Robust Contracting for Sequential Search*

Théo Durandard[†] Udayan Vaidya[‡] Boli Xu[§]

September 12, 2025

Click Here For Latest Version

Abstract

A principal contracts with an agent who sequentially searches over projects to generate a prize. The principal initially knows only one of the agent's available projects and evaluates a contract by its worst-case performance. We characterize the principal's robustly optimal contracts, which are all debt-like: the agent is only paid when the prize exceeds a threshold. Debt is optimal because it preserves the option value of continued exploration. Our characterization encompasses several common contract forms, including pure debt, debt-plus-equity, and capped-earnout debt. We identify settings in which each of these contracts is uniquely optimal.

Keywords: Robust contracts, sequential search, debt contracts, moral hazard.

JEL Codes: D82, D83, D86.

*We thank Nageeb Ali, Ian Ball, Dan Barron, Dan Bernhardt, Matteo Camboni, Alexis Ghersengorin, Mallesh Pai, Alessandro Pavan, Mike Powell, Debraj Ray, Luis Rayo, Pëllumb Reshidi, Ludvig Sinander, Bruno Strulovici, João Thereze, Jidong Zhou, and various audiences at 2024 Stony Brook ICGT, 2024 SEA, 2025 EAYE, 2025 SAET, 2025 EC, Duke, Northwestern, PSU, Stanford, TAMU, Toronto, U Iowa, and UNC.

[†]University of Illinois Urbana-Champaign. Email: theod@illinois.edu

[‡]Duke University. Email: udayan.vaidya@duke.edu

[§]University of Iowa. Email: boli-xu@uiowa.edu

1 Introduction

Innovation requires exploring ideas whose value is revealed only after investigation. Because this exploration typically demands expertise, a principal funding an innovation often contracts with an agent to carry it out. Such delegated search relationships arise in a range of settings, including the development of intellectual property, R&D partnerships, and the acquisition of early-stage startups. The principal's delegation to the agent introduces two core frictions: the agent's search process may be inherently difficult to monitor, and the agent may have better information about potentially fruitful ideas. Faced with limited information and limited control, