ambiguity as the function $x \mapsto [x, \bar{u}] \cap \text{supp } F$ (used to map threshold policies to their set analogues) is not injective for $x \in \mathbb{R}$, but is injective for $x \in \text{supp } F$.

Theorem 3 (Solution of the Main Contracting Problem) The optimal contract induces search policy $\omega_n^* = 1, \forall n \in \{1, ..., N\}$, the reporting policy $\{\xi_n^*\}$ that solves

$$V_n^P = \max_{\xi_n \in \text{supp } F} p \int_{\max\{\xi_n, \alpha V_{n+1}^P\}}^{\bar{u}} (u - \alpha V_{n+1}^P) dF(u) - z(\xi_n) + \alpha V_{n+1}^P,$$
(17)

with boundary condition $V_{N+1}^P = 0$, and the acceptance policy given by $\psi_n^* = r^* \alpha V_{n+1}^P \cap$ for all $n \in \{1, ..., N\}$.¹¹ The optimal contract $\{s_n^*(X^n)\}$ is then given by Theorem 1, as the one that implements $\{\omega_n^*\}, \{\xi_n^*\}$ at the lowest cost to the principal.

The optimal policies exhibit the following properties.

 $(\mathrm{i}) \quad \{\xi_n^*\} \text{ satisfies } \min\{\xi^{\mathrm{LC}}, \mathsf{P} \; \alpha V_{n+1}^P \; \mathsf{I}\} \leq \xi_n^* \leq \min\{\xi^{\mathrm{LC}}, \mathsf{P} \; \alpha V_{n+1}^P + e \; \mathsf{I}\}, \text{ for all } n \in \{1, ..., N\}.$

(ii) $\{\psi_n^*\}$ and $\{V_n^P\}$ are weakly decreasing in n.

(iii) $\{\xi_n^*\}$ is constant and equals ξ^{LC} up to a point, after which it is weakly decreasing in n.

(iv) Both the agent and the principal are less selective than the first-best policy, and so, for all $n \in \{1, ..., N\}$, it holds that $\xi_n^* \leq \xi_n^{\text{FB}}$ and $\psi_n^* \leq \xi_n^{\text{FB}}$.

(v) The policy $\{\psi_n^*\}$ crosses $\{\xi_n^*\}$ at most once from above as n increases. Furthermore, denote $M := \min\{n \in \mathbb{N} \mid \xi_n^* \ge \psi_n^*, n \le N\}$, if M > 1 we then have $\xi_{M-1}^* = \xi^{\text{LC}}$.

One salient feature of the theorem is that the search policy has a simple solution: it is optimal for the principal to have the agent search in all periods until an acceptable solution has been found. Observe that ω_n gives the decision to search or no in period *n* unless the search has reached a terminal state by the proposer accepting an alternative before that. Thus, the choice of ω_n becomes redundant on sample paths where T < n. In essence, the property of the canonical search model – that it is optimal to either search in all periods or to not search at all – extends to the delegated search setting. The solution where the principal hires the agent and then pays him a fixed wage only for the agent to never search does not appear in the theorem, as that one can not generate any value for the principal, so that situation would result in no contract being offered at all.

The other policies given in Theorem 3 are different than what the principal would prefer to do in the absence of agency frictions (the first-best case); in that event, the principal would adopt policy $\{\xi_n^{\text{FB}}\}$. The optimal reporting and evaluation policies are both threshold policies, with the threshold being below the one of $\{\xi_n^{\text{FB}}\}$. In other words: it is optimal for the agent to exhibit less than ideal selectivity when searching, and it is also optimal for the principal to accept most of what the agent delivers despite its quality being less than ideal.

¹¹ Optimal acceptance policies are not unique: when an agent uses reporting policy ξ_n , all evaluation policies $\psi_n \leq \xi_n$ are payoff equivalent (because him bringing an alternative valued below ξ_n is a zero-probability event). We focus on the particular optimal policy ψ_n^* as it is the only optimal policy in supp F that remains optimal even if the agent deviates from the induced reporting policy.



Figure 2 Optimal Threshold Policies (*x*-axis, time (period index); *y*-axis, value of alternatives)

Figure 2 reveals that, when the deadline of the search is near, the principal prefers the agent to bring most of the alternatives he finds. This situation results in minimal agency frictions because the agent's policies with the lowest threshold are those that can be induced via the smallest payment (Theorem 1). As the deadline gets further away, the principal prefers the agent to be more selective in his search. However, the resulting policies – which will have higher ξ_n – are more expensive to induce because they require that the agent be paid a higher risk premium.

When the search deadline is distant, it becomes attractive to the agent to not search, yet to report that nothing suitable was found. This incentive creates a need for the principal to micromanage the agent: she requires the agent to bring her even some alternatives that are of too low quality for her to accept and pays the agent a bonus for those alternatives nonetheless. Such micromanagement serves two functions: it allows the principals to more closely monitor whether the agent is really searching, and it allows the size of the bonus to be smaller than it would otherwise need to be, as the agent will receive it more often. (The agent's risk aversion causes him to prefer smaller bonuses paid more often, *provided* the expectation is the same.) This scenario can be observed in the region of Figure 2 left of the vertical dashed line, where $\{\xi_n^*\}$ is constant at ξ^{LC} while $\{\psi_n^*\}$ rises above it.

The expression for the optimal value function (17) is more complex than in the canonical search model (1) because there is a maximization embedded in the recursion. The source of this complexity is that, although the principal is willing to accept any alternatives with a value at least as high as the expectation for her search (αV_{n+1}^P) , she would actually prefer that the agent *not* bring her alternatives with values below $\alpha V_{n+1}^P + e$, because she incurs an evaluation cost for every delivered alternative. In the first-best solution, this problem is easily resolved by setting $\xi^{\text{FB}} = \beta \alpha V_{n+1}^P + e \beta$. Yet, in the optimal contract, another tradeoff needs to be considered: the payment needed to induce ξ_n increases as ξ_n gets larger. Hence the principal might prefer to set $\xi_n < \alpha V_{n+1}^P + e$, even though it will reduce the expected gain from search, because it will also reduce the expected payment to the agent. This is not a straightforward decision, as even the simplest single-period problem $V_N = \max_{\xi_N} p \int_{\xi_N}^{\bar{u}} x \, dF(x) - z(\xi_N)$ lacks a closed-form solution. Even so, the optimal reporting policy lies in a relatively narrow interval (per Theorem 3(i)). This behavior can also be observed on the right side of Figure 2, where $\{\xi_n^*\}$ is higher (but not more than *e* higher) than ψ_n^* .

2.5. Limiting Results

In the following proposition, we consider four special cases where the model becomes simpler.

Proposition 1 (Limiting Results)

(i) If the horizon is infinite $(N \to \infty)$, then the policies $\{\xi_n^{\text{FB}}\}, \{\xi_n^*\}, \{\psi_n^*\}$ are constant threshold policies such that $\xi_n^{\text{FB}} \ge \psi_n^*, \ \xi_n^{\text{FB}} \ge \xi_n^*$, and $\forall \ \psi_n^* + e \ \forall \ge \xi_n^*$. The bonus size under the optimal contract (s_n^B) is also constant.

(ii) In both the finite- and infinite-horizon cases, if there is no risk aversion $(r \rightarrow 0)$, then the optimal contract converges to the contract under which the agent's bonus exactly compensates for his expected costs:

$$s_n^B \to \frac{w}{\alpha} + \frac{c}{p\bar{F}(\xi_n)\alpha}, \qquad s_n^0 \to \frac{w}{\alpha}.$$

This contract achieves the first-best.

(iii) If there is no evaluation cost (e = 0), the optimal contract induces the agent to report everything he finds to the principal $(\xi_n^* = \min \operatorname{supp} F)$. This contract achieves the first-best only if p = 1.

(iv) If there is no opportunity cost for the agent (w = 0), the optimal contract reduces to one with only a bonus and no per-period fee $(s_n^0 = 0)$.

Many traditional search models in the literature consider an infinite horizon, in which case the optimal policy is a constant threshold, lending analytical tractability to those models. The same property extends to delegated search, where the optimal bonus and optimal delegation policies become time-invariant under an infinite horizon.

If there is no risk aversion, then the optimal contract offers a bonus of $c/(p\bar{F}(\xi_n)\alpha)$ any time the agent finds an alternative of value at least ξ_n . As this occurs with probability $p\bar{F}(\xi_n)$ (if the agent searches), it follows that he receives his search cost back in expected compensation – the discount factor α exactly compensates the agent for his search cost being incurred one period *before* receiving his compensation for that cost. Thus the agent's expected payment is the same regardless of the principal's policy preferences, which in turn allows her to achieve the first-best. Achieving the first-best in the absence of risk aversion is a standard result in contract theory (Bolton and Dewatripont 2005). What is noteworthy here is that the first-best can be achieved using the contract structure of Theorem 1 and without the need to breach limited liability. Part (iii) of Proposition 1 highlights the important role that evaluation cost e plays in the model. The principal would like the agent to make selection decisions (what to bring to the principal) in order to avoid this cost. In the absence of e, it is costless for the principal to verify whether the agent found anything in his search – removing the adverse selection component and reducing it to just dynamic moral hazard. Even in this case, the first-best outcome is not attained (unless p = 1), owing to the uncertainty of finding an alternative, which necessitates a risk premium for the agent.

3. The Decision to Outsource

Until this point, we have focused on how to maximize the efficiency of delegating a search process. Yet, another question remains: Is it a good idea to delegate at all, or would it be better to conduct the search on your own?

Consider the same model as in Section 2 *except* that now the principal can search on her own in lieu of contracting with an agent. If we assume that all search-related parameters are the same for both agent and principal, then a principal who conducts her own search will always be more efficient, thanks to the absence of agency frictions. We conclude that there must be some aspect(s) in terms of which the agent is more efficient at searching – that is, to compensate for those frictions – otherwise, it would never be a good idea to hire an agent. (This is proven in Proposition 2 to follow.)

Here we discuss three such aspects: the agent's search cost might be lower than the principal's, he could be faster at searching,¹² or he might have a better distribution of the alternatives. The agent's distribution of alternatives can be interpreted as him having access to certain *types* of opportunities to which the principal does not. For instance, an agent's knowledge of the language and the local stakeholders could gain him entry into a foreign market that is inaccessible to the principal.

These differences between the principal's and agent's search parameters capture the two reasons most often cited for outsourcing: a specialized firm can execute the task either better or at a lower cost. Denote the principal's and agent's search costs by c_P and c_A , and let p_P (resp., F_P) and p_A (resp., F_A) denote the probabilities of finding an alternative (resp., the distribution of alternatives). The supports of these distributions can also differ (supp $F_P \subset [\underline{u}_P, \overline{u}_P]$ and supp $F_A \subset [\underline{u}_A, \overline{u}_A]$).

Our next proposition identifies the conditions under which it is preferable to hire an agent. In the proposition, the requirement that the agent has access to a better distribution of the alternatives is quantified by the condition $p_A \bar{F}_A(x+e) \ge p_P \bar{F}_P(x)$. If the search speeds are the same $(p_A = p_P)$, then this condition is the conventional first-order stochastic dominance applied to the agent's distribution shifted to the left by e, which is a classic interpretation of the agent's distribution

¹² In search models, the probability of finding an alternative p is usually interpreted as search speed. That interpretation is perhaps more natural in search models which are continuous-time (where the search period's length approaches zero), in which p becomes the arrival rate of alternatives; see Zorc and Tsetlin (2020) for such a model.

being "strictly better." Yet, if the agent can search more rapidly than the principal $(p_A > p_P)$, then we have the weaker condition $\bar{F}_A(x+e) \ge \frac{p_P}{p_A} \bar{F}_P(x)$;¹³ here F_A can actually be worse than F_P in terms of first-order stochastic dominance (although not by an overly large margin). Thus the agent could be more effective at searching than the principal – even if he is accessing a worse distribution than she is – if he is sufficiently faster.

Proposition 2 (Optimal Delegation Decision) Let $p_A \bar{F}_A(x + e) \ge p_P \bar{F}_P(x)$ for all $x \ge 0$. Then the optimal decision of whether to delegate a search is given as follows.

(i) It is optimal for the principal to search herself if

$$p_P \int_{\underline{u}_P}^{\bar{u}_P} \bar{F}_P(x) \, dx - c_P \ge \max_{\xi \in \text{supp } F} p_A \int_{\xi}^{\bar{u}_A} x \, dF_A(x) - z(\xi). \tag{18}$$

(ii) It is optimal to delegate if $V^{\text{IN}} \leq V^{\text{OUT}}$, where

$$V^{\rm IN} = \frac{1}{1 - \alpha} \left(p_P \int_{\alpha V^{\rm IN}}^{\bar{u}_P} \bar{F}_P(x) \, dx - c_P \right) \quad \text{and}$$
$$V^{\rm OUT} = \frac{1}{1 - \alpha} \max_{\xi \in \text{supp } F} \left(p_A \int_{\max\{\xi, \alpha V^{\rm OUT}\}}^{\bar{u}_A} (x - \alpha V^{\rm OUT}) \, dF_A(x) - z(\xi) \right).$$

(iii) If neither (i) nor (ii) hold, then there exists an $N^* \in \mathbb{N}$ such that it is optimal to delegate if $N \leq N^*$ yet optimal not to delegate if $N > N^*$.

The first two parts of this proposition involve corner cases, where the decision of whether to hire the agent is clear. Part (i) demonstrates that the agent's search cost being smaller ($c_A < c_P$) is not a sufficient condition for delegating to be optimal. Not only must his search cost be lower than the principal's, it must also be low enough (a) that her complete hiring cost does not exceed her own search cost and (b) to compensate for agency frictions. However, this need for sufficiently low agent costs can be offset by the agent having access to a more promising distribution of alternatives.

Proposition 2(ii) concerns those situations in which it is always optimal to delegate. In such cases, V^{IN} and V^{OUT} are the limits of the principal's value functions—if she (respectively) conducts the search in-house and outsources — as N approaches infinity. These situations where $N \to \infty$ (ones where there is no pressure from impending deadline) are also the ones where the principal would like to be the most selective; however, (11) in Theorem 1 establishes that these are also the situations in which agency frictions are greatest. This result is intuitive because the principal's cost to contract with an agent increases with the quality threshold she wants him to apply, which

¹³ We could combine p and F into a single distribution G by defining G as a distribution that yields 0 (an alternative that cannot be brought to the principal) with probability 1 - p and otherwise yields a draw from distribution F with probability p. Given this definition of G, the condition $\overline{F}_A(x+e) \ge (p_P/p_A)\overline{F}_P(x)$ for all x is equivalent to $G_A(x+e)$ dominating $G_P(x)$ in terms of first-order stochastic dominance.

means that the optimal policy's cost increases with temporal distance from the deadline. So if the principal benefits from hiring an agent under these conditions, then she will also benefit from hiring an agent under any shorter deadline.

Part (iii) of the proposition states that when these two corner causes are excluded, the agent's hiring should depend on the length of the search horizon. The shorter the horizon, the more inclined the principal should be to hire an agent – with the decision to delegate always being optimal if there is but one period left. Underlying that claim is the greater expense of contracting with the agent when the principal wants him to search for a rare alternative. This connection is evident from (11), our expression for the optimal bonus, because it shows that the agent must be paid a larger risk premium when searching for rare alternatives; it is evident also from (15), according to which a search for rare alternatives entails more frequent (and hence more resources devoted to) evaluations by the principal if she is to avert the moral hazard problem.

In light of all these factors, conducting an in-house search is usually optimal when there is no time pressure. As for the search outcome's quality, the aphorism "if you want it done right, do it yourself" holds because the principal's sensitivity to quality leads to the optimality of in-house search. However, the same cannot be said of a principal who is – because of (say) a short deadline Nor a steep discount rate α – more concerned about time pressure than quality. Hence we can augment that aphorism as follows: "if you want it done *fast*, hire someone else to do it."

The condition $p_A \bar{F}_A(x+e) \ge p_P \bar{F}_P(x)$ in the proposition includes a horizontal shift from x to x + e. This shift is needed because there is no evaluation cost in the classical search model. An alternative way of modeling the same decision would be to model the in-house search as attaining the first-best outcome (under the principal's search parameters with the search cost increased by agent's wage w) and not as the canonical search model; then the evaluation cost is present in both in-house and delegated search, eliminating the need for an e-shift when comparing distributions. This model could be interpreted as an in-house search being conducted not by the principal herself but rather by an *in-house* agent whose actions are fully observable by the principal.

Remark 1. Proposition 2 extends to the variant case in which the principal's in-house search is modeled as attaining the first-best outcome (under her search parameters, but with the search cost increased by w). This extension follows because increasing the search cost by w and replacing the distribution F with G(x) := F(x+e) in the canonical value function (1) makes it equivalent to the first-best value function (13) in all cases where it is optimal to search.

Remark 2. If it is possible for the principal to "mix and match" – in other words, by searching in some periods herself and delegating the search in other periods – then cases (i) and (ii) of Proposition 2 are unaffected. In case (iii), however, there exists a time $T^* \in \{1, ..., N-1\}$ such that it is optimal for the principal to search on her own until T^* is reached and then to delegate the rest of the search. (This conclusion follows directly from the proof of Proposition 2, which derives this mix-and-match policy in Step 1 of the proof.)

4. Variants

4.1. Inability to Contract on Outcomes

The optimal contracts identified in Section 2 require that the value of the delivered alternative be contractible and thus verifiable by a third party. There are many business situations for which either this requirement is satisfied, or there is a good proxy for that value. For example, common contracts with real estate agents tie the agent's bonus to the final price of the purchased property, and bonuses for recruiters are similarly tied to the recruited candidate's salary. Yet, there are cases where contracting on the outcome is not feasible. For instance, most would view a design solution's *quality* as a subjective enough assessment that it could not be truly verifiable by a third party.

We explore the case of optimal contracting when the alternative's value is not contractible. Instead of X^n being the contractible variable, for this setting it is $Y^n := (Y_1, ..., Y_n)$, where

$$Y_n = \begin{cases} \emptyset & | \text{ If } no \text{ alternative was reported to the principal in period } n; \\ A & | \text{ If an alternative was reported to and accepted by the principal in period } n; \\ R & | \text{ If an alternative was reported to but rejected by the principal in period } n. \end{cases}$$

This type of delegated search is referred to as *unspecified search* (Lewis 2012). We therefore mark the associated policies with a superscript US to distinguish them from the policies of Section 2.

It might seem at first glance that the inability to contract on outcomes is not restrictive as the optimal payments prescribed in Theorem 1 do not depend on the value of the alternative. However, that value is still relevant (as a threshold) to the optimal contract $\{s_n^*(X^n)\}$: the agent is awarded a bonus for every delivered alternative with a value of at least ξ_n^* . Of course, such contracts cannot be implemented if the principal is unable to contract on the alternative's value. A major consequence is that the agent's actions will depend not only on the offered contract's stipulations (as in Theorem 1) but also on the principal's choice of acceptance policy. As a result, her ability to commit to an acceptance policy (which was inconsequential in Section 2) becomes relevant, as it effectively allows the principal to avoid her IC constraint. Next, we provide the solution to the variant where the principal is able to commit.

Theorem 4 (Unspecified Delegated Search with Commitment) Suppose that the principal announces and commits to an acceptance policy when offering the contract. Then the optimal contract consists of two parts, as follows. (i) In the first K time periods $(0 \le K < N)$, the principal pays the agent a fixed per-period wage s_n^0 as well as a bonus $s_n^B(\min \operatorname{supp} F)$ for every alternative that the agent delivers irrespective of whether (or not) the principal accepts it; here, s_n^0 and s_n^B are given by (11).

(ii) In the remaining N-K periods, the agent receives the same per-period wage, while the bonus of size $s_n^B(\psi_n^{\text{US}})$ is given whenever the principal accepts an alternative.

In either case, if we denote the principal's value function by V_n^{US} then the optimal acceptance policy $\{\psi_n^{\text{US}*}\}$ is the one that solves

$$V_{n}^{\rm US} = \max_{\psi_{n} \in \operatorname{supp} F} \left(p \int_{\psi_{n}}^{u} (x - \alpha V_{n+1}^{\rm US}) \, dF(x) - \min_{\xi \in \{\psi_{n}, \min \, \operatorname{supp} F\}} z(\xi) + \alpha V_{n+1}^{\rm US} \right), \tag{19}$$

with boundary condition $V_{N+1}^{\text{US}} = 0$. The optimal K is then given by $\max\{n \in \mathbb{N} \mid z(\psi_n^{\text{US}*}) \geq z(\min \operatorname{supp} F), n \leq N\}$ or by 0 if this set is empty. The optimal search policy is $\omega_n^{\text{US}*} = 1$, for all $n \in \{1, ..., N\}$. The optimal reporting policy $\{\xi_n^{\text{US}*}\}$ is

$$\xi_n^{\mathrm{US}*} = \begin{cases} \min \operatorname{supp} F & \text{if } n \le K, \\ \psi_n^{\mathrm{US}*} & \text{if } n > K. \end{cases}$$
(20)

This contract closely resembles the optimal contract when alternative values *are* contractible. The structure – which amounts to a fixed wage plus a bonus for delivery of some alternatives – remains, and our expression for the optimal payment size is unchanged. The optimality of threshold policies and having the agent search in every period extend to this setting as well.

We next examine how this contract and also its value to the principal differ from $\{s_n^*(X^n)\}$ (the optimal contract described in Section 2). The main difference is that a bonus being awarded now depends not on an alternative's value but rather on the principal's decision about accepting that alternative. Recall that, under $\{s_n^*(X^n)\}$, there is a time after which the agent presents the principal only with alternatives that she will be willing to accept. We use M to denote this time, and by Theorem 3, part (v), we have $M = \min\{n \in \mathbb{N} \mid \xi_n^* \geq \psi_n^*, n \leq N\}$. In Section 2, the principal would prefer that the agent bring her only those alternatives that are of at least slightly (less than e) higher quality than the minimum she is willing to accept. Enforcing that gap between the optimal reporting and acceptance policies is impossible without contracting on the value of alternatives, but the principal can still induce the agent to use the desired reporting policy by committing to an acceptance policy that equals the desired reporting policy. This possibility is illustrated in Figure 3, where the optimal induced reporting policy $\{\xi_n^{\text{US}*}\}$ coincides with the optimal acceptance policy $\{\psi_n^{\text{US}*}\}$ during the later part of the search.

From time M onward, the principal can induce the same actions at the same cost as when using the optimal contract of Section 2, *without* contracting on the value of the alternatives; the latter is possible via what is, in effect, the *purchase-right* contract described in Theorem 4(ii).



Figure 3 Optimal Policies for Unspecified Search (*x*-axis, time (period index); *y*-axis, value of alternatives)

When signing this contract, the principal commits to paying the agent a pre-determined bonus amount if, at any time from M onward, she accepts an alternative that he delivers. Thus, after M, there is no loss of efficiency if the principal is unable to contract on the outcomes; for $n \ge M$, we have $\xi_n^* = \xi_n^{\text{US}*} = \psi_n^{\text{US}*}$ and $V_n = V_n^{\text{US}}$. From a practical standpoint, this conclusion indicates that if the search deadline is near or if the principal is relatively less sensitive to the search outcome's quality,¹⁴ then there is no need for the contract to specify acceptable values of the search result. The principal's discretion to identify acceptable alternatives can be relied on without causing a conflict of interest between principal and agent.

However, that statement is not true if the horizon is far away or if the principal is sensitive to quality. Recall that under $\{s_n(X^n)\}$, before time M, the principal partially monitors the agent by requiring that he report any alternative of value at least ξ^{LC} , paying a bonus for each of them, but only accepting those with a value over ψ_n^* . That approach is not possible without the ability to contract on the value. Between times K and M, where such contracting is not an option, the principal will still want to use the same purchase-right contract that will be optimal later in the search (see Figure 3, where $\psi_n^{\text{US}*} = \xi_n^{\text{US}*}$ between K and M). However, this comes with a loss of efficiency compared to what is possible when the principal can contract on outcomes.

If the search horizon is extremely distant (i.e., prior to K), then the principal will opt for full monitoring of the agent, which she can achieve by offering a contract that requires him to report *all* found alternatives and awards a bonus for each delivery. Such monitoring prevents any moral hazard, but it does so at a high cost to the principal: this region is characterized by the greatest loss of efficiency as compared with the case where the principal can contract on value.

¹⁴ This could happen, for example, when the support of F_A is relatively narrow (i.e., when $\bar{u} - \underline{u}$ is small) – in which case the principal's payoff will not be much affected by which alternative was drawn from this distribution.

Theorem 4's main limitation is the difficulty of finding credible commitment mechanisms. Here, commitment can be made by public announcement of the principal's acceptance policy (which would impose a reputation loss should the principal deviate from the announced policy) or in situations involving repeated interaction between principal and agent (which would enable the latter to punish the former if she reneged on previously agreed-upon policies). Another commitment mechanism is hiring an intermediary to enforce the policy on which the principal and agent have agreed. As there exist situations where such mechanisms are not available, it would be useful to also consider what outcomes arise in the *absence* of an ability to commit.

However, in a setting of unspecified delegated search without commitment, we encounter methodological limitations. Plambeck and Zenios (2000) is designed for delegation of MDPs, such that the principal has no more decisions to make after signing the contract. This is not true in our setting, due to the principal's active decision making after being presented with an alternative. Nevertheless, we were able to show that Plambeck and Zenios (2000) apply to our setting so far. In Theorems 1-3, we were able to reduce the problem to a delegation of an MDP by showing that the optimal contract renders the principal's actions inconsequential to the agent, and the reduction to the delegation of an MDP was achieved in Theorem 4 through the power of commitment. Yet, such reduction is not possible in the setting of unspecified delegated search without commitment.

The source of complexity here is that in the absence of commitment, the choice of contract also directly determines the principal's acceptance policy. We can still think about this as delegation of MDPs (with principal effectively committed to policy determined by her IC constraint), but the issue is that the nature of the MDP being delegated changes based on the chosen incentive structure as the incentive structure determines the transition probabilities of the MDP.

4.2. Agent Chooses Search Intensity

One characteristic of our main model is that the agent can influence whether alternatives are generated (through his search policy ω_n), but he cannot influence the value of those alternatives through his decisions. This characteristic is the driving factor behind the bonus amount under optimal contracts being independent of the value of the delivered alternative (conditional on the alternative being of sufficiently high value). Yet, it is not hard to imagine situations where the agent *can* influence this value. E.g., a real estate agent who puts in more work into negotiating the price might find you the same selection of houses, but with a more favorable distribution of prices.

Consider an agent that decides not only on whether to search on not, but with which intensity to search. Thus, we assume $\omega_n \in \{0, i_L, i_H\}$, where i_H and i_L denote high and low intensity, respectively. Higher intensity will result in a better distribution of alternatives but also a higher search cost $(c_H > c_L)$. Assume $F(\cdot|i)$ is discrete uniform, but with support depending on the intensity: $\{0, 1, ..., \overline{u}_H\}$ if $i = i_H$ and $\{0, 1, ..., \overline{u}_L\}$ if $i = i_L$, where $\overline{u}_H, \overline{u}_L \in \mathbb{N}$ such that $\overline{u}_H > \overline{u}_L$. Here, the agent can influence the values through his actions, and the value of the delivered alternative provides information about the actions. Alternatives with a value $u > \overline{u}_L$ cannot be found unless the agent searches with high intensity; thus, those are completely informative. Alternatives with a value $u \leq \overline{u}_L$ are partially informative as they can be found with both low and high intensity but are more likely to be found under low intensity.

Here, we illustrate how the property that the optimal bonus size is independent of the value of the alternative can break under the assumption of multiple intensity levels that change the distribution. More detailed results for this variant are given in Appendix C.

Consider the situation where p = 1 and the principal would like the agent to search with high intensity but is willing to accept any outcome (so $\omega_n = i_H$, $\xi_n = 0$, $\psi_n = 0$). Here, inducing the agent to use non-zero intensity is trivial, as p = 1 implies that an outcome will always be generated under non-zero intensity. However, a flat fee for any outcome that the principal is willing to accept cannot provide an incentive to use high intensity because the probability that an outcome with a value of at least 0 is generated is 1, irrespective of the agent's intensity. Consequently, the only way to induce $\omega_n = i_H$, $\xi_n = 0$ is to make the value of the bonus depend on the value of the alternative delivered. Denoting $v(x) := -\exp[-r(1-\alpha)x]$ (this is the instantaneous utility function), and applying Proposition A3 (in Appendix C), the optimal way to accomplish this is by paying

$$s_n^L(i_H,0) = \frac{w + c_L}{\alpha},$$

if an alternative with a value no greater than \overline{u}_L is delivered, or

$$s_{n}^{H}(i_{H},0) = \frac{w}{\alpha} - \frac{1}{r(1-\alpha)} \ln\left(\frac{v(c_{L}/\alpha)(\overline{u}_{L}+1) - v(c_{H}/\alpha)(\overline{u}_{H}+1)}{(\overline{u}_{H} - \overline{u}_{L})}\right),^{15}$$

if an alternative of value more than \overline{u}_L is delivered. If nothing is delivered, the agent receives w/α .

4.3. Agent Reports the Value of Alternatives

One curious aspect of our model is that every time the agent finds an alternative, he is endowed with information about its value. This information could be valuable for the principal as it would allow her to make informed decisions on accepting or rejecting alternatives without the need to conduct an evaluation (thus avoiding the evaluation $\cot e$). Yet, there is no mechanism that allows the agent to transmit this information to the principal, as his reporting consists only of the binary decision to report the alternative (or no).

¹⁵ As in Theorem 1, the value inside the logarithm needs to be positive, otherwise $\omega_n = i_H$, $\xi_n = 0$ is not implementable.

Here we consider a variant that allows the agent to transmit the full information, or misrepresent it. In this case, the agent still decides whether to search in each period $\{\omega_n\}$, but the reporting policy takes a more complex form. In every period, the agent's reporting decision is a function $\mathcal{R}_n(x) : \operatorname{supp} F \cup \{\emptyset\} \to \operatorname{supp} F \cup \{\emptyset\}$, which maps each possible value found to value reported (here, \emptyset stands for no alternative found). Thus, the agent's reporting policy is $\{\mathcal{R}_n(\cdot)\}$. The principal then has the option to forego evaluation, making the call to either accept the alternative or continue the search based on the agent's report. Alternatively, the principal can pay e to reveal the alternative's value before making this decision. We allow the principal's evaluation frequency to also vary based on what the agent reported, so the principal's evaluation decisions can be described by functions $\phi_n(x) : \operatorname{supp} F \cup \{\emptyset\} \to [0, 1]$, which map the value of what agent reported to the probability of evaluating that report. We will refer to $\{\phi_n(\cdot)\}$ as the principal's *evaluation* policy.

The principal can also tie contractual payments to the value reported as well as the true value. Denote by $Z_n := (x_n, u_n, d_n)$ the public knowledge about search in period N. Here x_n is the reported value (or \emptyset if no report), u_n is the true value found via evaluation or NE if no evaluation was conducted, and $d_n \in \{A, R, \emptyset\}$ denotes whether principal accepted (A), rejected the alternative (R), or had no decision to make as there was no report (\emptyset). Then, the contractible variable here is $Z^n := \{Z_1, ..., Z_n\}$. Denoting $v(x) = -\exp[-r(1-\alpha)x]$, as in Section 4.2, we have the following result.

Proposition 3 Assume the principal has a costless mechanism available that induces the agent to be truthful. Denote the principal's value function under such mechanism by V_n^T . Thus defined V_n^T is an upper bound for the principal's value function under an optimal contract.

Let the implementability condition $p > 1 + v(c/\alpha)$ hold. The principal can get arbitrarily close to V_n^T in expectation by committing to conducting evaluations with frequency ϕ irrespective of the value reported (so $\phi_n(x) = \phi, \forall x, n$), and using a contract which consists of:

(i) Bonus when reporting that an alternative was found, and either no evaluation was conducted or evaluation found the report to be honest. The bonus size is given by

$$s_n^E := \frac{w+c}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{p}{1-(1-p)\exp[rc(1-\alpha)\alpha^{-1}]}.$$

(ii) Reduced (and potentially negative) pay when the evaluation was conducted and the report found to be dishonest, the value of this pay is given by

$$s_{n}^{P} := v^{-1} \left(\frac{1}{\phi} v(w/\alpha) - \frac{1-\phi}{\phi} v(s_{n}^{E}) \right).$$
(21)

(iii) Fixed pay $s_n^0 = w/\alpha$ if the agent reports that no alternative was found.

Denoting the value function of the principal using this contract by $V_n^{\text{ARV}}(x;\phi)$, we have that $\lim_{\phi\to 0} V_n^{\text{ARV}}(x;\phi) = V_n^T$. Thus the contract induces truthful revelation at an arbitrarily low cost to the principal and gives her an expected value equal to the one from the Section 2 model, where e = 0.

The intuition of Proposition 3 is that the principal has a simple mechanism available to induce honesty: she can randomly conduct evaluations (with frequency ϕ) and charge penalties to the agent if the reports were found to be dishonest. Because conducting evaluations is costly for the principal, she would prefer to do it as rarely as possible, which is not in conflict with the desire to induce truthfulness, as truthfulness can be induced for any non-zero ϕ by charging an appropriately sized penalty in case the agent is found to be dishonest, as given in part (ii) of Proposition 3.

This, in turn, allows the principal to avoid *all* inefficiency created by evaluation cost e, as she will only need to expand that cost with minuscule probability and is able to make decisions based on the agent's report in all other cases. This result is probably not realistic, as it is crucially dependent on the principal's ability to charge unboundedly high penalties for the agent if he is caught making a dishonest report. There are many refinements of this model possible that will eliminate this corner result by making truthful revelation costly to the principal.

Remark 3. For example, consider a setting where the principal cannot accept an alternative without evaluating it first. This is a relatively realistic assumption as, e.g., no corporate acquisition is finalized without due diligence (or no house purchase is finalized without conducting an inspection). Then, the contract of Proposition 3 will still incentivize the agent to be truthful (with identical proof), at no cost to the principal, but will no longer eliminate all system inefficiency due to evaluation cost e. Now, the principal will conduct an evaluation only once, before accepting the alternative, yielding an expected value to the principal equal to the one from the Section 2 model, but where there is no evaluation cost and F(x) is replaced by G(x) := F(x + e).

We can thus anticipate the properties that will hold for any refinement of this section's model that adds (either directly or indirectly) a cost to inducing truthfulness. First, the agent reporting the value of the alternative is at worst useless to the principal but never harmful. This is because the principal always has the option of ignoring any reports, evaluating all alternatives, and using the contract $\{s_n^*(X^n)\}$ given by Theorem 3, which will give her the same expected payout as in the Section 2 model. Second, the maximum possible benefit that the principal could attain through reporting is if she can induce truthful reporting without any cost, thus attaining value as if the evaluation cost did not exist (this is the case in Proposition 3). Any refinement should thus fall in between these two extremes, eliminating some but not all inefficiency from the evaluation cost e.

5. Conventional Contracts and Numerical Illustration

The structure of the optimal contracts we identify, which consists of a fixed per-period fee in every period and a bonus for successfully completing the search, does, in fact, correspond to some contracts observed in practice (e.g., contracts with consultants or lawyers). Moreover, if the agent does not forsake any other opportunities by signing a contract with the principal, as in part (iv) of Proposition 1, then the optimal contract still contains the bonus but omits a time fee. This type of contract is commonly observed with corporate recruiters (headhunters) and real estate agents. However, our proposed contract still diverges from those observed in practice in how the bonus is calculated and when it is awarded. In practice, a bonus tends to be either some fixed amount or simply a percentage of the final outcome.

Consider a company desiring to hire a seasoned executive. In practice, such hiring is often delegated to an outside headhunter (recruiter). So, let us examine in the context of our model, how would the optimal contract look like in this situation. While the parameter values (given in Table 1) are primarily chosen to illustrate the interesting behavior of the model, we also ensured that they are reasonable in the context of job search.

Table 1 Would parameters for the numeric indication		
Parameter	Value	Comments
Number of periods	N = 24	From 6-month recruiting cycle, and a period length of 1 week.
Probability of arrival	p = 0.6	Single-week probability of finding a candidate.
Risk aversion coefficient	r = 0.005 per \$	Consistent with a relatively high risk aversion amongst the
Discount factor	$\alpha = 0.999$	field estimates (Barseghyan et al. 2018). Weekly discount factor based on yearly discount factor .95.
Distribution of alternatives	$f(x) = \begin{cases} 1/\left(x\ln\left(\frac{\gamma}{\beta}\right)\right) & \beta \le x \le \gamma \\ 0 & \text{otherwis} \end{cases}$	γ Functional form used as the distribution of job market values for calibration of search models (Van den Berg 1990). Cal-
	x .	ibration using Financial Times EMBA job market salaries 2019, ¹⁶ adjusted for worker productivity (Kotlikoff and Gokhale 1992) yields parameters $\beta = \$60k, \gamma = \$1, 150k$. The resulting distribution has a mean of \$369,095 and a standard deviation of \$295,079. ¹⁷
Search costs	c = e = \$4k	Estimates from interviews with career services.
Opportunity cost	w = 0	Corresponds to the agent not forsaking other opportunities
		by taking the contract. Chosen to have a situation where the
		optimal contract differs from traditional ones only in how
		the bonus is determined and when it is awarded.

 Table 1
 Model parameters for the numeric illustration

Consider first what the bonus amounts are under the optimal policy - illustrated in the left panel of Figure 4. The bonus is constant up to period 14 and decreases afterward. The part where the bonus is constant are the periods for which $\xi_n^* = \xi^{\text{LC}}$, as earlier illustrated in Figure 2.

¹⁶ Total compensation around \$400k would be a typical compensation package for an EMBA holder from a top school, 3 years after graduation. "FT Global Executive MBA ranking 2019," *Financial Times*, 27th November, available online at http://rankings.ft.com/businessschoolrankings/executive-mba-ranking-2019.

¹⁷ In order to satisfy the assumptions of our model, we discretized this distribution using bracket-mean with bin probability $\mu = 10^{-6}$. We also obtained the numerical results with other discretization techniques (which do not satisfy the constant pmf assumption), as well as with the continuous distribution directly. In all of these cases, the results where indistinguishable as long as the discretization was sufficiently precise; this observation suggests that – consisten with the results of the continuous model in Appendix C.2 – the reliance on a specific discretization appears to be a purely technical limitation, which does not substantially affect the results.

Let us compare the performance of the optimal contract given in Theorem 3 to two contracts conventionally used in practice for headhunters. We can explore these suboptimal contracts in our framework by constraining the principal's choice of contracts to certain functional forms.

One conventional contact is a fixed bonus for successful completion of the search. Out of the history of the game X^n , this contract is contingent only on the value of the alternative reported in the current period, which we denote by x_n . Thus we can express this contract class as $\mathcal{F} := \{\{s_n(x_n)\}|s_n(x_n) = \mathbb{1}_{(x_n \in \Psi_n)} s, s \in \mathbb{R}^+\}$. The principal pays the agent a fixed amount s if the principal accepts the alternative brought by the agent.

Another conventional contract is a percentage of the value of the accepted alternative (e.g., percentage of the hire's salary). In our notation, this class of contracts is $\mathcal{P} := \{\{s_n(x_n)\}|s_n(x_n) = \mathbb{1}_{(x_n \in \Psi_n)} qx_n, q \in [0, 1]\}$, where q is the percentage of final value given as bonus.

Methodology for finding the optimal contract within \mathcal{F} and \mathcal{P} is given in Lemmas A4-A5 in Appendix B. Using the parameter values from Table 1 yields the optimal fixed bonus ~ \$219.9kand the optimal percentage bonus ~ 15.0%.

The comparison of the principal's value function under the optimal contract (given in Theorem 3) to these fixed-bonus and percentage-bonus contracts is given in the right panel of Figure 4. Using a percentage bonus contract instead of the optimal one results in the value function dropping by \$113,486 (13.9%) while fixed-bonus contracts perform even worse: a loss of \$197,010 (24.2%).



Figure 4 Left panel: optimal bonus size s_n^B . Right panel: principal's value functions under different contracts.

What is the reason behind this efficiency loss? Consider a fixed bonus contract paying \$10k. This bonus is more than sufficient to incentivize the agent to search in the last period. However, an agent with such contract is not incentivized to search in the period prior to the last – even though searching would be immediately profitable in the same period – because waiting for the last period and searching only then is even more profitable, due to the principal's decreased selectivity in the

last period. The only way to incentivize the agent to search in periods before the last is to make the bonus higher, but that compounds the problem of overpaying the agent in the last period.

In other words, if the bonus size is not decreasing over time, then the last few periods – in which the principal is not very selective – are more profitable for the agent. Thus, it becomes attractive for the agent not to search in early periods where it would be much less profitable and only start searching closer to the deadline. A curious consequence of this issue is that the maximum bonus that the optimal contract pays is \sim \$38.7k, yet the best fixed bonus contract pays \sim \$219.9k.

The percentage bonus contract partially resolves this issue, as the bonuses it pays for less valuable alternatives are smaller. However, percentage bonus contracts have an issue of their own: they impose an additional source of risk on the agent (the exact size of the bonus to be paid). This was not a serious issue in our chosen parameter set due to the relatively low risk aversion of the agent, but as risk aversion increases, the performance of percentage bonus contracts degrades.

One important property of the optimal contracts is that if it is optimal for the agent to search in one period, then it will be optimal in all of them. The same is not true for other contracts as contracts in \mathcal{F} and \mathcal{P} can fail to induce the agent to search. Indeed this can be seen in the right panel of Figure 4, where the agent under suboptimal contracts does not search in periods 2-5, leading to the flat value function for the principal there. Note, however, that the agent *does* search in the first period. This is driven by the incentive to search later (when it is more profitable) becoming weaker the longer the agent has to wait for the profitable situation (because of discounting).

It is worth noting that this inefficiency of conventional contract types is not driven by our assumptions about the agent's behavior:

Remark 4. Consider a multi-period problem (N > 1) with a risk-neutral agent. Then the contract of Theorem 3 will eliminate all agency frictions, achieving the first-best (shown in Proposition 1), however so will simple contracts such as the "sell the firm" contract (Bolton and Dewatripont 2005).¹⁸ Yet, contracts in \mathcal{F} and \mathcal{P} will not, as a necessary condition for the first-best to be achieved under these contracts is that the expected bonus that the agent is paid in the last period is equal to the search cost plus the value of the outside option, but such bonus will then be insufficient to incentivize the agent to search in all prior periods.

There are other dimensions in which optimal contracts differ from the traditional ones. First, the optimal contracts include a per-period component that compensates for the agent's opportunity cost. This is irrelevant for our example where w = 0, but becomes key if the opportunity cost w is high (e.g., increasing w to \$10,000 in our example will eliminate all value generated by conventional

¹⁸ This contract charges the agent a fixed amount in the beginning (enough to make the participation constraint binding) but then gives the agent all the value from successfully completing the search.

contracts but reduce the value from optimal contracts by only ~ 11%). Second, the optimal contract sometimes awards the bonus to the agent even for some of the alternatives that are not accepted (in this numerical example, this happens in periods 1-13), which reduces the amount of risk imposed on the agent, and becomes particularly important for high values of the risk aversion coefficient r, or when the alternatives that the principal is willing to accept can only be found with a small probability. This second difference can also be crucial, as there are situations where awarding bonuses for some alternatives that are not accepted is needed to gain any value from the contract (examples of such situations are given in the Proof of Proposition 4 in Appendix B). We now provide performance bounds for conventional contracts.

Proposition 4 (Efficiency of fixed and percentage bonus contracts) Let $V_n^{\mathcal{F}}$ (resp. $V_n^{\mathcal{P}}$) be the principal's value function when using the optimal contract within contract class \mathcal{F} (resp. \mathcal{P}), as given by Lemmas A4-A5 in Appendix B. Then, $V_n^{\mathcal{F}}$, $V_n^{\mathcal{P}} \leq V_n^{US}$, $\forall n \in \{1, ..., N\}$ (V_n^{US} is the principal's value function in the setting of Theorem 4). The efficiency gaps from the optimal contract, $V_n^{US} - V_n^{\mathcal{F}}$ and $V_n^{US} - V_n^{\mathcal{P}}$, can be unboundedly large.

The performance of the conventional contracts is bounded from above by what can be accomplished under unspecified search (US, the setting of Theorem 4). The optimal US contracts can be in the form of a bonus for successful completion of the search; it is just that this bonus is not stationary. However, in situations where N = 1 or $N \to \infty$, the optimal bonus under US *is* stationary; thus, in this situation, the fixed-bonus contract can be the same as the optimal US contract.

Other than this upper bound on the performance of conventional contracts, figuratively "anything can happen;" there is no lower bound on their performance, as the efficiency loss from use of suboptimal contracts can become unboundedly high, as stated in the last part of Proposition 4.

The existence of this efficiency gap suggests the possibility of improvement over existing contracts. One reason why we might not currently see the optimal contracts in practice is that they place a large knowledge burden on the principal to be able to set the decreasing bonus size in advance. Yet, we can expect this issue to diminish over time as the increasing availability of data helps the principal acquire the information needed.

For another possible reason, note that the discrepancy between conventional and optimal contracts is driven by periods in the proximity of the deadline. If there is no time pressure $(N \to \infty)$ and no opportunity cost (w = 0), the optimal delegation contracts in unspecified search can become exactly of the type currently used in practice: fixed bonus upon successful completion of the search. Thus, a plausible explanation for the prevalence of such contracts in practice is that time pressure from deadlines and the opportunity cost are negligibly small.

6. Discussion and Concluding Remarks

We first discuss the implications of our modeling approach, which should prove useful to scholars of dynamic contracting who must decide which methodology is most appropriate for their problems. We adopt the *dynamic principal-agent* (DPA) framework based on Plambeck and Zenios (2000); the other widely used method is the *promised future utility* (PFU) approach of Spear and Srivastava (1987). The commonality of these two approaches is that they eliminate the complexity of needing to deal with history-dependent contracts. PFU reduces history-dependence to but a single variable, while DPA eliminates it entirely; however, DPA does so by relying on a stronger set of assumptions.

One difference between the two approaches is the agent's limited liability, which is commonly present in PFU but would violate the assumptions of DPA. Here, we had the luxury of a nonbinding limited liability constraint: there are no negative payments in an optimal contract. Thus our solution remains valid even in the presence of this constraint. Moreover, this property will hold for any model in which an event that is profitable (from the principal's viewpoint) can occur only if the agent undertakes some costly action. (In our setting, as in the standard search model, an agent cannot find alternatives without paying the search cost.)

However, if there is a baseline rate at which such profitable events occur even without the agent's action (e.g., as in Sun and Tian 2018), then the optimal solution under a DPA approach will penalize the agent when the profitable event does not occur. If the same is modeled using PFU with limited liability, these penalties are not possible, and in lieu of the penalties, there is an initial period in which the agent does not receive compensation even for the profitable events – although their occurrence does reduce his time without compensation.

Our results are fairly robust to the approach used; the key difference is in how agents are penalized (direct monetary penalties in DPA vs. periods without pay in PFU). This robustness is perhaps best exemplified by the overlap between our solution in Section 2 and the solution of Lewis (2012), which occurs despite our substantially different search models and solution techniques.

When searching for a suitable alternative, such as a hire or a real estate property, a common conundrum is whether to search yourself or to hire an agent to search on your behalf. Seriously entertaining the latter requires that one consider how payments to the agent should be structured. The incentive structure affects not only the cost of hiring an agent but also the search outcome itself, as the decisions made during an agent's search process naturally depend on his incentives.

The optimal contracts we identify consist of a fixed wage (per period of time) and a bonus for delivering alternatives of sufficiently high value. This structure corresponds to conventional contracts encountered in the business world (e.g., those for recruiters, real estate agents, consultants, and lawyers). Thus our model offers a plausible rationale for why such contracts are so prevalent in practice, although our optimal contracts differ in how the bonus size is determined and when it is awarded. In our model, the size of the bonus is set when the contract is signed and decreases over time; it is also sometimes awarded even if the search is not successfully completed.

We show that the most common ways to determine those bonuses – namely, as a fixed value or as fraction of the accepted alternative's value – can be improved upon. In our numerical analysis, resorting to conventional contracts instead of the optimal ones results in a 14-24% loss of expectation for the principal. The reason why such bonus calculations are rare in practice may be that they place a large knowledge burden on the principal, who would need to know the exact distribution of possible values of the alternatives. Yet, as data become more accessible, it should become less difficult to implement these contracts, and so, we expect to see more of them. Another plausible reason for the scarcity of contracts in which the bonus decreases with time is that, when there is no deadline, this aspect of the optimal contract disappears (i.e., the bonus size is constant). Apparently, in practice, the time pressure due to deadlines seldom plays a leading role.

As for the question of whether to delegate search or not, the single greatest challenge related to that decision is the difficulty of inducing an agent to search for an alternative that is rare. The principal cannot easily distinguish between the agent not actually searching and him searching but without finding the desired alternative. This problem can be addressed only by the principal's costly monitoring in order to preclude moral hazard. It follows that a principal looking for rare, high-quality outcomes will prefer searching on her own – as will a principal with sufficient time available to explore different options. Conversely, a principal who is willing to accept a wide variety of alternatives or who faces a short deadline will prefer to delegate the search to an agent.

The principal's decisions concerning whether and how to outsource search are inherently complex, especially because conflicts of interest with the agent can easily arise. Our analysis sheds light on some aspects of this decision that have not been previously explored; however, our conclusions are not definitive because some of the complexities could not be incorporated into our framework. For instance, the business world is rife with complications, such as information asymmetries, that can arise at certain times and then continue to influence subsequent decisions – as when an agent is initially unclear about the target of his search but becomes more knowledgeable upon investigating the alternatives. There are also likely to be sources of agency frictions other than those that we consider; one example is not knowing the agent's attributes, so that the agent has private knowledge of his type. Also, there might be multiple agents that the principal could simultaneously hire, which would open up possibilities for different contracts (e.g., making the agents compete against each other). Addressing these topics would require an alternative modeling approach. In absence of those complexities, our analysis manages to identify which aspects of delegated search are the most consequential ones, and should help decision makers at large to decide when and how to delegate.

References

- Alizamir S, de Véricourt F, Sun P (2019) Search under accumulated pressure. *Operations Research* (forthcoming).
- Balseiro SR, Gurkan H, Sun P (2019) Multiagent mechanism design without money. *Operations Research* 67(5):1417–1436.
- Barseghyan L, Molinari F, O'Donoghue T, Teitelbaum JC (2018) Estimating risk preferences in the field. Journal of Economic Literature 56(2):501–64.
- Baucells M, Lippman SA (2004) Bargaining with search as an outside option: The impact of the buyer's future availability. *Decision Analysis* 1(4):235–249.
- Belloni A, Chen B, Sun P (2016) Computation of optimal dynamic mechanism with participation requirements. Duke working paper.
- Bolton P, Dewatripont M (2005) Contract theory (MIT Press).
- Chen M, Sun P, Xiao Y (2020) Optimal monitoring schedule in dynamic contracts. *Operations Research* 68(5):1285–1314.
- Erat S, Krishnan V (2012) Managing delegated search over design spaces. *Management Science* 58(3):606–623.
- Fudenberg D, Holmstrom B, Milgrom P (1990) Short-term contracts and long-term agency relationships. Journal of Economic Theory 51(1):1–31.
- Fuloria PC, Zenios SA (2001) Outcomes-adjusted reimbursement in a health-care delivery system. Management Science 47(6):735–751.
- Gerardi D, Maestri L (2012) A principal–agent model of sequential testing. *Theoretical Economics* 7(3):425–463.
- Gronau R (1971) Information and frictional unemployment. American Economic Review 61(3):290-301.
- Hart O, Holmström B (1987) The theory of contracts. Advances in economic theory: Fifth world congress, volume 1 (Cambridge University Press Cambridge).
- Hu Z, Tang W (2021) Size matters, so does duration: The interplay between offer size and offer deadline. Management Science 67(8):4643–5300.
- Kleinberg J, Kleinberg R (2018) Delegated search approximates efficient search. Proceedings of the 2018 ACM Conference on Economics and Computation, 287–302 (ACM).
- Kotlikoff LJ, Gokhale J (1992) Estimating a firm's age-productivity profile using the present value of workers' earnings. *Quarterly Journal of Economics* 107(4):1215–1242.
- Laffont JJ, Martimort D (2009) The theory of incentives: the principal-agent model (Princeton university press).

Lancaster T, Chesher A (1983) An econometric analysis of reservation wages. Econometrica 51(6):1661–76.

- Lewis TR (2012) A theory of delegated search for the best alternative. The RAND Journal of Economics 43(3):391-416.
- Li H, Zhang H, Fine CH (2013) Dynamic business share allocation in a supply chain with competing suppliers. Operations Research 61(2):280–297.
- Lippman SA, Mamer JW (2012) Exploding offers. Decision Analysis 9(1):6–21.
- Lippman SA, McCall J (1976) The economics of job search: A survey. *Economic Inquiry* 14(2):155–189.
- Mai Y, Pekeč D (2018) Information acquisition and technology adoption in a partnership. Duke Working Paper, SSRN 3221843.
- Malcomson JM, Spinnewyn F (1988) The multiperiod principal-agent problem. The Review of Economic Studies 55(3):391–407.
- Massala O, Tsetlin I (2015) Search before trade-offs are known. Decision Analysis 12(3):105–121.
- McCall JJ (1970) Economics of information and job search. The Quarterly Journal of Economics 84(1):113–26.
- McCardle KF (1985) Information acquisition and the adoption of new technology. *Management Science* 31(11):1372–1389.
- Merton RC (1990) Continuous-time finance (Blackwell Boston).
- Mortensen DT (1986) Job search and labor market analysis. Handbook of Labor Economics 2:849–919.
- Plambeck EL, Zenios SA (2000) Performance-based incentives in a dynamic principal-agent model. Manufacturing & Service Operations Management 2(3):240–263.
- Plambeck EL, Zenios SA (2003) Incentive efficient control of a make-to-stock production system. Operations Research 51(3):371–386.
- Poblete J, Spulber D (2017) Managing innovation: Optimal incentive contracts for delegated R&D with double moral hazard. *European Economic Review* 95:38–61.
- Radner R (1985) Repeated principal-agent games with discounting. *Econometrica* 1173–1198.
- Rahmani M, Ramachandran K (2021) Delegating innovation projects with deadline: Committed versus flexible stopping. *Management Science* 67(10):5969–6627.
- Rubinstein A, Yaari ME (1983) Repeated insurance contracts and moral hazard. *Journal of Economic Theory* 30(1):74–97.
- Samuel-Cahn E (1984) Comparison of threshold stop rules and maximum for independent nonnegative random variables. The Annals of Probability 12(4):1213–1216.
- Smith JE (1998) Evaluating income streams: A decision analysis approach. *Management Science* 44(12-part-1):1690–1708.

- Spear SE, Srivastava S (1987) On repeated moral hazard with discounting. *The Review of Economic Studies* 54(4):599–617.
- Sun P, Tian F (2018) Optimal contract to induce continued effort. Management Science 64(9):4193-4217.
- Tang W, Bearden JN, Tsetlin I (2009) Ultimatum deadlines. Management Science 55(8):1423-1437.
- Ulbricht R (2016) Optimal delegated search with adverse selection and moral hazard. Theoretical Economics 11(1):253-278.
- Ulu C, Smith JE (2009) Uncertainty, information acquisition, and technology adoption. *Operations Research* 57(3):740–752.
- Van den Berg GJ (1990) Nonstationarity in job search theory. The Review of Economic Studies 57(2):255–277.
- Weitzman ML (1979) Optimal search for the best alternative. Econometrica 47(3):641–54.
- Zhang H (2012a) Analysis of a dynamic adverse selection model with asymptotic efficiency. Mathematics of Operations Research 37(3):450–474.
- Zhang H (2012b) Solving an infinite horizon adverse selection model through finite policy graphs. Operations Research 60(4):850-864.
- Zhang H, Zenios SA (2008) A dynamic principal-agent model with hidden information: Sequential optimality through truthful state revelation. *Operations Research* 56(3):681–696.
- Zorc S, Tsetlin I (2020) Deadlines, offer timing, and the search for alternatives. *Operations Research* 68(3):927–948.

Electronic Companion: Appendices

Appendix A summarizes the principal notation. Appendix B contains proofs of the formal claims in the main paper. Appendix C provides supplementary content for the paper: Appendix C.1 gives further details on discretization techniques that satisfy our assumption of a constant pmf, Appendix C.2 contains an alternate version of our model where the distribution of alternatives is assumed to be continuous, and Appendix C.3 provides additional results for the variable search intensity variant of our model, as given in Section 4.2.

Appendix A: Notation

Table EC.1 N	otation
--------------	---------

Notation	Definition
Ν	Maximum number of periods in the model (search horizon)
T	Stochastic stopping time of the search process (acceptance time or N , whichever is first)
p	Probability of finding an alternative within a single period when searching
$F(\cdot)$	Distribution from which search alternatives are drawn
P(S)	Probability of drawing an outcome in set S when drawing from F
$\mu(\cdot)$	Probability mass function of $F(\cdot)$; with property that $\mu(x) = 1/ \operatorname{supp} F $, for all $x \in \operatorname{supp} F$
c	Search cost
e	Evaluation cost
α	Discount factor
$\{\bar{\xi}_n\}$	Canonical threshold search policy, equal to $\{\alpha V_n\}$, where V_n is given by (1)
$\{\omega_n\}$	Agent's search policy
$\{\Xi_n\}$	Agent's reporting policy
$\{\xi_n\}$	Agent's reporting policy in case it is a threshold policy $(\Xi_n = [\xi_n, \bar{u}] \cap \operatorname{supp} F, \forall n)$
$\{\gamma_n\}$	Agent's consumption policy
$\{\xi_n^{ m FB}\}$	First-best (threshold) reporting policy given by Lemma A2
$\{\Psi_n\}$	Principal's acceptance policy
$\{\psi_n\}$	Principal's acceptance policy in case it is a threshold policy $(\Psi_n = [\psi_n, \bar{u}] \cap \operatorname{supp} F, \forall n)$
X_n	Common knowledge about the search in period n
X^n	Common knowledge about the search process up to the end of period n, i.e., $(X_0,, X_n)$
	Rounding up to an element in supp F ($f' x \cong \min\{u \in \operatorname{supp} F u \ge x\}$)
$s = \{s_n(X^n)\}$	Agent's contract
$\{s_n^*(X^n)\}$	The optimal (second-best) contract, as given in Theorem 3
$V_n^P(\cdot) (V_n^A(\cdot))$	Principal's (agent's) value function in period n
r	Agent's risk aversion coefficient
W_n	Agent's wealth in period n , immediately following the contractual payments in that period
w	Agent's opportunity cost (per-period amount earned <i>absent</i> the principal)
$s_n^{\scriptscriptstyle D}(\Xi_n) \ (s_n^{\scriptscriptstyle D}(\xi_n))$	Single-period payment to the agent when an alternative of value in Ξ_n (at least ξ_n)
	is delivered to the principal under a Theorem 1-type contract that induces $\{\Xi_n\}$ $(\{\xi_n\})$
$z(\xi_n)$	Expected single-period principal costs when inducing the agent to use ξ_n (Lemma 1)
$\xi^{\mu\nu}$	Value of ξ_n that minimizes the principal's expected single-period costs $z(\xi_n)$ (Lemma 1)
$\{\xi_n^*\}\;(\{\psi_n^*\})$	Optimal threshold reporting (acceptance) policy as given in Theorem 3

Appendix B: Proofs of Formal Statements

The following lemma, stated here in the appendix, establishes a property of search policies that will be used in several of the proofs that follow for results from the main paper. **Lemma A1.** In the canonical model, if there exists a period $n \in \{1, ..., N\}$ in which it is optimal to pay the search cost, then it is optimal to pay the search cost in every period.

Proof. Because $\{V_n\}$ is positive (from the boundary condition) and decreasing in n, it follows that $p \int_{\alpha V_{n+1}}^{\bar{u}} \bar{F}(x) dx$ is increasing in n. So if there exists a $k \in \{1, ..., N\}$ such that $p \int_{\alpha V_{n+1}}^{\bar{u}} \bar{F}(x) dx - c \ge 0$, then $p \int_{\alpha V_{n+1}}^{\bar{u}} \bar{F}(x) dx - c \ge 0$ for all $h \in \{k, ..., N\}$. Therefore, the optimality of paying the search cost in period k implies the optimality of paying it in every subsequent period.

We demonstrate that the same holds for preceding periods by contradiction. Assume there exists a period m such that (a) it is optimal for the agent not to pay the search cost in period m but (b) it is optimal to pay that cost in period m + 1. Then, under an optimal policy, this agent's NPV in period m is αV_{n+1} . Yet, it is possible to use the same search strategy *as if* it were currently period m + 1; doing so would yield a higher value (V_{n+1}) in expectation – contradicting the optimality of not searching in period m.

Proof of Theorem 1.

Assume (temporarily) that the principal commits to an acceptance policy $\{\Psi_n\}$ when offering the contract to the agent. Then, the principal has no decisions to make once the contract is accepted, which reduces this problem to delegation of control over a Markov decision process. It follows from Plambeck and Zenios (2000, Thm. 1) that the problem of finding a contract that implements any policy pair ($\{\omega_n\}, \{\Xi_n\}$) at the lowest possible cost to the principal can be decomposed into N single-period optimization problems (for every $n \in \{1, ..., N\}$), where each problem is given by

$$\underset{s_n^A(u),s_n^R(u),s_n^0}{\operatorname{arg\,min}} \quad p\omega_n \sum_{u \in \Xi_n \cap \Psi_n} s_n^A(u)\mu(u) + p\omega_n \sum_{u \in \Xi_n \setminus \Psi_n} s_n^R(u)\mu(u) + \left(1 - p\omega_n \sum_{u \in \Xi_n} \mu(u)\right) s_n^0; \tag{EC.1}$$

$$-p\omega_{n}\sum_{u\in\Xi_{n}\cap\Psi_{n}}\exp[-r(1-\alpha)(s_{n}^{A}(u)-\alpha^{-1}c)]\mu(u) - p\omega_{n}\sum_{u\in\Xi_{n}\setminus\Psi_{n}}\exp[-r(1-\alpha)(s_{n}^{R}(u)-\alpha^{-1}c)]\mu(u)$$
(EC.2)

$$-\left(1 - p\omega_{n}\sum_{u\in\Xi_{n}}\mu(u)\right)\exp[-r(1-\alpha)(s_{n}^{\circ} - \alpha^{-1}c)] = -\exp[-r\alpha^{-1}(1-\alpha)w];$$

$$-p\sum_{u\in\Xi\cap\Psi_{n}}\exp[-r(1-\alpha)(s_{n}^{A}(u) - \alpha^{-1}c)]\mu(u) - p\sum_{u\in\Xi\setminus\Psi_{n}}\exp[-r(1-\alpha)(s_{n}^{R}(u) - \alpha^{-1}c)]\mu(u)$$

$$-\left(1 - p\sum_{u\in\Xi}\mu(u)\right)\exp[-r(1-\alpha)(s_{n}^{0} - \alpha^{-1}c)] \le -\exp[-r\alpha^{-1}(1-\alpha)w] \quad \forall\Xi\in\mathcal{F};$$

$$-\exp[-r(1-\alpha)s_{n}^{0}] \le -\exp[-r\alpha^{-1}(1-\alpha)w]. \quad (EC.4)$$

In these expressions, $s_n^A(u)$ is the agent's reward for bringing an alternative of value u to the principal when the principal accepts that alternative; his reward is $s_n^R(u)$ when delivering an alternative of value uthat the principal rejects. Finally, s_n^0 is the agent's reward when he does not present the principal with any alternatives. The objective function (EC.1) minimizes the expected sum of these rewards conditional on the agent using the desired policy Ξ_n . Equation (EC.2) is the (binding) participation constraint. The incentive compatibility constraint (EC.3) ensures that the agent can not do better by searching and using any other reporting policy, while (EC.4) ensures that he can not do better by not searching at all (notice that the reporting policy is redundant if not searching). The recurring $-\exp[-r(1-\alpha)x]$ expression in the system (EC.1)-(EC.4) is the amount by which the agent's expected wealth-independent utility component changes (under an optimal consumption policy) if the agent receives payment x (this also follows from Theorem 1 of Plambeck and Zenios 2000).

If the principal would like to implement $\omega_n = 0$ – i.e., induce the agent to *not* search – a simple solution of the problem (EC.1)–(EC.4) is $s_n^R(u) = s_n^A(u) = s_n^0 = w/\alpha, \forall u$.

We now turn to the case where $\omega_n = 1$. We can make several observations that will allow us to simplify the system of equations (EC.1)–(EC.4).

First of all: to solve (EC.1)–(EC.4) while ignoring the incentive compatibility constraint (EC.3), we will show that we can restrict our attention to $s_n^A(u)$ and $s_n^R(u)$, which are step functions that pay a fixed value s_n^A (resp. s_n^R) if the reported alternative is in Ξ_n and pay 0 otherwise. To do that, let $s_n^{A*}(u)$, $s_n^{R*}(u)$, and s_n^{0*} be solutions of the system (EC.1),(EC.2),(EC.4) and denote $\varphi(x) := -\exp -r(1-\alpha)(x-c/\alpha)$. We can define $s_n^{A**}(u)$ as $s_n^{A**}(u) := s_n^{A**}$ if $u \in \Xi_n$ and $s_n^{A**}(u) := 0$ otherwise, where the constant s_n^{A**} is the solution of $\sum_{u \in \Xi_n \cap \Psi_n} \varphi(s_n^{A**})\mu(u) = \sum_{u \in \Xi_n \cap \Psi_n} \varphi(s_n^{A*}(u))\mu(u)$. In other words, $s_n^{A**}(u)$ is constructed so that it is a step function of the desired form and so that the agent's expected payoff when using Ξ_n is the same under $s_n^{A*}(u)$ and $s_n^{A**}(u)$. Then, noticing that $\varphi(\cdot)$ is an increasing and concave function and using Jensen's inequality yields $\varphi(\mathbb{E}_u[s_n^{A**}(u)|u \in \Xi_n \cap \Psi_n]) = \phi(s_n^{A**}) = \mathbb{E}_u[\varphi(s_n^{A**}(u))|u \in \Xi_n \cap \Psi_n] = \mathbb{E}_u[\varphi(s_n^{A*}(u))|u \in \Xi_n \cap \Psi_n]$. Observing that the objective function in (EC.1) depends on the choice of $s_n^A(u)$ only through its first term which is equal to $pP(\Xi_n \cap \Psi_n)\mathbb{E}_u[s_n^A(u), s_n^{O*})$ follows, showing that we can restrict attention to $s_n^A(u)$, $s_n^{O*}(u), s_n^{O*})$ compared to $(s_n^{A*}(u), s_n^{R*}(u), s_n^{O*})$ follows, showing that we can restrict attention to $s_n^A(u)$ functions of the desired form without loss of optimality. The same property for $s_n^R(u)$ follows analogously.

The intuition of the result in the above paragraph is straightforward. By choosing policy Ξ_n , the agent ensures that all alternatives in Ξ_n that are found are brought to the principal. Yet, the agent has no control over the exact value of these alternatives, and from this it follows that basing his payment on such exact values – and not simply on whether an alternative's value is in Ξ_n – serves only as an additional source of risk for the agent, which in turn reduces his expected utility due to the agent's risk aversion.

Returning to the problem given by the system (EC.1),(EC.2),(EC.4): the participation constraint is binding, and so, the agent is indifferent between the principal accepting the alternative (and thereby ending the search process) or rejecting it (thus continuing the search). By the same Jensen's inequality argument as before, we also have $s_n^A = s_n^R$. With this in mind, a relaxed version of the problem (EC.1)-(EC.4) that drops constraint (EC.3) is the solution of

$$\underset{s_{A}^{A}, s_{D}^{0}}{\arg\min pP(\Xi_{n})s_{n}^{A} + (1 - pP(\Xi_{n}))s_{n}^{0}};$$
(EC.5)

$$-pP(\Xi_n)\exp[-r(1-\alpha)(s_n^A - \alpha^{-1}c)] - (1-pP(\Xi_n))\exp[-r(1-\alpha)(s_n^0 - \alpha^{-1}c)] = -\exp[-r\alpha^{-1}(1-\alpha)w];$$

(EC.6)

$$-\exp[-r(1-\alpha)s_n^0] \le -\exp[-r\alpha^{-1}(1-\alpha)w].$$
(EC.7)

We proceed to solve this relaxed version of the problem, after which we will demonstrate that the solution does not violate (EC.3) and is thus also a solution to the original problem (EC.1)-(EC.4). Substituting

 $u_n^A := \exp\left[-r(1-\alpha)s_n^A\right], u_n^0 := \exp\left[-r(1-\alpha)s_n^0\right], k := \exp\left[r(1-\alpha)\alpha^{-1}c\right], o := \exp\left[-r\alpha^{-1}(1-\alpha)w\right)\right], \text{ the problem (EC.5)-(EC.7) simplifies to}$

$$\underset{u_{n}^{A}, u_{n}^{0}}{\arg\min(r(1-\alpha))^{-1} \left(-(1-pP(\Xi_{n}))\ln u_{n}^{0} - pP(\Xi_{n})\ln u_{n}^{A} \right)};$$
(EC.8)

$$-(1 - pP(\Xi_n))u_n^0 k - pP(\Xi_n)u_n^A k + o = 0;$$
(EC.9)

$$-u_n^0 + o \le 0. \tag{EC.10}$$

Here, the objective function is convex and the constraints are affine, so the Karush–Kuhn–Tucker (KKT) conditions are both necessary and sufficient. From KKT conditions, the solution only exists if $1 - k(1 - pP(\Xi_n)) > 0$ (this is the implementability condition (10)), and is given by $u_n^0 = o, u_n^A = o(1 - k(1 - pP(\Xi_n)))/(kpP(\Xi_n))$. Reverting the substitution then gives

$$s_n^A = \frac{w+c}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{pP(\Xi_n)}{1 - (1 - pP(\Xi_n)) \exp\left[rc(1-\alpha)\alpha^{-1}\right]}, \qquad s_n^0 = \frac{w}{\alpha},$$
 (EC.11)

as the solution. We find it useful to resort to slightly different notation in the main text compared this this proof, writing $s_n^B(\Xi_n)$ rather than s_n^A in (11) to emphasize its dependence on Ξ_n as well as the fact that whether the bonus is awarded does not depend on whether the alternative is accepted (thus the supersript A might be misleading). We still need to verify if this solution adheres to the dropped constraint. Plugging the solution of (EC.5)-(EC.7) into the constraint (EC.3) and expanding its RHS yields

$$pP(\Xi_n \cap \Xi)s_n^A + (1 - pP(\Xi_n \cap \Xi))s_n^0 \le pP(\Xi_n)s_n^A + (1 - pP(\Xi_n))s_n^0, \forall \Xi \in \mathcal{F},$$

which holds because it follows from (EC.11) that $s_n^A \ge s_n^0$. Thus, this is the solution to the original problem (EC.1)-(EC.4) as well. That this contract is renegotiation proof follows from Plambeck and Zenios (2000, Cor. 1).

Then, observe from (EC.11) that the solution depends neither on the acceptance policy nor on the agent's knowledge of that policy. Furthermore, the solution does not depend on the principal adhering to the announced policy; the reason is that, because the single-period participation constraint is binding, the agent is always indifferent between the principal accepting or rejecting the alternative he delivers.¹⁹

It remains to find the optimal consumption policy under contracts given by this theorem. Thus, let the agent hold the contract given by this theorem which implements policies $\{\omega_n\}$, $\{\Xi_n\}$, and let the principal use acceptance policy $\{\Psi_n\}$. From (5) it follows that the agent's optimal consumption policy solves

$$V_{n}^{A}(W_{n}) = \max_{\gamma_{n}} - \exp(-r\gamma_{n}) + \omega_{n}p \int_{\Xi_{n} \cap \Psi_{n}} -\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_{n} - c - \gamma_{n}}{\alpha} + s_{n}^{B}(\Xi_{n})\right)\right)\right] dF(u) + \omega_{n}p \int_{\Xi_{n} \setminus \Psi_{n}} \alpha V_{n+1}^{A}\left(\frac{W_{n} - \gamma_{n} - c}{\alpha} + s_{n}^{B}(\Xi_{n})\right) dF(u) + (1-\omega_{n}pP(\Xi_{n}))\alpha V_{n+1}^{A}\left(\frac{W_{n} - \gamma_{n} - \omega_{n}c}{\alpha} + s_{n}^{0}\right),$$
(EC.12)

¹⁹ The optimal contract identified here is not unique. In particular, the contract is not sensitive to the payment linked to bringing the principal an undesirable alternative $(s_n^A(u) \text{ and } s_n^R(u) \text{ for } u \notin \Xi_n)$. We arbitrarily set this value to 0 in the proof so as to simplify the expressions, but any payment that does not exceed w/α can accomplish the same result.

with boundary condition $V_{N+1}^A(W_{N+1}) = -\frac{1}{1-\alpha} \exp[-r(w+(1-\alpha)W_{N+1})]$. We will show – by backwards induction – that in any period, the optimal consumption is given by $\gamma_n = W_n(1-\alpha) + w$ and that the agent's value function is $V_n^A(W_n) = -\frac{1}{1-\alpha} \exp[-r(w+(1-\alpha)W_n)]$. Assume there is a period n+1 such that $V_{n+1}^A(W_{n+1}) = -\frac{1}{1-\alpha} \exp[-r(w+(1-\alpha)W_{n+1})]$. For the induction base, this property holds for n = N, directly from the boundary condition.

For the induction step, in period n using $V_{n+1}^A(W_{n+1}) = -\frac{1}{1-\alpha} \exp[-r(w+(1-\alpha)W_{n+1})]$ yields

$$V_n^A(W_n) = \max_{\gamma_n} - \exp(-r\gamma_n) + \omega_n p \int_{\Xi_n} -\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - c - \gamma_n}{\alpha} + s_n^B(\Xi_n)\right)\right)\right] dF(u) - \left(1 - \omega_n p P(\Xi_n)\right) \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - \omega_n c - \gamma_n}{\alpha} + s_n^0\right)\right)\right].$$
(EC.13)

If $\omega_n = 0$, this reduces to $V_{n+1}^A(W_{n+1}) = \max_{\gamma_n} - \exp(-r\gamma_n) - \frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n + w - \gamma_n}{\alpha}\right)\right)\right]$. Here the objective is concave and solving for its FOC yields $\gamma_n = w + (1-\alpha)W_n$ as the sole maximizer and $V_n^A(W_n) = -\frac{1}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)W_n\right)\right]$ as the maximum. For the other case, if $\omega_n = 1$, (EC.13) can be decomposed as

$$V_n^A(W_n) = \max_{\gamma_n} - \exp(-r\gamma_n) - \frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\frac{W_n - \gamma_n}{\alpha}\right)\right] \\ \times \left(pP(\Xi_n)\exp\left[-r(1-\alpha)(s_n^B(\Xi_n) - c/\alpha)\right] + (1-pP(\Xi_n))\exp\left[-r(1-\alpha)\left(s_n^0 - c/\alpha\right)\right]\right).$$

Using (EC.6), the second line of this expression is equal to $\exp[-r\alpha^{-1}(1-\alpha)w]$, which reduces the problem to $V_{n+1}^A(W_{n+1}) = \max_{\gamma_n} - \exp(-r\gamma_n) - \frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n + w - \gamma_n}{\alpha}\right)\right)\right]$, which is the same as in the $\omega_n = 0$ case. Thus, $\gamma_n = w + (1-\alpha)W_n$ and $V_n^A(W_n) = -\frac{1}{1-\alpha}\exp\left[-r\left(w + (1-\alpha)W_n\right)\right]$ hold in both cases, completing the induction. \Box

Proof of Theorem 2 and Lemma 1.

We prove these two results concurrently.

Part 1: the Cost Function. This part is an intermediate step that is needed for the other parts of the proof. Summing the different costs in (12), and using $s_n^0 = w/\alpha$ (shown in Theorem 1), the principal's expected single-period costs to induce the agent to search and use reporting policy Ξ_n are:

$$\zeta(\Xi_n) = w + pP(\Xi_n)(\alpha s_n^B(\Xi_n) + e - w).$$
(EC.14)

Applying (11) to expand $s_n^B(\Xi_n)$ then yields

$$\zeta(\Xi_n) = w + pP(\Xi_n) \left(c + \frac{\alpha}{r(1-\alpha)} \ln \frac{pP(\Xi_n)}{1 - (1-pP(\Xi_n)) \exp\left[rc(1-\alpha)\alpha^{-1}\right]} + e \right).$$
(EC.15)

As $\zeta(\Xi_n)$ depends on Ξ_n only through $P(\Xi_n)$, for any two sets $\Xi_n, \Xi'_n \in \mathcal{F}$ such that $P(\Xi_n) = P(\Xi'_n)$ it also holds that $\zeta(\Xi_n) = \zeta(\Xi'_n)$.

Part 2: Optimality of History-Independent Policies. This part establishes the first sentence of Theorem 2. Applying Theorem 1, the contracting problem (6)-(9) reduces to an MDP, after which from (12) we have that the principal's value function in any non-terminal state is given by:

$$V_n^P = \max_{\omega_n \in \{0,1\}, \Xi_n, \Psi_n \in \mathcal{F}} \omega_n p P(\Xi_n) \left(\alpha s_n^0 - e - \alpha s_n^B(\Xi_n) \right) + \omega_n p \int_{\Xi_n \cap \Psi_n} (u - \alpha V_{n+1}^P) dF(u) - \alpha s_n^0 + \alpha V_{n+1}^P,$$

with boundary condition $V_{N+1}^P = 0$. Here, the objective function is independent of the history and thus the optimal choices of ω_n, Ξ_n, Ψ_n need not depend on it either. Note that the property that the optimal policy of an MDP needs to depend only on the current state and not on the full history is universal to most dynamic programs (see the adequacy of Markov policies in ch. 1.1.3 of Bertsekas 2012); however, as in the canonical search problem with no recall, the optimal policy here can also be independent of the current state

Part 3: Optimality of Threshold Policies. This part establishes the second sentence of Theorem 2. If $\omega_n = 0$, the choice of Ξ_n and Ψ_n is redundant, so consider the case where $\omega_n = 1$. From (12), the choice of Ψ_n affects the principal's value function only via the $\sum_{u \in \Xi_n \cap \Psi_n} (u - \alpha V_{n+1}^P) \mu(u)$ term, thus the threshold acceptance policy $\Psi_n^* = [\alpha V_{n+1}^P, \bar{u}] \cap \text{supp } F$ is optimal (as this sum is maximized by including all points for which the summands are positive).²⁰

Assume, by contradiction, there exists an optimal Ξ_n^* that is not a threshold one, denote $k := |\Xi_n^*| (|\cdot|)$ is the cardinality of a set), and define Ξ_n^{**} as the set of k largest elements of supp F. We then have $P(\Xi_n^*) = P(\Xi_n^{**})$, as $|\Xi_n^*| = |\Xi_n^{**}|$ and $\mu(u)$ is constant. Thus, part 1 yields $\zeta(\Xi_n^*) = \zeta(\Xi_n^{**})$. Denote by $V_n^P(\omega_n, \Xi_n, \Psi_n)$ the principal's expected value when using $(\omega_n, \Xi_n, \Psi_n)$ in period n and proceeding optimally. From the objective function in (12) we then have $V_n^P(1, \Xi_n^{**}, \Psi_n^*) - V_n^P(1, \Xi_n^*, \Psi_n^*) = p\mu \sum_{u \in \Xi_n^* \cap \Psi_n^*} (u - \alpha V_{n+1}^P) - p\mu \sum_{u \in \Xi_n^* \cap \Psi_n^*} (u - \alpha V_{n+1}^P) \ge 0$. (Here, the last inequality follows from noticing that the first sum has at least as many summands as the second, all of which are positive, and that the *i*-th biggest element of the first sum is greater than or equal to the *i*-th biggest element of the second.) Thus, either both Ξ_n^* and Ξ_n^{**} are optimal, or we have a contradiction to optimality of Ξ_n^* . Either way, the optimality of threshold reporting policies follows.

Part 4: the Cost Function for Threshold Policies. This part establishes the statement of Lemma 1. Defining $z(\xi_n) := \zeta([\xi_n, \bar{u}] \cap \text{supp } F)$ yields (14). Because $\zeta(\Xi_n)$ depends on Ξ_n only through the probability of finding an alternative in Ξ_n , we can also consider a cost function $\Pi(\pi)$ defined directly on these probabilities, which is then given by

$$\Pi(\pi) = w + \pi \left(c + \frac{\alpha}{r(1-\alpha)} \ln \frac{\pi}{1 - (1-\pi) \exp\left[rc(1-\alpha)\alpha^{-1}\right]} + e \right).$$

From (14), $\Pi(p\bar{F}(\xi_n)) = z(\xi_n)$ for all ξ_n that satisfy the implementability condition of Theorem 1. Consider $\Pi(\pi)$ on the domain $(\tau, p]$ where $\tau := 1 - 1/\exp[rc(1-\alpha)/\alpha]$ (this is the implementability threshold). We then have

$$\frac{d^2\Pi(\pi)}{d\pi^2} = \frac{\alpha \left(e^{\left(\frac{1}{\alpha}-1\right)cr}-1\right)^2}{(1-\alpha)r\pi \left((\pi-1)e^{\left(\frac{1}{\alpha}-1\right)cr}+1\right)^2} > 0,$$

thus $\Pi(\pi)$ is strictly convex. Consequently, as $\bar{F}(\xi_n)$ is weakly decreasing, $z(\xi_n) = \Pi(p\bar{F}(\xi_n))$ is quasi-convex. Because $\bar{F}|_{\text{supp }F}(\xi_n)$ is strictly decreasing, it follows from strict convexity of $\Pi(\pi)$ that $z|_{\text{supp }F}(\xi_n)$ is strictly increasing on $\xi_n \leq \xi^{\text{LC}}$ and strictly decreasing on $\xi_n > \xi^{\text{LC}}$. \Box

The following lemma, which characterizes the first-best policy, will be used in several of the proofs.

²⁰ Note that because summation here is taken over $\Xi_n \cap \Psi_n$, Ψ_n^* is not *uniquely* optimal. Every Ψ_n' such that that $\Xi_n \cap \Psi_n^* = \Xi_n \cap \Psi_n'$ will also be optimal. However, Ψ_n^* is the only policy which is optimal irrespective of which Ξ_n the agent uses, which is why we focus on it.

Lemma A2. The principal's first-best value function is given by

$$V_{n}^{\rm FB} = \left(p \int_{\alpha V_{n+1}^{\rm FB} + e}^{\bar{u}} \bar{F}(u) du - c \right)^{+} - w + \alpha V_{n+1}^{\rm FB}, \tag{EC.16}$$

with boundary condition $V_{N+1}^{\text{FB}} = 0$. The first-best reporting and acceptance policies are $\{\xi_n^{\text{FB}}\} = \{\psi_n^{\text{FB}}\} = \{\uparrow \alpha V_{n+1}^{\text{FB}} + e \uparrow\}$. For all $n \in \{1, ..., N\}$, the first-best search policy is

$$\omega_n^{\rm FB} = \begin{cases} 1 & \text{if } p \int_{\alpha V_{n+1}^{\rm FB} + e}^{\bar{u}} \bar{F}(x) \, dx - c > 0; \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Noting that under the FB, the agent's costs are just passed on to the principal and the agent receives no additional compensation, from (4) we have that the FB policies solve the Bellman equation

$$V_{n}^{\text{FB}} = \max_{\omega_{n} \in \{0,1\}, \xi_{n}, \psi_{n} \in \text{supp}} \omega_{F} \left(p \int_{\max\{\xi_{n}, \psi_{n}\}}^{\bar{u}} (u - e - \alpha V_{n+1}^{\text{FB}}) dF(u) - peP([\xi_{n}, \psi_{n})) - c \right) - w + \alpha V_{n+1}^{\text{FB}}, \quad (\text{EC.17})$$

with boundary condition $V_{N+1}^{\text{FB}} = 0$. The choice of ξ_n, ψ_n affects the objective function of (EC.17) in two ways: a) through the non-positive $-peP([\xi_n, \psi_n))$ term (which is then maximized by any $\xi_n \geq \psi_n$, and b) the integral term $\int_{\max\{\xi_n, \psi_n\}}^{\bar{u}} (u - e - \alpha V_{n+1}^{\text{FB}}) dF(u)$, which is maximized by $\max\{\xi_n, \psi_n\} = \alpha V_{n+1}^{\text{FB}} + e$ (thus also for $\max\{\xi_n, \psi_n\} = \vec{r} \alpha V_{n+1}^{\text{FB}} + e \uparrow$, as F(u) has no probability mass in $(\alpha V_{n+1}^{\text{FB}} + e, \vec{r} \alpha V_{n+1}^{\text{FB}} + e \uparrow))$. Consequently, setting $\xi_n^{\text{FB}} = \psi_n^{\text{FB}} = \vec{r} \alpha V_{n+1}^{\text{FB}} + e \uparrow$ is optimal. Statement of the lemma then follows from using the identity $\int_{\vec{r} \alpha V^{\text{FB}} + e^{\uparrow}} (u - e - \alpha V_{n+1}^{\text{FB}}) dF(u) = \int_{\alpha V^{\text{FB}} + e}^{\bar{u}} (u - e - \alpha V_{n+1}^{\text{FB}}) dF(u) = \int_{\alpha V^{\text{FB}} + e}^{\bar{u}} F(u) du$ to shorten the integral in (EC.17) and noticing that $\omega_n = 1$ is optimal if and only if the expression in large brackets in (EC.17) is positive. \Box

Proof of Theorem 3.

As shown in part 2 of the proof for Theorem 2, $\psi_n^* = r^* \alpha V_n^P$ is an optimal acceptance policy, inserting which into (16) yields

$$V_{n}^{P} = \max_{\omega_{n} \in \{0,1\}, \xi_{n} \in \text{supp } F} \omega_{n} \left(p \int_{\max\{\xi_{n}, \alpha V_{n+1}^{P}\}}^{\bar{u}} (x - \alpha V_{n+1}^{P}) dF(x) - z(\xi_{n}) + w \right) - w + \alpha V_{n+1}^{P}, \quad (\text{EC.18})$$

with boundary condition $V_{N+1}^P = 0$.

We will now demonstrate that it is optimal to search (set $\omega_n = 1$) in every period. There has to be at least one period in which it is optimal to search, otherwise the principal would pay cost w in any period but never get any benefits from search, in which case it would better to offer no contract at all (in violation of the $V_1^P > 0$ assumption). Assume there are two consecutive periods, k and k + 1, such that that under the optimal policies $\{\omega_n^*\}, \{\xi_n^*\}, \{\psi_n^*\}$ the agent searches in period k + 1 but not in k. Then, from (EC.18), we have

$$V_{k}^{P} = -w + \alpha \left(p \int_{\max\{\xi_{k+1}^{*}, \psi_{k+1}^{*}\}}^{\bar{u}} (x - \alpha V_{k+2}^{P}) dF(x) - z(\xi_{k+1}^{*}) + \alpha V_{n+1}^{P} \right).$$

Using an alternate policy which switches the order of strategies for these two periods (searching and using $\xi_{k+1}^*, \psi_{k+1}^*$ in period k and then not searching in period k+1) yields

$$V_k^{P\prime} := p \int_{\max\{\xi_{k+1}^*, \psi_{k+1}^*\}}^{\bar{u}} \left(x - \alpha^2 V_{k+2}^P + \alpha w\right) dF(x) - z(\xi_{k+1}^*) - \alpha w + \alpha^2 V_{n+1}^P.$$

Looking at the difference between these two value functions gives

$$\begin{split} V_k^{P'} - V_k^P &= p \int_{\max\{\xi_{k+1}^*, \psi_{k+1}^*\}}^u \left((1-\alpha)x + \alpha w \right) dF(x) - (1-\alpha)z(\xi_{k+1}^*) + (1-\alpha)w, \\ &= (1-\alpha) \left[p \int_{\max\{\xi_{k+1}^*, \psi_{k+1}^*\}}^{\bar{u}} x \, dF(x) - z(\xi_{k+1}^*) \right] + (1-\alpha + p\alpha \bar{F}(\max\{\xi_{k+1}^*, \psi_{k+1}^*\}))w, \\ &\geq (1-\alpha) \left[p \bar{F}(\max\{\xi_{k+1}^*, \psi_{k+1}^*\}) \alpha V_{k+2}^P - w \right] + (1-\alpha + p\alpha \bar{F}(\max\{\xi_{k+1}^*, \psi_{k+1}^*\}))w, \\ &= p\alpha \bar{F}(\max\{\xi_{k+1}^*, \psi_{k+1}^*\}) \left((1-\alpha) V_{k+2}^P + w \right) > 0. \end{split}$$

Here, the first inequality follows from the assumed optimality of $\omega_{k+1}^*, \xi_{k+1}^*, \psi_{k+1}^*$, while the second follows from $-w/(1-\alpha)$ being a strict lower bound for the value function under an optimal policy (it is the NPV of paying -w in every period ad infinitum). Consequently, $V_k^{P'} > V_k^P$, contradictory to optimality of $\{\omega_n^*\}, \{\xi_n^*\}, \{\psi_n^*\}$. Thus, if there are periods in which the agent does not search, they have to come at the end of the search process, or in other words, there exists $K \in \{1, ..., N\}$ such that $\omega_n^* = 1$ for $n \leq K$ and $\omega_n^* = 0$ for n > K.

We note that there exists at least one period *i* such that $p \int_{\max\{\xi_i^*, \psi_i^*\}}^{\bar{u}} x \, dF(x) - z(\xi_i^*) > 0$, otherwise the principal could not create any value through delegation, which would be in contradiction to the $V_1^P > 0$ assumption. But then if K < N, in period K + 1, setting $\omega_{K+1} = 1$, $\xi_{K+1} = \xi_i^*$, and $\psi_{K+1} = \psi_i^*$ would yield higher expectation than using $\omega_{K+1}^* = 0$, contradictory to optimality of $\{\omega_n^*\}$. Consequently K = N, and thus for each $n \in \{1, ..., N\}$ we have $\omega_n^* = 1$. Using this, (EC.18) simplifies to

$$V_n^P = \max_{\xi_n \in \text{supp } F} p \int_{\max\{\xi_n, \alpha V_{n+1}^P\}}^u (x - \alpha V_{n+1}^P) \, dF(x) - z(\xi_n) + \alpha V_{n+1}^P, \tag{EC.19}$$

with boundary condition $V_{N+1}^P = 0$. The optimal contract is then given by Theorem 1, and is the one that implements the above-given $\{\omega_n^*\}$ alongside with $\{\xi_n^*\}$ that solves (EC.19).

It remains to prove the properties stated in parts (i)–(v) of the theorem.

Part (i): $\min\{\xi^{LC}, \uparrow \alpha V_{n+1}^P \ \uparrow\} \leq \xi_n^* \leq \min\{\xi^{LC}, \uparrow \alpha V_{n+1}^P + e \ \uparrow\}, \ \forall n \in \{1, ..., N\}.$ Consider the objective function in (EC.19) as a function of ξ_n on domain supp F. In situations where $\xi^{LC} \leq \alpha V_{n+1}^P$, the optimal policy is $\xi_n^* = \xi^{LC}$ – as is evident from (EC.19), where the integral term is constant in ξ_n for $\xi_n \leq \alpha V_{n+1}^P$ (and strictly decreasing in ξ_n afterwards) while the cost term $z(\xi_n)$ is minimized at $\xi_n = \xi^{LC}$ (per Lemma 1). Consider now the situations where $\xi^{LC} > \alpha V_{n+1}^P$. Then, it has to be that $\xi_n^* \leq \xi^{LC}$ because higher values would both weakly increase the costs (as $z(\xi_n)$ is minimized at ξ^{LC}) and strictly decrease the value from search (the integral term). In that case, we also have $\xi_n^* \geq \alpha V_{n+1}^P$, because the integral term in (EC.19) is constant for $\xi_n \leq \alpha V_{n+1}^P$ whereas the costs are strictly decreasing in ξ_n on the same region (by Lemma 1, as $\xi^{LC} > \alpha V_{n+1}^P$). Thus, in the case where $\xi^{LC} > \alpha V_{n+1}^P$, we can expand $z(\xi_n)$ and rewrite (EC.19) as

$$V_{n}^{P} = \max_{\xi_{n} \in \text{supp } F \mid \alpha V_{n+1}^{P} \le \xi_{n} \le \xi^{\text{LC}}} p \int_{\xi_{n}}^{\bar{u}} (x - e - \alpha V_{n+1}^{P}) dF(x) - w - p\bar{F}(\xi_{n})(\alpha s_{n}^{B}(\xi_{n}) - w) + \alpha V_{n+1}^{P}, \quad (\text{EC.20})$$

from which we see that the integral term is strictly increasing in ξ_n for $\xi_n \leq \alpha V_{n+1}^P + e$ (and strictly decreasing afterwards), while the term $p\bar{F}(\xi_n)(\alpha s_n^B(\xi_n) - w)$ is strictly increasing in ξ_n . Therefore, if $[\alpha V_{n+1}^P, \alpha V_{n+1}^P + e] \neq \emptyset$ then $\xi_n^* \in [\alpha V_{n+1}^P, \alpha V_{n+1}^P + e] \cap \text{supp } F$, otherwise $\xi_n^* = \uparrow^* \alpha V_{n+1}^P + e \uparrow$. Finally, because $\xi^{\text{LC}} \leq \alpha V_{n+1}^P$ implies $\xi_n^* = \xi^{\text{LC}}$ and $\xi^{\text{LC}} > \alpha V_{n+1}^P$ implies $\xi_n^* \in [\alpha V_{n+1}^P, \min\{\xi^{\text{LC}}, \uparrow^* \alpha V_{n+1}^P + e \uparrow\}] \cap \text{supp } F$, noticing that $\xi_n \geq x$ (for any x) also implies $\xi_n \geq \uparrow^* x \uparrow$ we obtain $\min\{\xi^{\text{LC}}, \uparrow^* \alpha V_{n+1}^P + e \uparrow\} \leq \xi_n^* \leq \min\{\xi^{\text{LC}}, \uparrow^* \alpha V_{n+1}^P + e \uparrow\}$.

 $\begin{aligned} Part \text{ (ii): } \{\psi_n^*\}, \{V_n^P\} \downarrow n. \text{ At any time } m, \text{ following the policy } (\{\xi_n^{\text{shifted}}\}, \{\psi_n^{\text{shifted}}\}) \text{ given by } \forall n: \xi_n^{\text{shifted}} := \\ \xi_{n+1}^*, \psi_n^{\text{shifted}} &:= \psi_{n+1}^* \text{ yields } V_{m+1}^P \text{ in expectation}, \\ ^{21} \text{ implying that the optimal policy has to yield at least that much, i.e., } \\ V_m^P \ge V_{m+1}^P \text{ and thus } \{V_n^P\} \downarrow n. \text{ Then, } \{\psi_n^*\} \downarrow n \text{ follows from } \\ \psi_n^* = \upharpoonright \alpha V_{n+1}^P \stackrel{\text{``}}{\rightarrow}. \end{aligned}$

Part (iii): $\{\xi_n^*\}$ is constant and equal to ξ^{LC} up to a point, after which it is decreasing in n. We have $\{V_n^P\} \downarrow n$ from part (ii) of the theorem and that $\alpha V_{n+1}^P \ge \xi^{\text{LC}}$ implies $\xi_n^* = \xi^{\text{LC}}$ from part (i). Hence there exists $K \ge 0$ such that the first K elements of $\{\xi_n^*\}$ are equal to ξ^{LC} and, for all n > K, $\xi^{\text{LC}} > \alpha V_{n+1}^P$. It follows that $\{\xi_n^*\}$ solves (EC.20) for n > K + 1. Denote the objective function of (EC.20) as $ob(\xi, n) := p \int_{\xi}^{\bar{u}} (x - e - \alpha V_{n+1}^P) dF(x) - w - p\bar{F}(\xi)(\alpha s_n^B(\xi) - w) + \alpha V_{n+1}^P$. Thus defined $ob(\xi, n)$ is submodular, and so, by Topkis's theorem (Topkis 1978), $\{\xi_n^*\} \downarrow n$.

 $\begin{array}{l} Part \mbox{ (iv): } for \mbox{ all } n \in \{1,...,N\}, \mbox{ } \xi_n^* \leq \xi_n^{\rm FB} \mbox{ and } \psi_n^* < \xi_n^{\rm FB}. \mbox{ By Lemma A2, } \xi_n^{\rm FB} = \vec{r} \mbox{ } \alpha V_{n+1}^{\rm FB} + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } \alpha V_{n+1}^P + e \mbox{ } 1 \geq \vec{r} \mbox{ } 0 = \vec{r} \mbox{ }$

Part (v): Single crossing of ψ_n^* and ξ_n^* . From part (iii), ξ_n^* is constant and equal to ξ^{LC} up to a certain period, after which it is weakly decreasing in n. Then, from part (i), in the non-constant region of ξ_n^* we have that $\xi_n^* \ge \alpha V_n^P$, thus also $\xi_n^* \ge r^* \alpha V_n^P \stackrel{h}{=} \psi_n^*$ (so no crossings within this region). Hence, $\psi_n^* > \xi_n^*$ can only be in the region where $\xi_n^* = \xi^{\text{LC}}$, from which the $\xi_{M-1}^* = \xi^{\text{LC}}$ property immediately follows. Using part (ii), ψ_n^* is weakly decreasing and thus crosses the constant function at most once.

Proof of Proposition 1.

In the first-best solution (analogously, for the principal who delegates search), V_n^{FB} (resp., V_n^P) is the value of the rest of the search under an optimal policy. As $N \to \infty$ this value must be equal in all periods (Proposition 5.4.1 of Bertsekas 2017). Hence, as $\xi_n^{\text{FB}} = \uparrow \alpha V_{n+1}^{\text{FB}} + e \uparrow$ (Lemma A2) and $\psi_n^* = \uparrow \alpha V_{n+1}^{\text{FB}} \uparrow$ (Theorem 3), those policies are equal in all periods as well. Then, stationarity of ξ_n^* follows from (17), while $\xi_n^{\text{FB}} \ge \xi_n^*$ and $\xi_n^{\text{FB}} \ge \xi_n^*$ follow from Theorem 3(iv). Because $\psi_n^* = \uparrow \alpha V_{n+1}^P \uparrow$, we have $\uparrow \psi_n^* + e \uparrow = \uparrow \uparrow \alpha V_{n+1}^P \uparrow + e \uparrow \geq \uparrow \alpha V_{n+1}^P + e \uparrow \geq \xi_n^*$, where the last inequality follows from Theorem 3(i). The time invariance of s_n^B then follows from equation (11) because ξ_n^* is constant.

Now we turn to the risk-neutral case. Using the expression for s_n^B from (11) gives us the equality $\lim_{r\to 0} s_n^B(\xi_n) = w/\alpha + c/(p\bar{F}(\xi_n)\alpha)$. Observe that the expected single-period payment to the agent is now independent of ξ_n , which follows more explicitly from substituting this expression for s_n^B into (14). Thus the principal can induce any policy (including the first-best) for the same expected cost. Given that the participation constraint is binding (as shown in the proof of Theorem 1), the principal is also able to extract all value from the agent; hence the contract achieves the first-best.

For part (iii) of the proposition, $z(\xi)$ is an increasing function of ξ when e = 0, so $\xi^{\text{LC}} = \text{min supp } F$ and thus also $\xi_n^* = \text{min supp } F$ by Theorem 3(i). First-best being attained for p = 1 (but not otherwise) follows

²¹ The shifted policy leaves undefined what to do in the last period (as ξ_{N+1}^* is undefined). However, we do not need to know what happens in the last period, only that there is some policy we can use there that does not result in a negative payoff. As we demonstrated earlier in the proof, a policy pair (ξ, ψ) exists for which $p \int_{\max\{\xi,\psi\}}^{\bar{u}} x \, dF(x) - z(\xi) > 0$ which can then be used in the last period.

from $s_n^B(\min \operatorname{supp} F) = (c+w)/\alpha$ when p = 1 (but not when p < 1), which then makes the first-best value function (13) equivalent to the one under the optimal contract, as given by (17).

Lastly, part (iv) of the proposition follows directly from Theorem 3 under parameter value w = 0.

Proof of Proposition 2.

Denote by V_n^{IN} the principal's value function (her NPV under an optimal policy) in period n if she searches on her own (in-house); V_n^{IN} is given by (1) with $c = c_P$, $p = p_P$, and $F(\cdot) = F_P(\cdot)$. From Lemma A1 it follows that the bracketed term in (1) is positive, so

$$V_n^{\rm IN} = p_P \int_{\alpha V_{n+1}^{\rm IN}}^{\bar{u}_P} \bar{F}_P(x) \, dx - c_P + \alpha V_{n+1}^{\rm IN}, \tag{EC.21}$$

with boundary condition $V_{N+1}^{\text{IN}} = 0$. We similarly use V_n^{OUT} to denote the value function for the principal who delegates (outsources) search. Applying (17) now yields

$$V_n^{\text{OUT}} = \max_{\xi_n \in \text{supp } F} p_A \int_{\max\{\xi_n, \alpha V_{n+1}^{\text{OUT}}\}}^{\bar{u}_A} (x - \alpha V_{n+1}^{\text{OUT}}) \, dF_A(x) - z(\xi_n) + \alpha V_{n+1}^{\text{OUT}}, \tag{EC.22}$$

with boundary condition $V_{N+1}^{\text{OUT}} = 0$. The rest of our proof proceeds in three steps: after showing that these two value functions do not cross more than once, we characterize situations in which they do not cross at all.

Step 1: Single-crossing property. Consider the problem faced by a principal who can, in each period, choose whether to search herself in that period or to delegate the search task. Thus she need not commit to one of these options at the beginning of the search and can mix and match as needed, getting the best of both options. (Such principal is also the subject of Remark 2 in the main text.) Our motive for considering such a principal is to establish that optimal single-period outsourcing decisions depend primarily on the value of the rest of the search (i.e., from period n + 1 onward): delegation is optimal when that value is low, but an in-house search is optimal when it is high.

Let \hat{V}_n denote this principal's value function. To express that function concisely, we first define

$$\hat{\xi}(V) = \min \underset{\xi \in \operatorname{supp} F}{\operatorname{arg\,max}} p_A \int_{\max\{\xi, \alpha V\}}^{\bar{u}_A} (x - \alpha V) \, dF_A(x) - z(\xi) + \alpha V. \tag{EC.23}$$

This function gives the optimal threshold (reporting policy) $\hat{\xi}$ that the principal should induce when delegating, and her value function in the next period will be V. The property given in part (i) of Theorem 3 can also be extended, via an identical proof, to show that $\hat{\xi}(V) \in [\min\{\xi^{LC}, \uparrow^{\alpha} V \uparrow\}, \min\{\xi^{LC}, \uparrow^{\alpha} V + e \uparrow\}] \cap \text{supp } F$. Using (EC.21) and (EC.22), the principal's value function is given by

$$\hat{V}_{n} = \max\left\{p_{P} \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_{P}} \bar{F}_{P}(x) \, dx - c_{P}, \\ p_{A} \int_{\max\{\hat{\xi}(\hat{V}_{n+1}), \alpha \hat{V}_{n+1}\}}^{\bar{u}_{A}} (x - \alpha \hat{V}_{n+1}) \, dF_{A}(x) - z(\hat{\xi}(\hat{V}_{n+1}))\right\} + \alpha \hat{V}_{n+1},$$
(EC.24)

with boundary condition $\hat{V}_{N+1} = 0$. We can also express (EC.24) as

$$\begin{split} \hat{V}_n &= \max_{h_n \in \{0,1\}} h_n \left[p_A \int_{\max\{\hat{\xi}(\hat{V}_{n+1}), \alpha \hat{V}_{n+1}\}}^{\bar{u}_A} (x - \alpha \hat{V}_{n+1}) dF_A(x) - p_P \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_P} \bar{F}_P(x) dx + c_P - z(\hat{\xi}(\hat{V}_{n+1})) \right] \\ &+ p_P \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_P} \bar{F}_P(x) dx - c_P + \alpha \hat{V}_{n+1}, \end{split}$$

where $h_n = 1$ (resp. $h_n = 0$) stands for the decision to delegate search (resp. not delegate) in period n. It is then optimal to delegate if and only if the expression in brackets next to h_n is positive. We abbreviate that expression as VD_n, the value of delegation, where

$$VD_{n}(\hat{V}_{n+1}) = p_{A} \int_{\max\{\hat{\xi}(\hat{V}_{n+1}),\alpha\hat{V}_{n+1}\}}^{\bar{u}_{A}} (x - \alpha\hat{V}_{n+1}) dF_{A}(x) - p_{P} \int_{\alpha\hat{V}_{n+1}}^{\bar{u}_{P}} \bar{F}_{P}(x) dx + c_{P} - z(\hat{\xi}(\hat{V}_{n+1})). \quad (EC.25)$$

Next we show that $VD_n(\hat{V}_{n+1})$ is a weakly decreasing function of \hat{V}_{n+1} . Consider first the case when $\alpha \hat{V}_{n+1} \ge \xi^{\text{LC}}$. In this case, we have that $\hat{\xi}(\hat{V}_{n+1}) = \xi^{\text{LC}}$, which enables the following simplification:

$$\begin{aligned} \mathrm{VD}_{n}(\hat{V}_{n+1}) &= p_{A} \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_{A}} (x - \alpha \hat{V}_{n+1}) \, dF_{A}(x) - p_{P} \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_{P}} \bar{F}_{P}(x) \, dx + c_{P} - z(\xi^{\mathrm{LC}}) \\ &= \int_{\alpha \hat{V}_{n+1}}^{\max\{\bar{u}_{A},\bar{u}_{P}\}} (p_{A} \bar{F}_{A}(x) - p_{P} \bar{F}_{P}(x)) \, dx + c_{P} - z(\xi^{\mathrm{LC}}), \end{aligned}$$

which is decreasing in \hat{V}_{n+1} because the integrand is positive.

Now consider the case where $\alpha \hat{V}_{n+1} \leq \xi^{\text{LC}}$ (thus $\hat{\xi}(\hat{V}_{n+1}) \geq \hat{\gamma} \alpha \hat{V}_{n+1}$ $\hat{\gamma}$), and let $\xi^{\text{LC}}/\alpha \geq \hat{V}_{n+1}^{H} > \hat{V}_{n+1}^{L}$. We will show that the following difference is positive

$$\begin{split} \mathrm{VD}_{n}(\hat{V}_{n+1}^{L}) - \mathrm{VD}_{n}(\hat{V}_{n+1}^{H}) &= p_{A} \int_{\hat{\xi}(\hat{V}_{n+1}^{L})}^{\tilde{u}_{A}} (x - \alpha \hat{V}_{n+1}^{L}) dF_{A}(x) - p_{A} \int_{\hat{\xi}(\hat{V}_{n+1}^{H})}^{\tilde{u}_{A}} (x - \alpha \hat{V}_{n+1}^{H}) dF_{A}(x) \\ &\quad - z(\hat{\xi}(\hat{V}_{n+1}^{L})) + z(\hat{\xi}(\hat{V}_{n+1}^{H})) - p_{P} \int_{\alpha \hat{V}_{n+1}^{L}}^{\alpha \hat{V}_{n+1}^{H}} \bar{F}_{P}(x) dx \\ &\geq p_{A} \int_{\alpha \hat{V}_{n+1}^{L}+e}^{\tilde{u}_{A}} (x - \alpha \hat{V}_{n+1}^{L}) dF_{A}(x) - p_{A} \int_{\hat{\xi}(\hat{V}_{n+1}^{H})}^{\tilde{u}_{A}} (x - \alpha \hat{V}_{n+1}^{H}) dF_{A}(x) \\ &\quad - z(\hat{\xi}(\hat{V}_{n+1}^{L})) + z(\hat{\xi}(\hat{V}_{n+1}^{H})) - p_{P} \int_{\alpha \hat{V}_{n+1}^{L}}^{\alpha \hat{V}_{n+1}^{H}} \bar{F}_{P}(x) dx \\ &= p_{A} \int_{\alpha \hat{V}_{n+1}^{L}+e}^{\tilde{u}_{A}} (x - \alpha \hat{V}_{n+1}^{L} - e) dF_{A}(x) - p_{A} \int_{\hat{\xi}(\hat{V}_{n+1}^{H})}^{\tilde{u}_{A}} (x - \alpha \hat{V}_{n+1}^{H} - e) dF_{A}(x) \\ &\quad + p_{A} e \bar{F}_{A}(\alpha \hat{V}_{n+1}^{L} + e) - p_{A} e \bar{F}_{A}(\hat{\xi}(\hat{V}_{n+1}^{H})) \\ &\quad - z(\hat{\xi}(\hat{V}_{n+1}^{L})) + z(\hat{\xi}(\hat{V}_{n+1}^{H})) - p_{P} \int_{\alpha \hat{V}_{n+1}^{H}}^{\alpha \hat{V}_{n+1}^{H}} \bar{F}_{P}(x) dx. \end{split}$$

Then, as $\hat{\xi}(\hat{V}_{n+1}^H) \leq \hat{r} \alpha \hat{V}_{n+1}^H + e^{\gamma}$ we have $\int_{\hat{\xi}(\hat{V}_{n+1}^H)}^{\bar{u}_A} (x - \alpha \hat{V}_{n+1}^H - e) dF_A(x) = \int_{\hat{r} \alpha \hat{V}_{n+1}^H + e^{\gamma}}^{\bar{u}_A} (x - \alpha \hat{V}_{n+1}^H - e) dF_A(x) + \int_{x \in [\hat{\xi}(\hat{V}_{n+1}^H), \hat{r} \alpha \hat{V}_{n+1}^H + e^{\gamma})} (x - \alpha \hat{V}_{n+1}^H - e) dF_A(x) \leq \int_{\alpha \hat{V}_{n+1}^H + e}^{\bar{u}_A} (x - \alpha \hat{V}_{n+1}^H - e) dF_A(x) = \int_{\alpha \hat{V}_{n+1}^H + e^{\gamma}}^{\bar{u}_A} \bar{F}_A(x) dx = \int_{\alpha \hat{V}_{n+1}^H}^{\bar{u}_A} \bar{F}_A(x + e) dx$. Inserting this inequality into the equation above yields

$$\begin{split} \mathrm{VD}_{n}(\hat{V}_{n+1}^{L}) - \mathrm{VD}_{n}(\hat{V}_{n+1}^{H}) &\geq p_{A} \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_{A}} \bar{F}_{A}(x+e) dx - p_{A} \int_{\alpha \hat{V}_{n+1}}^{\bar{u}_{A}} \bar{F}_{A}(x+e) dx \\ &+ p_{A} e \bar{F}_{A}(\alpha \hat{V}_{n+1}^{L}+e) - p_{A} e \bar{F}_{A}(\hat{\xi}(\hat{V}_{n+1}^{H})) \\ &- z(\hat{\xi}(\hat{V}_{n+1}^{L})) + z(\hat{\xi}(\hat{V}_{n+1}^{H})) - p_{P} \int_{\alpha \hat{V}_{n+1}}^{\alpha \hat{V}_{n+1}^{H}} \bar{F}_{P}(x) dx \\ &= \int_{\alpha \hat{V}_{n+1}^{L}}^{\alpha \hat{V}_{n+1}^{H}} \left(p_{A} \bar{F}_{A}(x+e) - p_{P} \bar{F}_{P}(x) \right) dx \\ &+ p_{A} e \bar{F}_{A}(\alpha \hat{V}_{n+1}^{L}+e) - p_{A} e \bar{F}_{A}(\hat{\xi}(\hat{V}_{n+1}^{H})) - z(\hat{\xi}(\hat{V}_{n+1}^{L})) + z(\hat{\xi}(\hat{V}_{n+1}^{H})) \end{split}$$

Expanding the $z(\cdot)$ terms yields

$$\begin{split} \mathrm{VD}_{n}(\hat{V}_{n+1}^{L}) - \mathrm{VD}_{n}(\hat{V}_{n+1}^{H}) &\geq \int_{\alpha \hat{V}_{n+1}^{L}}^{\alpha \hat{V}_{n+1}^{H}} \left(p_{A} \bar{F}_{A}(x+e) - p_{P} \bar{F}_{P}(x) \right) dx \\ &\quad - p_{A} \bar{F}_{A}(\hat{\xi}(\hat{V}_{n+1}^{L})) (\alpha s_{n}^{B}(\hat{\xi}(\hat{V}_{n+1}^{L})) - w) + p_{A} \bar{F}_{A}(\hat{\xi}(\hat{V}_{n+1}^{H})) (\alpha s_{n}^{B}(\hat{\xi}(\hat{V}_{n+1}^{H})) - w) \\ &\quad + p_{A} e(\bar{F}_{A}(\alpha \hat{V}_{n+1}^{L} + e) - \bar{F}_{A}(\hat{\xi}(\hat{V}_{n+1}^{L}))). \end{split}$$

Here, the first line is weakly positive as the integrand is weakly positive, the second is weakly positive becaue $p_A \bar{F}_A(x)(\alpha s_n^B(x) - w) \uparrow x$ and $\hat{\xi}(V) \uparrow V$, and the last is also weakly positive as $\hat{\xi}(\hat{V}_{n+1}^L)) \leq \vec{r} \alpha \hat{V}_{n+1}^L + e \uparrow$, thus $\bar{F}_A(\hat{\xi}(\hat{V}_{n+1}^L)) \geq \bar{F}_A(\vec{r} \alpha \hat{V}_{n+1}^L + e \uparrow) = \bar{F}_A(\alpha \hat{V}_{n+1}^L + e)$. Consequently, $\mathrm{VD}_n(\hat{V}_{n+1}^L) - \mathrm{VD}_n(\hat{V}_{n+1}^H) \geq 0$, that is $\mathrm{VD}_n(\hat{V}_{n+1})$ is weakly decreasing.

As a result, the optimal control sequence $\{h_n^*\}$ is weakly increasing in n. So if there exists k such that $h_k^* = 1$, then for every $n \in \{k, ..., N\}$ we have $h_n^* = 1$ and hence $\hat{V}_n = V_n^{\text{OUT}}$ for $n \ge k$. We conclude that, if the principal who solves (EC.24) should delegate, then (a) so should the main model's principal and (b) their value functions are identical. So if there exists period n in which $V_n^{\text{OUT}} \ge V_n^{\text{IN}}$, then the same holds for any subsequent period. Conversely, if there exists n, in which $V_n^{\text{OUT}} \le V_n^{\text{IN}}$, then the same holds in every preceding period, completing the single-crossing property.

Step 2: Situations where it is optimal not to delegate. Consider the problem in the last period N. The principal receives $V_N^{IN} = p_P \int_{\underline{u}_P}^{\overline{u}_P} \overline{F}_P(x) dx - c_P$ if searching on her own or $V_N^{OUT} = \max_{\xi \in \text{supp } F} p_A \int_{\xi}^{\overline{u}_A} x \, dF_A(x) - z(\xi)$ if she delegates. So in the last period, the principal is better-off conducting her own search if and only if $V_N^{IN} \ge V_N^{OUT}$; expanding this expression gives the condition (18). Finally, if this condition holds then it is optimal to conduct an in-house search not only in the last period but in every preceding period as well, a result that follows from Step 1 of this proof and completes part (i) of the proposition.

Step 3: Situations where not delegating is optimal. Notice that $V_n^{\text{IN}} \ge V_n^{\text{OUT}} \Rightarrow V_{n-1}^{\text{IN}} \ge V_{n-1}^{\text{OUT}}$, which was shown in part 1, also implies that if it is optimal to delegate if there are N periods until the horizon, it is also optimal to delegate if there are K periods until the horizon, for every $K \in \{1, ..., N\}$. Thus we can identify situations in which delegation is optimal – irrespective of the number of periods – by considering $N \to \infty$, in which case the optimal policies are stationary (see Lippman and McCall 1976). Then, by (EC.21) and (EC.22), the value functions in any period n are given by

$$V^{\mathrm{IN}} = \frac{1}{1-\alpha} \left(p_P \int_{\alpha V^{\mathrm{IN}}}^{u_P} \bar{F}_P(x) \, dx - c_P \right) \quad \text{and}$$
$$V^{\mathrm{OUT}} = \frac{1}{1-\alpha} \max_{\xi \in \mathrm{supp}\, F} \left(p_A \int_{\max\{\xi, \alpha V^{\mathrm{OUT}}\}}^{\bar{u}_A} (x - \alpha V^{\mathrm{OUT}}) \, dF_A(x) - z(\xi) \right).$$

Therefore, if $V^{\text{OUT}} \ge V^{\text{IN}}$ then it is optimal to delegate for every possible horizon length $N \in \mathbb{N}$. This completes part (ii) of the proposition, and also part (iii) by applying the single-crossing property (Step 1 of this proof). \Box

Proposition A1 and Lemma A3 (to follow) are the analogues of (respectively) Theorems 1 and 2 and Lemma 1 for the setting in which u is not contractible. Both of these results are used in our proof of Theorem 4.

Proposition A1 (Implementable policies under unspecified search) If Y^n is the sole contractible variable and the principal commits to an acceptance policy $\{\Psi_n\}$ when offering the contract, then the principal can implement any policy pair $(\{\omega_n\}, \{\Xi_n\})$ that satisfies the following conditions for every period $n \in$ $\{1, ..., N\}$: either $\omega_n = 0$ or $\Psi_n \subseteq \Xi_n$ and $pP(\Psi_n)^{1_{\Xi_n \neq \text{supp }F}} > 1 - 1/\exp[rc(1-\alpha)\alpha^{-1}]$. The contract that implements $\{\Xi_n\}$ at the lowest possible cost to the principal has two possible components. In periods during which the agent does not bring all alternatives to the principal $(\Xi_n \neq \text{supp }F)$, the contract pays the agent a bonus of $s_n^B(\Psi_n)$ if the principal accepts the delivered alternative or of s_n^0 if she does not accept it or no alternative was delivered; here $s_n^B(\cdot)$ and s_n^0 are given by (11). In periods where $(\Xi_n = \text{supp }F)$ the bonus is equal to $s_n^B(\text{supp }F)$ and is paid for any alternative presented to the principal – that is, irrespective of her decision to accept or reject it. This contract is renegotiation proof.

Proof. From Theorem 1 of Plambeck and Zenios (2000) it follows that the problem of finding a contract that implements any admissible search policy $\{\omega_n\}$ and reporting policy $\{\Xi_n\}$ at the lowest possible cost to the principal can be decomposed into N single-period problems, each given by

$$\underset{s_n^A, s_n^R, s_n^0}{\operatorname{arg\,min}} p\omega_n P(\Psi_n \cap \Xi_n) s_n^A + p\omega_n P(\Xi_n \setminus \Psi_n) s_n^R + (1 - p\omega_n P(\Xi_n)) s_n^0;$$
(EC.26)

$$-p\omega_{n}P(\Psi_{n}\cap\Xi_{n})\exp[-r(1-\alpha)(s_{n}^{A}-\alpha^{-1}c)] - p\omega_{n}P(\Xi_{n}\setminus\Psi_{n})\exp[-r(1-\alpha)(s_{n}^{R}-\alpha^{-1}c)] -(1-p\omega_{n}P(\Xi_{n}))\exp[-r(1-\alpha)(s_{n}^{0}-\alpha^{-1}c)] = -\exp[-r\alpha^{-1}(1-\alpha)w];$$
(EC.27)

$$-pP(\Psi_n \cap \Xi) \exp[-r(1-\alpha)(s_n^A - \alpha^{-1}c)] - pP(\Xi \setminus \Psi_n) \exp[-r(1-\alpha)(s_n^R - \alpha^{-1}c)]$$
(EC.28)

$$-(1-pP(\Xi))\exp[-r(1-\alpha)(s_n^0-\alpha^{-1}c)] \le -\exp[-r\alpha^{-1}(1-\alpha)w] \quad \forall \Xi \in \mathcal{F};$$

$$-\exp[-r(1-\alpha)s_n^0] \le -\exp[-r\alpha^{-1}(1-\alpha)w].$$
(EC.29)

As before, s_n^A (resp. s_n^R) is the agent's reward for bringing an alternative to the principal that she accepts (resp. rejects), and s_n^0 is his reward if he does not bring her any alternative. The objective function (EC.26) minimizes the expected sum of these rewards conditional on the agent using the desired policies ω_n and Ξ_n . Equation (EC.27) is the (binding) participation constraint. The incentive compatibility constraints (EC.28) and (EC.29) ensure that the agent could not do better by choosing a different reporting policy or by declining to pay the search cost. As in Proposition 1, if $\omega_n = 0$ then (EC.26)-(EC.29) has a trivial solution of $s_n^A =$ $s_n^R = s_n^0 = w/\alpha$, so we proceed to examine the case when $\omega_n = 1$.

Notice from the incentive compatibility constraint (EC.28) that if $s_n^A, s_n^R > s_n^0$, then it is strictly optimal for the agent to report all alternatives, if $s_n^A > s_n^0 > s_n^R$ it is strictly optimal for him to report only the ones that the principal will accept (and the exact value of s_n^R is redundant), and if $s_n^R > s_n^0 > s_n^A$ it is strictly optimal for the agent to report all that the principal will *not* accept (and the exact value of s_n^A is redundant). Thus, unless $s_n^A = s_n^0$ or $s_n^R = s_n^0$, the only search policies which can be induced are $\Xi_n = \text{supp } F$, $\Xi_n = \Psi_n$, and $\Xi_n = \text{supp } F \setminus \Psi_n$. Thus, we will separate this problem into 5 exhaustive cases and solve each case individually.

Case 1. Consider first the contract parameters $s_n^R = s_n^0$. Then we also need $s_n^A > s_n^R$, otherwise the contract would induce $\omega_n = 0$. Such contract parameters will induce the agent to report any alternative in Ψ_n to the principal, and make him indifferent about reporting other alternatives. We first solve relaxed version of the problem, where the constraint (EC.28) is replaced by $s_n^R = s_n^0$, which is then given by

$$\arg\min_{s_n^A, s_n^0} pP(\Psi_n) s_n^A + (1 - pP(\Psi_n)) s_n^0;$$
(EC.30)

$$-pP(\Psi_n)\exp[-r(1-\alpha)(s_n^A-\alpha^{-1}c)] - (1-pP(\Psi_n))\exp[-r(1-\alpha)(s_n^0-\alpha^{-1}c)] = -\exp[-r\alpha^{-1}(1-\alpha)w]; \quad (\text{EC.31})$$

$$-\exp[-r(1-\alpha)s_n^0] \le -\exp[-r\alpha^{-1}(1-\alpha)w].$$
(EC.32)

This problem is identical to (EC.5)–(EC.7), which was solved in our proof of Theorem 1. It is thus solvable only if the implementability condition $pP(\Psi_n) > 1 - 1/\exp[rc(1-\alpha)\alpha^{-1}]$ holds, and its solution is

$$s_n^A = \frac{w+c}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{pP(\Psi_n)}{1 - (1 - pP(\Psi_n)) \exp[rc(1-\alpha)\alpha^{-1}]}, \qquad s_n^0 = s_n^R = \frac{w}{\alpha}.$$
 (EC.33)

Plugging this solution back into the incentive compatibility constraint (EC.28), we obtain

$$-pP(\Psi_n \cap \Xi) \exp[-r(1-\alpha)(s_n^A - \alpha^{-1}c)]$$
$$-(1-pP(\Psi_n \cap \Xi)) \exp[-r(1-\alpha)(s_n^0 - \alpha^{-1}c)]$$
$$\leq -\exp[-r\alpha^{-1}(1-\alpha)w], \forall \Xi \in \mathcal{F}.$$

This condition is satisfied for any Ξ_n such that $\Psi_n \subseteq \Xi_n$. This is because the LHS of the condition is maximized at $\Xi = \Psi_n$ (as $s_n^A > s_n^0$) and because, by (EC.31), the value of the left-hand side when evaluated at $\Xi = \Psi_n$ is equal to that of the right-hand side. Hence (EC.33) is a *candidate* solution to the original system (EC.26)–(EC.29) for any Ξ_n that satisfies $\Psi_n \subseteq \Xi_n$. The solution is only a candidate solution as the optimization restricts the contract parameters to those that satisfy $s_n^R = s_n^0 < s_n^A$. Thus, we still need to verify that none of the other cases are capable of implementing the same reporting policy at a lower cost. Note that the expression in (EC.33) for optimal payments is independent of the desired policy Ξ_n ; in fact, the optimal incentive structure renders the agent indifferent between all policies that satisfy $\Psi_n \subseteq \Xi_n$.

Case 2. Consider the contract parameters $s_n^R < s_n^0$. Then we also need $s_n^A > s_n^0$, otherwise the contract would induce $\omega_n = 0$. As noted earlier, an optimal strategy for the agent would then be to report everything that the principal will accept and nothing else; this is a situation which can be solved analogously to Case 1 and yields the same candidate solution for s_n^A and s_n^0 as in (EC.33), while s_n^0 can be any $s_n^0 < s_n^0$. Another difference from Case 1 is that here we can only implement $\Xi_n = \Psi_n$. This similarity of solutions between Case 1 and 2 is driven by insensitivity to parameter s_n^R when $s_n^R < s_n^0$, as under any induced policy reporting something that the principal will reject is a zero probability event.

Case 3. Consider the contract parameters $s_n^A = s_n^0$; then we also need $s_n^R > s_n^A$, lest we induce $\omega_n = 0$. We can solve this analogously to Case 1, which is possible if the implementability condition $pP(\operatorname{supp} F \setminus \Psi_n) > 1 - 1/\exp[rc(1-\alpha)\alpha^{-1}]$ holds, and yields

$$s_n^A = \frac{w+c}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{pP(\operatorname{supp} F \setminus \Psi_n)}{1 - (1 - pP(\operatorname{supp} F \setminus \Psi_n)) \exp[rc(1-\alpha)\alpha^{-1}]}, \qquad s_n^0 = s_n^R = \frac{w}{\alpha}.$$
 (EC.34)

This is a candidate solution to problem (EC.26)–(EC.29) for any Ξ_n that satisfies $\Xi_n \subseteq \text{supp } F \setminus \Psi_n$. Note that the contract from this case will never be a part of an optimal contract, as it generates an acceptable outcome for the principal with probability 0, but is more expensive to implement than $\omega_n = 0$.

Case 4. Consider the contract parameters $s_n^A < s_n^0$; then we also need $s_n^R > s_n^A$, lest we induce $\omega_n = 0$. Solving it (analogously to Case 1) yields s_n^R and s_n^0 given by (EC.34) as the candidate solution (s^A can be anything as long as $s_n^A < s_n^0$), and can only implement reporting policies such that $\Xi_n = \operatorname{supp} F \setminus \Psi_n$. As in Case 3, the contract from this case will never be a part of an optimal contract, as it generates an acceptable outcome for the principal with probability 0, but is more expensive to implement than $\omega_n = 0$.

Case 5. Consider the contract parameters $s_n^A, s_n^R > s_n^0$. This contract will induce the agent to report all found alternatives to the principal. Using the same Jensen's inequality argument as in Theorem 1, the contract then also has to satisfy $s_n^A = s_n^R$, otherwise it would impose an additional source of risk on the risk-averse agent, requiring the principal to pay an increased risk premium in expectation. This is also solvable analogously to Case 1, which is possible if the implementability condition $p > 1 - 1/\exp[rc(1-\alpha)\alpha^{-1}]$ holds, and yields

$$s_n^A = s_n^R = \frac{w+c}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{p}{1-(1-p)\exp[rc(1-\alpha)\alpha^{-1}]} \quad \text{and} \quad s_n^0 = \frac{w}{\alpha}.$$
 (EC.35)

This is the solution to problem (EC.26)–(EC.29) for $\Xi_n = \operatorname{supp} F$ – even though the contract from Case 1 can also implement this reporting policy – as it can be observed from inserting (EC.35) and (EC.33) into (EC.26) that (EC.35) implements it for a lower cost.

Finally, the renegotiation-proof nature of this contract in all cases follows from Corollary 1 of Plambeck and Zenios (2000), the statement of the proposition follows from combining the exhaustive cases. \Box

Lemma A3. If Y^n is the sole contractible variable and the principal commits to an acceptance policy, we can restrict attention to history-independent policies without loss of optimality and under an optimal contract, threshold reporting and accepting policies are optimal, so $\Psi_n^{\text{US}*} = [\psi_n^{\text{US}*}, \bar{u}] \cap \text{supp } F$ and $\Xi_n^{\text{US}*} = [\xi_n^{\text{US}*}, \bar{u}] \cap$ supp F, for all $n \in \{1, ..., N\}$. The optimal contract implements either $\xi_n^{\text{US}*} = \psi_n^{\text{US}*}$ or $\xi^{\text{US}*} = \min \text{supp } F$ for all $n \in \{1, ..., N\}$. The optimal contract incurred by the principal when implementing reporting policy $\xi_n^{\text{US}*}$ are $z(\xi_n^{\text{US}*})$ as given by (14).

Proof. Part 1: The agent should report either all alternatives or all acceptable ones. If $\omega_n = 0$ the choice of threshold and reporting policies is redundant (any policy is optimal), so consider periods where $\omega_n = 1$. Notice that for any implementable Ξ_n other than $\Xi_n = \operatorname{supp} F$, the amount paid to the agent is the same (from Proposition A1), as is the principal's upside (she will accept any alternative with a value in Ψ_n), but the principal's evaluation costs are incurred every time an alternative is delivered to her and are thus increasing in $P(\Xi_n)$. Consequently, under an optimal contract, it has to be that $\Xi_n^{\mathrm{US}*} = \Psi_n^{\mathrm{US}*}$ or $\Xi_n^{\mathrm{US}*} = \operatorname{supp} F$, and similarly to the proof of Lemma 1, the principal's expected single-period costs to induce the agent to search and use reporting policy Ξ_n^{US} are given by $\zeta(\Xi_n^{\mathrm{US}})$ as defined in (EC.14).

Part 2: Optimality of threshold policies. Assume, by contradiction, there exists an optimal $\Psi_n^{\text{US}*}$ that is not a threshold one in a period where $\omega_n^{\text{US}} = 1$, and let $\psi_n^{\text{US}*}$ be such that $P([\psi_n^{\text{US}*}, \bar{u}] \cap \text{supp } F) = P(\Psi_n^{\text{US}*})$ (the existence of such $\psi_n^{\text{US}*}$ is guaranteed by the probability mass function μ being constant). Then, from (EC.14), the costs to the principal will be the same under both policies, acceptable alternatives will be delivered to the principal with the same probability, but the value of delivered alternatives will be better (in terms of weak first-order stochastic dominance) under $[\psi_n^{\text{US}*}, \bar{u}] \cap \text{supp } F$. Thus, using $[\psi_n^{\text{US}*}, \bar{u}] \cap \text{supp } F$ gives a higher expectation to the principal, in contradiction to optimality of $\Psi_n^{\text{US}*}$, so optimality of threshold acceptance policies follows. Then, optimality of threshold reporting policies follows from part 1 of this proof. Thus, either $\Xi_n^{US} = \Psi_n^{US}$ is optimal or $\Xi_n^{US} = \operatorname{supp} F$ is. Lastly, the principal cost function being given by (14) follows by inserting the optimal threshold policies into (EC.14).

Part 3: Optimality of History-Independent Policies. Proposition A1 reduces the contracting problem to an MDP in which the principal's choice is over the policies $\{\Xi_n\}, \{\Psi_n\}, \{\omega_n\}$. Applying part 1 of this proof to (4), the resulting MDP can be expressed as

$$V_n^{\mathrm{US}} = \max_{\omega_n \in \{0,1\}, \Psi_n \in \mathcal{F}, \Xi_n \in \{\Psi_n, \mathrm{supp}\, F\}} \left(\omega_n p \int_{\Psi_n} (x - \alpha V_{n+1}^{\mathrm{US}}) \, dF(x) - \omega_n \zeta(\Xi_n) - (1 - \omega_n) w + \alpha V_{n+1}^{\mathrm{US}} \right),$$

with boundary condition $V_{N+1}^{\text{US}} = 0$. Here, the objective function is independent of the state history, including the current state, thus so are the optimal choices of ω_n, Ξ_n , and Ψ_n .

Proof of Theorem 4.

By Lemma A3, history-independent threshold policies are optimal, and in every period it is optimal for the principal to induce a reporting policy which is equal to the announced acceptance policy $(\xi_n = \psi_n^{\text{US}})$ or to induce $\xi_n = \text{supp } F$. We can therefore express the principal's problem as the dynamic program

$$V_n^{\mathrm{US}} = \max_{\omega_n \in \{0,1\}, \psi_n \in \mathrm{supp}\, F} \left(\omega_n p \int_{\psi_n}^{\bar{u}} (x - \alpha V_{n+1}^{\mathrm{US}}) \, dF(x) - \omega_n \min\{z(\psi_n), z(\min \operatorname{supp} F)\} - (1 - \omega_n)w + \alpha V_{n+1}^{\mathrm{US}} \right),$$

with boundary condition $V_{N+1}^{\text{US}} = 0$. Here, optimality of setting $\omega_n = 1$ in every period follows analogously to the same statement in the proof of Theorem 3, simplifying the equation to

$$V_{n}^{\text{US}} = \max_{\psi_{n} \in \text{supp } F} \left(p \int_{\psi_{n}}^{\bar{u}} (x - \alpha V_{n+1}^{\text{US}}) \, dF(x) - \min\{z(\psi_{n}), z(\min \operatorname{supp} F)\} + \alpha V_{n+1}^{\text{US}} \right).$$
(EC.36)

Note that this reduces the problem to a univariate one: the choice of acceptance policy. Denoting $\operatorname{ob}(\psi, n) := p \int_{\psi}^{\bar{u}} \left(x - \alpha V_{n+1}^{\mathrm{US}}\right) dF(x) - \min\{z(\psi), z(\min\operatorname{supp} F)\} + \alpha V_{n+1}^{\mathrm{US}}$, it is easily verifiable that $\operatorname{ob}(\psi, n)$ is submodular, thus by Topkis's theorem the optimal $\{\psi_n^{\mathrm{US}}\}$ is weakly decreasing in n. From the shape of the cost function z given in the proof of Lemma 1 (it is quasi-convex with minimum at ξ^{LC} , such that it is strictly decreasing on $[\underline{u}, \xi^{\mathrm{LC}}] \cap \operatorname{supp} F$ and strictly increasing and unbounded for $\xi_n > \xi^{\mathrm{LC}}$), it follows that $\exists \psi^{\mathrm{crit}} \in \operatorname{supp} F$ s. t. $\psi^{\mathrm{crit}} \ge \xi^{\mathrm{LC}}, z(\psi^{\mathrm{crit}}) \ge z(\min\operatorname{supp} F)$ and $\forall \psi \in \operatorname{supp} F$ such that $\psi > \psi^{\mathrm{crit}} : z(\psi) > z(\min\operatorname{supp} F)$, and $\forall \psi \in \operatorname{supp} F$ s. t. $\psi < \psi^{\mathrm{crit}} : z(\psi) < z(\min\operatorname{supp} F)$. Thus, there exists $0 \le K \le N$ such that for K periods the principal induces policy $\xi_n = \min\operatorname{sup} F$ in the agent, while in the remaining periods she induces ψ_n^{US} . Parts (i) and (ii) of the theorem then follow from Proposition A1, which gives the optimal contracts for implementing those policies. \Box

Proof of Proposition 3.

Consider the setting where agent always reports truthfully, i.e., $\mathcal{R}_n(u) = u, \forall u, n$. This eliminates the adverse selection dimension of the contracting problem, but the moral hazard part remains. It is then optimal for the principal to never evaluate the agent's reports ($\phi_n(x) = 0, \forall x, n$) and from Theorem 1 of Plambeck and Zenios (2000) it follows that the problem of finding a contract that implements any admissible search policy $\{\omega_n\}$ at the lowest possible cost to the principal can be decomposed into N single-period problems, each given by

$$\underset{s_n^A(u), s_n^R(u), s_n^0}{\operatorname{arg\,min}} \quad p\omega_n \sum_{u \in \Psi_n} \left(s_n^A(u)\mu(u) \right) + p\omega_n \sum_{u \in \operatorname{supp} F \setminus \Psi_n} \left(s_n^R(u)\mu(u) \right) + (1 - p\omega_n) s_n^0; \tag{EC.37}$$

$$-p\omega_{n}\sum_{u\in\Psi_{n}}\left(\exp[-r(1-\alpha)(s_{n}^{A}(u)-\alpha^{-1}c)]\mu(u)\right) - p\omega_{n}\sum_{u\in\operatorname{supp}F\setminus\Psi_{n}}\left(\exp[-r(1-\alpha)(s_{n}^{R}(u)-\alpha^{-1}c)]\mu(u)\right) - (1-p\omega_{n})\exp[-r(1-\alpha)(s_{n}^{0}-\alpha^{-1}c)] = -\exp[-r\alpha^{-1}(1-\alpha)w];$$
(EC.38)
$$-\exp[-r(1-\alpha)s_{n}^{0}] \leq -\exp[-r\alpha^{-1}(1-\alpha)w].$$
(EC.39)

Notice that the system above is equivalent to the system (EC.1),(EC.2),(EC.4) (solved as part of the proof for Theorem 1) when $\Xi_n = \operatorname{supp} F$. Thus, the solution is also the same: for $\omega_n = 1$ the solution exists only if $p > 1 + v(c/\alpha)$ and it is to pay the agent a bonus $s_n^B(\operatorname{supp} F)$ as given by (11) when any alternative is reported, or $s_n^0 = w/\alpha$ when agent reports that no alternative was found. For $\omega_n = 0$, the solution is to pay the agent w/α irrespective of the report. By Theorem 3, the same contract is optimal in the Section 2 model when e = 0, and the principal's value function in that setting is equal to V_n^T .

Now we will demonstrate how to achieve the same with an agent which is not necessarily truthful. Consider a principal who conducts evaluations with frequency ϕ . Applying Theorem 1 of Plambeck and Zenios (2000) it follows that the problem of finding a contract that implements any admissible search policy $\{\omega_n\}$ and a truthful reporting policy $(\mathcal{R}_n(u) = u, \forall u, n)$ at the lowest possible cost to the principal can be decomposed into N single-period problems, each given by

$$\underset{s_n(x,u)}{\operatorname{arg\,min}} \left[(1 - \omega_n p) \Big(\phi s_n(\emptyset, \emptyset) + (1 - \phi) s_n(\emptyset, \operatorname{NE}) \Big) + \omega_n p \sum_{u \in \operatorname{supp} F} \Big(\phi s_n(u, u) + (1 - \phi) s_n(u, \operatorname{NE}) \Big) \mu(u) \right]; \quad (\operatorname{EC.40}) \\
(1 - \omega_n p) \Big(\phi v(s_n(\emptyset, \emptyset)) + (1 - \phi) v(s_n(\emptyset, \operatorname{NE})) \Big) + \omega_n p \sum_{u \in \operatorname{supp} F} \Big(\phi v(s_n(u, u)) + (1 - \phi) v(s_n(u, \operatorname{NE})) \Big) \mu(u) \\
= -v(w/\alpha)/v(-c/\alpha);$$

$$p \sum_{u \in \text{supp } F} \left(\phi v(s_n(\mathcal{R}(u), u)) + (1 - \phi) v(s_n(\mathcal{R}(u), \text{NE})) \right) \mu(u)$$
(EC.42)

$$+ (1-p) \Big(\phi v(s_n(\mathcal{R}(\emptyset), \emptyset)) + (1-\phi) v(s_n(\mathcal{R}(\emptyset), \operatorname{NE})) \Big)$$
(EC.42)

$$\leq -v(w/\alpha)/v(-c/\alpha) \quad \forall \mathcal{K} : \operatorname{supp} F \cup \{\emptyset\} \to \operatorname{supp} F \cup \{\emptyset\};$$

$$\phi v(s_n(\emptyset, \emptyset)) + (1 - \phi)v(s_n(\emptyset, \operatorname{NE})) \leq v(w/\alpha).$$
(EC.43)

Here, $s_n(x, u)$ denotes the payment to the agent in the scenario where he reports that an alternative with value x was found (or that no alternative was found, in which case $x = \emptyset$), and the principal's evaluation finds the value of the alternative to be u (u = NE if the principal did not conduct an evaluation). Note that the formulation (EC.40)-(EC.43) does not allow for payment to differ based on the principal's decision to accept or reject the alternative. This restriction is made without loss of optimality, which follows from the same Jensen's inequality argument as in the proof of Theorem 1. (Intuition: doing otherwise would impose an additional source of risk on the agent which is outside his control, necessitating a higher payment in

(EC.41)

expectation.) The objective function (EC.40) minimizes the principal's expected costs, subject to the binding participation constraint (EC.41), and the incentive compatibility constraints (EC.42)-(EC.43). For $\omega_n = 0$, (EC.40)-(EC.43) has a trivial solution of $s(x, u) = w/\alpha, \forall x, u$.

For $\omega_n = 1$, consider a contract which pays $s_n(u, u) = s_n(u, \text{NE}) = s_n^B(\text{supp } F), \forall u \in \text{supp } F$ (this is the bonus s_n^E given in part (i) of the proposition) and $s_n(\emptyset, \emptyset) = s_n(\emptyset, \text{NE}) = w/\alpha$ (this is the fixed pay given in part (iii) of the proposition), where $s_n^B(\text{supp } F)$ is given by (11). Notice that this contract solves the problem (EC.40),(EC.41),(EC.43), irrespective of how $s_n(x, u)$ are defined in the case where $x \neq u$ and $u \neq \text{NE}$. This is essentially the same solution as for (EC.37)-(EC.39), which ensures that the agent is incentivized to search if he is truthful, but not necessarily that truthfulness is incentive compatible, as it ignores (EC.42).

We will now show that this contract can be made incentive compatible, so also to satisfy (EC.42), by properly defining $s_n(x, u)$ in the case where $x \neq u$. Consider the point-wise incentive compatibility constraint

$$\phi v(s_n(u,u)) + (1-\phi)v(s_n(u,\text{NE})) \ge \phi v(s_n(x,u)) + (1-\phi)v(s_n(x,\text{NE})), \quad \forall x, u \in \text{supp} \, F \cup \emptyset, \text{ s.t. } x \neq u.$$
(EC.44)

This constraint ensures that for every possible u the agent gets higher expected utility by honestly reporting it (the LHS of of (EC.44)), than by making any dishonest report x (the RHS). Notice that the condition (EC.44) implies (EC.42). Setting $s_n(x,u) = s_n^P, \forall x \neq u, u \neq \text{NE}$, where s_n^P is given by (21), ensures that (EC.44) holds. This is the reduced pay if caught making dishonest reports, as given in part (ii) of the proposition. Because (EC.44) is satisfied, this contract solves (EC.40)-(EC.43) and is thus optimal.

Finally, the identity $\lim_{\phi\to 0} V_n^{\text{ARV}}(x;\phi) = V_n^T$ follows from comparing the solution of (EC.37)-(EC.39) to the one of (EC.40)-(EC.43) and noticing that these contracts induce the same actions and pay the same amount to the agent; yet in the second of these two cases the principal bears an additional cost from random evaluations, but this cost vanishes as $\phi \to 0$. \Box

Lemma A4. Let the agent hold contract $\{s_n(x)\} \in \mathcal{F} \cup \mathcal{P}$ and the principal use threshold acceptance policy $\{\psi_n\}$. Then, it is optimal for the agent to use policies $\{\tilde{\gamma}_n\}, \{\tilde{\omega}_n\}, \{\tilde{\xi}_n\}$ which solve the recursion:

$$\begin{split} U_n^A(\{s_n(x)\},\{\psi_n\}) &= \max_{\gamma_n,\mathbb{R},\xi_n \in \{\xi \in \text{supp} F | \xi \geq \psi_n\},\omega_n \in \{0,1\}} - \exp[-r\gamma_n] \\ &+ \omega_n p \alpha \int_{\xi_n}^{\bar{u}} \exp[-r(1-\alpha)(s_n(x) - (\gamma_n + c)/\alpha)] dF(x) U_{N+1}^A(\{s_n(x)\},\{\psi_n\}) \\ &+ (1 - \omega_n p \bar{F}(\xi_n)) \alpha \exp[-r(1-\alpha)(-(\gamma_n + \omega_n c)/\alpha)] U_{n+1}^A(\{s_n(x)\},\{\psi_n\}) \end{split}$$

with boundary condition

$$U_{N+1}^{A}(\{s_{n}(x)\},\{\psi_{n}\}) = -\frac{1}{1-\alpha}\exp(-rw).$$

If $\{s_n(x)\} \in \mathcal{F}$, the agent reports only those alternatives that the principal is willing to accept $(\{\tilde{\xi}_n\} = \{\psi_n\})$. *Proof.* Follows directly from applying Lemma 1 of Plambeck and Zenios (2000) to the agent's value function (5). \Box

Lemma A5. Let $\mathcal{C} \in \{\mathcal{F}, \mathcal{P}\}$ and $\{s_n(x)\} \in \mathcal{C}$. It will then be optimal for the principal to use a threshold acceptance policy $\psi_n(\{s_n(x)\}) = \stackrel{\sim}{\cap} \alpha V_{n+1}^P(\{s_n(x)\}) + \alpha s \cap if \mathcal{C} = \mathcal{F}$, or $\psi_n(\{s_n(x)\}) = \stackrel{\sim}{\cap} \alpha V_{n+1}^P(\{s_n(x)\})/(1 - 1)$

 αq) \neg if $\mathscr{C} = \mathscr{P}$. Here, $V_n^P(\{s_n(x)\})$ is the principal's value function. Denoting by $\{\tilde{\omega}_n\}$ the agent's optimal search policy, as given by Lemma A4, $V_n^P(\{s_n(x)\})$ is given recursively by

$$V_n^P(\{s_n(x)\}) = \alpha V_{n+1}^P(\{s_n(x)\}) + p\tilde{\omega}_n \int_{\psi_n(\{s_n(x)\})}^{\tilde{u}} \left(x - \alpha s_n(x) - e - \alpha V_{n+1}^P(\{s_n(x)\})\right) dF(x),$$

with boundary condition $V_{N+1}^P(\{s_n(x)\}) = 0.^{22}$ The optimal contract then solves

$$\max_{\{s_n(x)\}\in\mathscr{C}} V_1^P(\{s_n(x)\}) \tag{EC.45}$$

s.t.
$$U_0^A(\{s_n(x)\}, \{\alpha V_{n+1}^P(\{s_n(x)\})\}) \ge -\frac{1}{1-\alpha}\exp(-rw).$$
 (EC.46)

This optimization reduces to optimizing over $s \in \mathbb{R}$ if $\mathscr{C} = \mathscr{F}$ or optimizing over $q \in [0,1]$ if $\mathscr{C} = \mathscr{P}$.

Proof. Let $\{s_n(x)\} \in \mathcal{C}$ (resp. $\{s_n(x)\} \in \mathcal{F}$). Each time the principal is presented with an alternative of value x, the evaluation cost is already sunk when she finds out x, which leaves the principal with two options: a) accept the alternative, receiving its value x, but paying a bonus s (resp. qx) in the beginning the of next period, or b) reject it, in which case the search proceeds which will yield $\alpha V_n^P(\{s_n(x)\})$ in expectation. Optimality of threshold acceptance policy $\psi_n(\{s_n(x)\}) = \mathbb{P} \alpha V_{n+1}^P(\{s_n(x)\}) + \alpha s \, \forall, \forall n \text{ (resp. } \psi_n(\{s_n(x)\}) = \mathbb{P} \alpha V_{n+1}^P(\{s_n(x)\}) / (1 - \alpha q) \, \forall, \forall n)$ directly follows. The expression for the principal's value function then follows by inserting the optimal $\{\psi_n\} = \{\xi_n\}$ into (4). The system (EC.45)-(EC.46) is then a simplified version of the contracting problem (6)-(9), which follows from inserting the expressions for optimal policies for any given $\{s_n\}$ (as given by Lemma A4 and this proof) and from applying Lemma 1 of Plambeck and Zenios (2000) to replace the agent's participation constraint by its wealth-independent component. \Box

Proof of Proposition 4

Part 1: $V_n^{\mathfrak{F}}$, $V_n^{\mathfrak{F}} \leq V_n^{US}$. Consider two different principals, each with their own agent. The first principal's agent holds $\{s_n(x)\} \in \mathfrak{F} \cup \mathfrak{F}$. Denote by $\{\tilde{\psi}_n\}$ the first principal's optimal acceptance policy and by V_n^1 her value function (both given by Lemma A5). Also denote the optimal searching and reporting policy for this principal's agent by $\{\tilde{\omega}_n\}$ and $\{\tilde{\xi}_n\}$ (both given by Lemma A4).

The second principal commits to the same acceptance policy $\{\tilde{\psi}_n\}$ and desires to induce the agent to use the same reporting and search policies $\{\tilde{\omega}_n\}, \{\tilde{\xi}_n\}$ but can optimize over which contract $\{s_n(Y^n)\}$ to use, without being bound to a particular contract class (thus the second principal is in the setting of Section 4.1). Then, denoting the second principal's value function by V_n^2 , it follows from Proposition A1 that if $\{s_n(x)\} \in \mathcal{F}$, then $V_n^2 \geq V_n^1, \forall n$ as in any single period, the contract of Proposition A1 is the most cost-efficient way of inducing those policies, optimized over the space of all contracts, including the ones in \mathcal{F} .

However, the percentage-bonus contracts of \mathscr{P} require contracting on the value of the alternative, which the second principal is not able to do (as Y^n is his sole contractable variable, and it does not include that information). Thus, to show $V_n^2 \ge V_n^1$ in the case where $\{s_n(x)\} \in \mathscr{P}$, an additional step is needed: to construct a contract $\{s_n(Y^n)\}$ (in the setting of Proposition A1) which induces $\{\tilde{\omega}_n\}$ and $\{\tilde{\xi}_n\}$ (so the same policies as $\{s_n(x)\} \in \mathscr{P}$ with percentage bonus q), and yields at least V_n^1 to the principal. To do that, let the

²² Note that $\tilde{\omega}_n$ here depends on $\{\psi_n\}$ and $\{s_n(x)\}$, thus this recursion needs to be solved jointly with the agent's problem in Lemma A4.

principal instead commit to acceptance policy $\{\tilde{\psi}_n^*\}$ given by $\tilde{\psi}_n^* = \max\{\tilde{\psi}_n, \tilde{\xi}_n\}, \forall n \in \{1, ..., N\}$, and use a contract in $\{s_n^*(x)\} \in \mathcal{F}$ with fixed bonus size s^* that solves

$$\int_{\tilde{\psi}_{n}^{*}}^{\bar{u}} \exp[-r(1-\alpha)qx] dF(x) = \int_{\tilde{\psi}_{n}^{*}}^{\bar{u}} \exp[-r(1-\alpha)s^{*}] dF(x).$$
(EC.47)

Solving it yields $s^* = -\frac{1}{r(1-\alpha)} \ln \left(\int_{\tilde{\psi}_n^*}^{\tilde{u}} \exp[-r(1-\alpha)qx] dF(x)/\bar{F}(\tilde{\psi}_n^*) \right)$. Then, applying Lemma A4, because of (EC.47), $\{s_n^*(x)\}$ also induces $\{\tilde{\omega}_n\}$ and $\{\tilde{\xi}_n\}$ and yields the same value to the agent as $\{s_n(x)\}$. However the expected cost for the principal in period n under $\{s_n(x)\}$ is $\tilde{\omega}_n p \int_{\tilde{\psi}_n^*}^{\tilde{u}} qx dF(x)$ while under $\{s_n^*(x)\}$ it is $\tilde{\omega}_n ps^*\bar{F}(\tilde{\psi}_n^*)$. Here, from concavity of $\exp[-r(1-\alpha)x]$ in x and (EC.47) we have that $\int_{\tilde{\psi}_n^*}^{\tilde{u}} qx dF(x) \ge s^*\bar{F}(\tilde{\psi}_n^*)$ (as for every concave u, random variable X, and constant y, it holds that $\mathbb{E}[u(X)] = u(y) \Rightarrow \mathbb{E}[X] \ge y$), Thus, denoting the principal's value function under $\{s_n^*(x)\}$ by V_n^{1*} , we have that $V_n^2 \ge V_n^{1*} \ge V_n^1, \forall n.^{23}$

Lastly, because $V_n^2 \ge V_n^1, \forall n$ holds in both cases, that the truly optimal contract yields an even higher expectation $(V_n^{US} \ge V_n^2 \ge V_n^1, \forall n)$ follows from Bellman's principle of optimality.

Part 2: $V_n^P - V_n^{\mathcal{F}}$ is unbounded. The idea of the proof is to construct a situation in which no contract in \mathcal{F} can create positive value for the principal, but the optimal contract can generate unboundedly high value. Consider a single period model (N = 1) where $\underline{u} = \min \operatorname{supp} F = 0$, max supp $F < \overline{u}$, p = 1, e = 0, and w = 0. Under a contract which pays a fixed bonus *s*, the principal will use threshold acceptance policy $\psi_1 = \overrightarrow{r} \alpha s \ \neg$ and the agent will use reporting policy $\xi_1 = \overrightarrow{r} \alpha s \ \neg$ (by Lemmas A5 and A4 respectively). As there is no opportunity cost, this contract will always satisfy the participation constraint. Using Proposition 5 of Plambeck and Zenios (2000), we can more concisely formulate the agent's search decision: he should search if and only if the following inequality holds

$$-\bar{F}(\alpha s)\exp(-r(1-\alpha)(s-c/\alpha)) - (1-\bar{F}(\alpha s))\exp(-r(1-\alpha)(-c/\alpha)) \ge -1.$$
 (EC.48)

Here the first term is the agent's instantaneous utility of paying search cost c but then finding an alternative which results in bonus s, the second term is the scenario where search cost is expanded but no suitable alternative is found, and the RHS is the utility from not searching. Rearranging (EC.48) yields

$$\bar{F}(\alpha s)\exp(-r(1-\alpha)s) \le -(1-\bar{F}(\alpha s)) + \frac{1}{\exp(-r(1-\alpha)(-c/\alpha))}.$$

Here, the LHS is positive, so the RHS also needs to be positive, which yields the following necessary (but not sufficient) condition for the contract to induce the agent to search:

$$1 - \bar{F}(\alpha s) \le \frac{1}{\exp(-r(1-\alpha)(-c/\alpha))}.$$
(EC.49)

Define $c^* := \max\{c \in \mathbb{R}^+ | 1 - \bar{F}(\alpha c) \le 1/\exp(-r(1-\alpha)(-c/\alpha))\}$. Existence and uniqueness of c^* is guaranteed by $1 - \bar{F}(\alpha c)$ being an weakly increasing left-continuous function with $1 - \bar{F}(0) = 0$ and $1 - \bar{F}(\bar{u}) = 1$, while $1/\exp(-r(1-\alpha)(-c/\alpha))$ is strictly decreasing, continuous, and evaluates to 1 at c = 0. Let the agent's costs be c^* . Then, setting the bonus to any amount $s > c^*$ will violate the condition (EC.49) and thus induce the

²³ From this construction, it may appear that the best contract in \mathcal{F} outperforms all contracts in \mathcal{P} (so $V_n^{\mathcal{F}} \ge V_n^{\mathcal{P}}$). However, this is not necessarily true (a counterexample is provided in Figure 4), as attaining V_n^{1*} also requires commitment to an acceptance policy whereas attaining $V_n^{\mathcal{F}}$ does not.

agent not to search, but setting $s \leq c^*$ will trivially violate (EC.48) (bonus is not enough to recoup even the cost of search), thus also induce the agent not to search. Consequently, $V_n^{\mathcal{F}} = 0$.

Now consider $F_y(\cdot)$ such that $F_y(u) = F(u), \forall u \leq c^*$ and the expectation of drawing from that distribution is $y > c^*$. Notice that the optimal contract s^* given by Theorem 3 remains the same if the agent is drawing from F_y instead of F, and the $V_n^{\mathcal{F}} = 0$ property is preserved. Applying Theorem 3, the optimal contract in this case is to pay the agent c^*/α for any alternative *delivered*, rather than only accepted alternatives. Using this contract gives the principal expectation of $y - c^*$ which can be unboundedly high as y increases.

Part 3: $V_n^P - V_n^{\mathcal{P}}$ is unbounded. As in part 2, we aim to construct a situation in which no contract in \mathcal{P} can create positive value for the principal, but the optimal contract can generate unboundedly high value. Consider a single period model (N = 1) where $\underline{u} = \min \operatorname{supp} F = 0$, p = 1, e = 0, and w = 0. Applying Lemmas A4 and A5, under any contract in \mathcal{P} we have $\xi_1 = \psi_1 = 0$. Analogously to (EC.48), it will be optimal for the agent holding such contract to search if and only if the following equation holds:

$$-\int_0^{\bar{u}} \exp(-r(1-\alpha)(qu-c/\alpha))dF(u) \ge -1.$$

Here the LHS is the instantaneous utility from searching, while the RHS is the utility of not searching. Now consider a contract $s_b(u)$, which pays the agent $c/(2\alpha)$ for delivery of any alternative with a value below c/2, and pays $b > \bar{u}$ for delivery of alternatives with a value at least equal to c/2. Note that this contract does not belong to \mathcal{P} , but is preferred by the agent over any contract in \mathcal{P} as the distribution of payouts to the agent when searching will dominate the distribution of payouts under all contracts in \mathcal{P} by first-order stochastic dominance. Under $\hat{s}(u)$, the agent should search if and only if

$$-\bar{F}(c/2)\exp(-r(1-\alpha)(b-c/\alpha)) - (1-\bar{F}(c/2))\exp(-r(1-\alpha)(-c/(2\alpha))) \ge -1$$

The utility on LHS in increasing in b, and we can finds its supremum by considering $\lim_{b\to\infty}$ which yields

$$-(1 - \bar{F}(c/2)) \exp(-r(1 - \alpha)(-c/(2\alpha))) \ge -1,$$

$$1 - \bar{F}(c/2) \le 1/\exp(r(1 - \alpha)c/(2\alpha)).$$
 (EC.50)

This condition is of interest to us, as the agent not being incentivized to search here will also imply that no contract in \mathscr{P} will incentivize him to search either. Because the RHS of (EC.50) is <1, there are many distributions for which this condition is not satisfied. So, let $F(\cdot)$ be one such distribution for which $1 - \bar{F}(c/2) > 1/\exp(r(1-\alpha)c/(2\alpha))$. Then, no matter the value of $b, s_b(u)$ will incentivize the agent not to search; no contract in \mathscr{P} will incentivize him to search either; and thus we have $V_n^{\mathscr{P}} = 0$.

The intuition for why this occurs is as follows: the scenarios where u < c/2 create a loss (and the agent is loss averse) as the payment is not enough to recoup cost c, the situations where $u \ge c/2$ create a gain, but this gain is bounded as exponential utility is bounded on the right. Consequently, if the unfavorable scenario is common enough and favorable scenario is rare enough, there is no payment for the favorable scenario that is high enough to make this an attractive prospect. (The same phenomenon is the driver of implementability condition (10) in Theorem 1.)

Now consider a distribution $F_y(\cdot)$ such that $F_y(u) = F(u), \forall u \leq c/2$ and the expectation of drawing from that distribution is y > c/2. Notice that if (EC.50) does not hold for $F(\cdot)$ then it does not hold for $F_y(\cdot)$

either. Applying Theorem 3, the optimal contract in this case is to pay the agent c/α for any alternative *delivered*, rather than only accepted alternatives. Using the optimal contract yields the principal expectation y - c which can be unboundedly high as y increases. Thus, when the distribution is $F_y(u)$, we have that V_n^P can be unboundedly high (as $y \to \infty$) but $V_n^{\mathcal{P}} = 0$. \Box

Appendix C: Additional results

C.1. Discretization techniques

Here, we illustrate bracket-mean and bracket-median, which are two commonly used discrete approximations of continuous distributions (Miller III and Rice 1983, Smith 1993); these algorithms are also used in computer science where they are referred to as equal-frequency binning (Kotsiantis and Kanellopoulos 2006). Denote by $G(\cdot)$ a continuous probability distribution with support $[\underline{u}, \overline{u}]$, which we will approximate with a discrete distribution $F_n(\cdot)$. The bracket-mean technique starts with dividing $G(\cdot)$ into n equally probable intervals. Denoting $x_0^n := \underline{u}, x_n^n := \overline{u}$, the separation into intervals consists of finding points $x_1^n, x_2^n, ..., x_{n-1}^n$ such that $G(x_i^n) - G(x_{i-1}^n) = 1/n$, for all $i \in \{1, ..., n\}$. The discrete approximation $F_n(\cdot)$ is then given as the one with n elements in its support, each being equal to the mean of one of the n just-defined subintervals of G(x), and each having a probability mass of 1/n. Formally, for a random variable X distributed according to $G(\cdot)$, for every $i \in \{1, ..., n\}$, let $y_i^n := \mathbb{E}(X | x_{i-1}^n \le X \le x_i^n)$; the discrete approximation $F_n(\cdot)$ is then defined by its support $\{y_1^n, ..., y_n^n\}$ and a constant pmf $f_n(x) := 1/n$.

We can also define the discretization procedure via the discretization function $D_n: [\underline{u}, \overline{u}] \to [\underline{u}, \overline{u}]$, given by

$$D_n(x) := \begin{cases} y_1^n & \text{if } x \in [x_0^n, x_1^n) \\ y_2^n & \text{if } x \in [x_1^n, x_2^n) \\ \dots & \dots \\ y_n^n & \text{if } x \in [x_{n-1}^n, x_n^n]. \end{cases}$$

The bracket-median technique is almost identical, differing only the way y_i are defined: in bracket-median, they are medians (rather than means) of each of the *n* intervals. Importantly for our purposes, both of these approximation techniques yield discrete approximations that satisfy our assumption that all elements in the support of the distribution of alternatives have the same probability mass. In the lemma below, we show that for both these approximation techniques, the approximation can be made arbitrarily close to the true distribution by increasing *n*.

Lemma A6 (Convergence of approximations) Let $\{F_n\}$ be a sequence of increasingly fine approximations of a G obtained through either bracket-mean or bracket-median. This sequence converges in distribution to G: $\lim_{n\to\infty} F_n(x) = G(x), \forall x \in [\underline{u}, \overline{u}]$. Let X be a random variable with cdf G. Then, $D_n(X)$ has cdf F_n and the sequence $\{D_n(X)\}$ converges to X in the L^1 norm: $\lim_{n\to\infty} \mathbb{E}[|D_n(X) - X|] = 0$.

Proof Sketch. Denote $x_{y\downarrow}^n := \max\{x \in \{x_1^n, x_2^n, ..., x_n^n\} | x \le y\}$ and $x_{y\uparrow}^n := \min\{x \in \{x_1^n, x_2^n, ..., x_n^n\} | x \ge y\}$. From the construction of F_n , we have that $G(x_{y\downarrow}^n) \le F_n(y) \le G(x_{y\uparrow}^n)$ and $G(x_{y\downarrow}^n) \le G(y) \le G(x_{y\uparrow}^n), \forall y \in [\underline{u}, \overline{u}]$. Then, $\lim_{n\to\infty} G(x_{y\uparrow}^n) - G(x_{y\downarrow}^n) = \lim_{n\to\infty} 1/n = 0$ so $\lim_{n\to\infty} F_n(y) = G(y)$, thus convergence in distributions. To show L^1 convergence, we have $0 \le \lim_{n\to\infty} \mathbb{E}[|D_n(X) - X|] \le \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^n (x_i - x_{i-1}) = \lim_{n\to\infty} (\overline{u} - \underline{u})/n = 0$. \Box

C.2. Continuous distribution of alternatives

Consider a version of our main model (as in Section 2), such that the distribution of alternatives $F(\cdot)$ is replaced by a continuous distribution $G(\cdot)$, with supp $G = [\underline{u}, \overline{u}]$. Thus, in this model, P(S) is the probability of drawing an outcome in set S, when drawing from a distribution with cdf G. This change will cause the state variable X_n will become uncountably infinite, which gives rise to measurability-related technical issues. We can ensure our model is well-behaved and eliminate the measurability related problems by tightening some of the other assumptions. Thus, assume that \mathcal{F} (the choice set for Ξ_n and Ψ_n) is the set all sets that can be formed by finite unions of closed intervals in $[\underline{u}, \overline{u}]$, and that for every n, $s_n(X^n)$ is bounded and almost-everywhere continuous.²⁴ A deeper issue still remains in that the uncountably-infinite state space violates the assumptions of several papers whose results we leverage in the main paper: Plambeck and Zenios (2000), Fudenberg et al. (1990), and Smith (1998). Yet, there are considerable results we can obtain *without* relying on that literature. First, the proposition below serves as an analogue of Theorem 1 for this setting.

Proposition A2 (Implementable policies under a continuous distribution) The principal can implement any policy pair $(\{\omega_n\}, \{\Xi_n\})$, provided it satisfies the implementability condition (10), for all $n \in \{1, ..., N\}$. This can be done by offering the agent a contract that pays a base pay of s_n^0 in every period, replaced by bonus $s_n^B(\Xi_n)$ if the agent delivers an alternative with a value in Ξ_n in a period in which the principal would like him to search ($\omega_n = 1$), where s_n^0 and $s_n^B(\Xi_n)$ are given by (11). Under any such contract, the agent's optimal consumption policy is given by $\gamma_n(X^{n-1}, W_n) = (1 - \alpha)W_n + w$, for all $n \in \{1, ..., N\}$. In all periods, the agent's value function under such contract is $V_n^A(W_n, X^{n-1}) = -\frac{1}{1-\alpha} \exp[-r(w+(1-\alpha)W_n)]$, rendering the agent indifferent between termination and continuation.

Proof Sketch. Consider the agent holding the contract given by this proposition, i.e., one designed to induce policies $(\{\omega_n\}, \{\Xi_n\})$. From (5) we have that the agent's optimal policies solve the recursion

$$\begin{aligned} V_n^A(W_n, X^{n-1}) &= \max_{\bar{\Xi}_n, \bar{\omega}_n, \bar{\gamma}_n} - \exp[-r\bar{\gamma}_n] \\ &+ \bar{\omega}_n \omega_n p P\left(\bar{\Xi}_n \cap \Xi_n(X^{n-1}) \cap \Psi_n(X^{n-1})\right) \left(-\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - c - \bar{\gamma}_n}{\alpha} + s_n^B(\Xi_n(X^{n-1}))\right)\right)\right)\right]\right) \\ &+ \bar{\omega}_n \omega_n p P\left(\bar{\Xi}_n \cap \Xi_n(X^{n-1}) \setminus \Psi_n(X^{n-1})\right) \alpha V_n^A \left(\frac{W_n - c - \bar{\gamma}_n}{\alpha} + s_n^B(\Xi_n(X^{n-1}))\right) \\ &+ \bar{\omega}_n \omega_n p P\left(\bar{\Xi}_n \cap \Psi_n(X^{n-1}) \setminus \Xi_n(X^{n-1})\right) \left(-\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - c - \bar{\gamma}_n}{\alpha} + s_n^0\right)\right)\right)\right]\right) \\ &+ \left(1 - \bar{\omega}_n \omega_n p P(\bar{\Xi}_n \cap (\Psi_n(X^{n-1}) \cup \Xi_n(X^{n-1}))\right) \alpha V_n^A \left(\frac{W_n - c \bar{\omega}_n - \bar{\gamma}_n}{\alpha} + s_n^0\right), \end{aligned}$$
(EC.51)

with boundary condition $V_{N+1}^A(W_{N+1}, X^N) = -\frac{1}{1-\alpha} \exp[-r(w + (1-\alpha)W_{N+1})]$. Note that s_n^B is well-defined if and only if (10) holds. The proof is based on showing by induction that $V_n^A(W_n, X^{n-1}) = -\frac{1}{1-\alpha} \exp[-r(w + (1-\alpha)W_n)]$ (this will also make the participation constraint satisfied and binding). All of the other results –

²⁴ The measurability issues can also be resolved using weaker assumptions by following Bertsekas and Shreve (2004), i.e., through the use of either outer integration or by constraining all policy choices to universally measurable ones. Using that approach requires introduction of considerable addition notation and a relatively heavy measure-theoretic apparatus.

i.e., that this contract really induces the desired policies $(\{\omega_n\}, \{\Xi_n\})$ as well as $\gamma_n(X^{n-1}, W_n) = (1-\alpha)W_n + w$ – are obtained as steps in this proof. The boundary condition of (EC.51) also serves as the induction base. For induction step, assume $V_{n+1}^A(W_{n+1}, X^n) = -\frac{1}{1-\alpha} \exp[-r(w + (1-\alpha)W_{n+1})]$, using which, from (EC.51) we have

$$V_n^A(W_n, X^{n-1}) = \max_{\bar{\Xi}_n, \bar{\omega}_n, \bar{\gamma}_n} -\exp[-r\bar{\gamma}_n] + \bar{\omega}_n \omega_n p P\left(\bar{\Xi}_n \cap \Xi_n(X^{n-1})\right) \left(-\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - c - \bar{\gamma}_n}{\alpha} + s_n^B(\Xi_n(X^{n-1}))\right)\right)\right)\right]\right)$$
(EC.52)
+ $\left(1 - \bar{\omega}_n \omega_n p P\left(\bar{\Xi}_n \cap \Xi_n(X^{n-1})\right)\right) \left(-\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - \bar{\omega}_N c - \bar{\gamma}_n}{\alpha} + s_n^0\right)\right)\right]\right).$

Here, $\bar{\Xi}_n = \Xi_n[X^{n-1}]$ is optimal as V_n^A depends on the choice of $\bar{\Xi}_n$ only through $P\left(\bar{\Xi}_n \cap \Xi_n(X^{n-1})\right)$, and is increasing in that probability as $s_n^B(\Xi_n(X^{n-1})) > s_n^0$ by (11). If $\omega_n = 0$, the value function becomes $V_n^A(W_n, X^{n-1}) = \max_{\bar{\omega}_n, \bar{\gamma}_n} - \exp[-r\bar{\gamma}_n] + (-\alpha(1-\alpha)^{-1}\exp[-r(w+(1-\alpha)((W_n-\bar{\omega}_n c-\bar{\gamma}_n)/\alpha+s_n^0))]))$, where $\bar{\omega}_n = 0$ is optimal (it only appears in the $-\bar{\omega}_n c$ term), reducing the equation to

$$V_n^A(W_n, X^{n-1}) = \max_{\bar{\gamma}_n} - \exp\left[-r\bar{\gamma}_n\right] + \left(-\frac{\alpha}{1-\alpha} \exp\left[-r\left(w + (1-\alpha)\left(\frac{W_n - \bar{\gamma}_n}{\alpha} + s_n^0\right)\right)\right]\right), \quad (EC.53)$$

Solving (EC.53) for $\bar{\gamma}_n$ yields $\bar{\gamma}_n = (1 - \alpha)W_n + w$; plugging that back into (EC.53) gives $V_n^A(W_n, X^{n-1}) = -\frac{1}{1-\alpha} \exp[-r(w + (1-\alpha)W_n)].$

Now we turn to the case where $\omega_n = 1$. Here, from (EC.52), not searching also yields the RHS of (EC.53), so $-\frac{1}{1-\alpha} \exp[-r(w + (1-\alpha)W_n)]$. Expanding the expressions for bonus sizes in (EC.52) using (11) and simplifying the resulting expression yields that searching (choosing $\bar{\omega}_n = 1$) also gives the agent the RHS of (EC.53). So searching is optimal (weakly, as the agent is indifferent), the optimal consumption is $\bar{\gamma}_n =$ $(1-\alpha)W_n + w$ and $V_n^A(W_n, X^{n-1}) = -\frac{1}{1-\alpha} \exp[-r(w + (1-\alpha)W_n)]$ in this case as well, completing the induction. \Box

Proposition A2 shows that the contracts of Theorem 1 are incentive compatible and individually rational in this setting as well. They are also subject to same implementability condition, induce the same consumption policy, and have the same property that the agent's participation constraint is binding in all periods, not just at contract signing. The important gap from Theorem 1 is that, here, these contracts are not guaranteed to be the lowest cost ones over the space of all contracts (even existence of such minimum is not guaranteed in this setup).

The results of Proposition A2, combined with those that we can approximate the true distribution with arbitrarily fine precision with discrete approximations (shown in Lemma A6) and that for all the discrete approximations, the contract class of Proposition A2 is guaranteed to be optimal (shown in Theorem 1), provide a reasonable justification to focus on this contract class.

Then, as in (12), restricting the principal's choice of contracts to ones given in Proposition A2 reduces to the principal's problem to

$$V_n^P = \max_{\omega_n \in \{0,1\}, \Xi_n, \Psi_n \in \mathcal{F}} \omega_n p P(\Xi_n) \left(\alpha s_n^0 - e - \alpha s_n^B(\Xi_n) \right) + \omega_n p \int \left(u - \alpha V_{n+1}^P \right) dG(u) - \alpha s_n^0 + \alpha V_{n+1}^P, \quad (\text{EC.54})$$

with boundary condition $V_{N+1}^P = 0$. This is now a well-behaved dynamic program which is amenable to the same methods of analysis used in the main paper, thus our main results (Lemma 1, Theorem 2, Theorem 3, Proposition 1, and Proposition 2) should be obtainable in this setting as well, by following analogous proofs – or course, with the limitation that the principal be restricted in the form of contracts she can offer to the contracts of Proposition A2.

As a final reality check that the discrete approximation is reasonable, we show that as the precision of the approximation increases the principal's value function converges to the value function in the continuous version of the model.

Lemma A7 (Convergence of value functions) Let V_k^P be the value function of a principal with a continuous distribution of alternatives G with $\operatorname{supp} G = [\underline{u}, \overline{u}]$, when she uses acceptance policy $\{\Psi_k\}$ and induces the agent to use policies $\{\Xi_k\}, \{\omega_n\}$ using a contract given by Proposition A2. Let ${}^nV_k^P$ be the value function of a principal with a discrete distribution F_n (approximation of continuous distribution G obtained through either bracket-mean or bracket-median method, with n brackets), who uses acceptance policy $\{\Psi_k \cap \operatorname{supp} F_n\}$ and induces the agent to use policies $\{\Xi_k \cap \operatorname{supp} F_n\}, \{\omega_n\}$ using a contract given by Theorem 1. Then, $\lim_{n\to\infty} {}^nV_k^P = V_k^P$.

Proof Sketch. Proof is by induction. The basis ${}^{n}V_{N+1}^{P} = V_{N+1}^{P} = 0$ follows directly from boundary conditions. Assume by the way of induction that $\lim_{n\to\infty} {}^{n}V_{k+1}^{P} = V_{k+1}^{P}$. Denote by P(S;D) the probability of drawing an outcome in set S when drawing from a distribution with cdf D. Then, from (4) and Theorem 1 we have ${}^{n}V_{k}^{P} = \omega_{k}p \int_{\Xi_{k}\cap\Psi_{k}\cap\operatorname{supp} F_{n}} (u-e-\alpha s_{k}^{B}(\Xi_{k}\cap\operatorname{supp} F_{n})) dF_{n}(u) + \omega_{k}p \int_{(\Xi_{k}\setminus\Psi_{k})\cap\operatorname{supp} F_{n}} (-e-\alpha s_{k}^{B}(\Xi_{n}\cap\operatorname{supp} F_{n}) + {}^{n}V_{k+1}^{P}\alpha) dF_{n}(u) + (1-\omega_{n}pP(\Xi_{n};F_{n}))(-\alpha s_{k}^{0} + {}^{n}V_{k+1}^{P}\alpha) = \omega_{k}p \int_{\Xi_{k}\cap\Psi_{k}\cap\operatorname{supp} F_{n}} (u-{}^{n}V_{k+1}^{P}\alpha) dF_{n}(u) + \omega_{k}pP(\Xi_{k};F_{n})(-e-\alpha s_{k}^{B}(\Xi_{k}\cap\operatorname{supp} F_{n}) + \alpha s_{k}^{0}) - \alpha s_{k}^{0} + {}^{n}V_{k+1}^{P}\alpha.$ Taking a limit of this expression when $n \to \infty$ and using the convergence in distribution result of Lemma A6 yields $\lim_{n\to\infty} {}^{n}V_{k}^{P} = V_{k}^{P}$.

C.3. Variable search intensity model

Here we provide more detailed results for the model of Section 4.2, where the agent decides on search intensity. Proposition A3 is an analogue of Theorem 1 for this setting.

Proposition A3 Any implementable policy pair $(\{\omega_n\}, \{\xi_n\})$ satisfies the condition $p\bar{F}(\xi_n|\omega_n) > 1 - 1/\exp[rc(i_n)(1-\alpha)\alpha^{-1}]$ for all $n \in \{1, ..., N|\omega_n \neq 0\}$.²⁵ The contract which implements $(\{\omega_n\}, \{\xi_n\})$ at the lowest possible cost to the principal consists of a sequence of outcome-dependent payments. Such a contract pays $s_n^H(\omega_n, \xi_n)$ for period n if the agent delivers an alternative with a value of at least $\max\{\xi_n, \overline{u}_L\}$, pays $s_n^L(\omega_n, \xi_n)$ if the agent delivers an alternative with a value in $[\xi_n, \max\{\xi_n, \overline{u}_L\}]$ and pays s_n^0 in any other scenario. The exact values of those payments depend on the policy which they are designed to induce as follows:

(i) If $\omega_n = 0$, then $s_n^H(0, \xi_n) = s_n^L(0, \xi_n) = s_n^0 = \frac{w}{\alpha}$.

²⁵ This implementability condition used to be both sufficient and necessary in the basic model, but here it is only necessary, as there are additional conditions on implementability of i_H : either (EC.57) needs to hold or the argument of the logarithm in (EC.59) needs to be positive.

(ii) If $\omega_n = i_L$ then

$$s_n^L(i_L,\xi_n) = \frac{w+c_L}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{p\bar{F}(\xi_n|i_L)}{1 - (1 - p\bar{F}(\xi_n|i_L)) \exp[rc_L(1-\alpha)\alpha^{-1}]} \quad , \quad s_n^H(i_L,\xi_n) = s_n^0 = \frac{w}{\alpha}. \quad (\text{EC.55})$$

(iii) If $\omega_n = i_H$ then first consider the contract with parameters

$$s_n^L(i_H,\xi_n) = s_n^H(i_H,\xi_n) = \frac{w+c_H}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{p\bar{F}(\xi_n|i_H)}{1 - (1-p\bar{F}(\xi_n|i_H))\exp[rc_H(1-\alpha)\alpha^{-1}]} \quad , \quad s_n^0 = \frac{w}{\alpha}.$$
(EC.56)

This is the solution if the contract really induces i_H , that is, if it satisfies the constraint

$$\frac{\overline{u}_{H} + 1 - \xi_{n}}{\overline{u}_{H} + 1} v \left(s_{n}^{H}(i_{H}, \xi_{n}) - \frac{c_{H}}{\alpha} \right) + \frac{\xi_{n}}{\overline{u}_{H} + 1} s_{n}^{0} \ge \frac{(\overline{u}_{L} + 1 - \xi_{n})^{+}}{\overline{u}_{L} + 1} v \left(s_{n}^{H}(i_{H}, \xi_{n}) - \frac{c_{L}}{\alpha} \right) + \frac{\min\{\xi_{n}, \overline{u}_{L} + 1\}}{\overline{u}_{L} + 1} s_{n}^{0}.$$
(EC.57)

Otherwise, the solution is given by

$$s_n^L(i_H,\xi_n) = \frac{w+c_L}{\alpha} + \frac{1}{r(1-\alpha)} \ln \frac{p\bar{F}(\xi_n|i_L)}{1-(1-p\bar{F}(\xi_n|i_L))\exp[rc_L(1-\alpha)\alpha^{-1}]} \quad , \quad s_n^0 = \frac{w}{\alpha},$$
(EC.58)

$$s_n^H(i_H,\xi_n) = \frac{w}{\alpha} - \frac{1}{r(1-\alpha)} \ln\left(-\frac{1-p}{p} + \frac{v(c_L/\alpha)(\overline{u}_L+1) - v(c_H/\alpha)(\overline{u}_H+1)}{(\overline{u}_H - \overline{u}_L)p}\right).$$
 (EC.59)

Proof Sketch. The proof follows the same steps as the proof of Theorem 1, with the following key differences. After applying Theorem 1 of Plambeck and Zenios (2000), the problem reduces to (EC.1)-(EC.4) with an additional IC constraint that ensures that the agent cannot do better by choosing a different search intensity. For the case where the principal desires to induce i_H , this constraint is given by (EC.57). Separation into three different bonus levels follows from the observation that all alternatives in $\{0, 1, ..., \bar{u}_L\}$ carry the same information about the agent's search intensity, as do ones in $\{\bar{u}_L + 1, ..., \bar{u}_H\}$. Part (i) of the proposition is the same as in Theorem 1, as the new constraint is irrelevant when $\omega_n = 0$. Part (ii) follows from the observation that changing s_n^H while keeping other contract parameters the same only affects the agent's payoff if he uses i_H , which reduces the problem to the one solved in Theorem 1 if the principal desires to induce i_L . For part (iii), the constrained problem is solved using KKT conditions (KKT is both sufficient and necessary here, as in Theorem 1), which yields (EC.56) if (EC.57) is not binding, or (EC.58)-(EC.59) if it is. \Box

This result breaks into a number of cases, yet, most of them closely parallel our main model: (EC.55) and (EC.56) are slight variations of our basic result. The sole situation where the results here substantially differ from the main model arises when the principal desires to induce $\omega_n = i_H$ combined with $\xi_n \leq \overline{u}_L$. This difference occurs because here it is possible that the traditional flat contract for everything delivered, as given by (EC.56), induces $\omega_n = i_L$ instead.

If this happens, the solution involves modifying the values in (EC.56) simultaneously rising $s_n^H(i_H, \xi_n)$ (which makes the contract more attractive for the agent who is using high intensity) and reducing the $s_n^L(i_H, \xi_n)$ component (which makes the contract less attractive to both high and low intensity, but more so for the low intensity, who will experience these outcomes more often).

References

- Bertsekas DP (2012) Dynamic programming and optimal control, 4th edition, volume 2 (Athena Scientific Belmont, MA).
- Bertsekas DP (2017) Dynamic programming and optimal control, 4th edition, volume 1 (Athena Scientific Belmont, MA).
- Bertsekas DP, Shreve S (2004) Stochastic optimal control: the discrete-time case.
- Fudenberg D, Holmstrom B, Milgrom P (1990) Short-term contracts and long-term agency relationships. Journal of Economic Theory 51(1):1–31.
- Kotsiantis S, Kanellopoulos D (2006) Discretization techniques: A recent survey. GESTS International Transactions on Computer Science and Engineering 32(1):47–58.
- Lippman SA, McCall J (1976) The economics of job search: A survey. *Economic Inquiry* 14(2):155–189.
- Miller III AC, Rice TR (1983) Discrete approximations of probability distributions. *Management Science* 29(3):352–362.
- Plambeck EL, Zenios SA (2000) Performance-based incentives in a dynamic principal-agent model. Manufacturing & Service Operations Management 2(3):240–263.
- Smith JE (1993) Moment methods for decision analysis. Management Science 39(3):340–358.
- Smith JE (1998) Evaluating income streams: A decision analysis approach. *Management Science* 44(12-part-1):1690–1708.
- Topkis DM (1978) Minimizing a submodular function on a lattice. Operations Research 26(2):305–321.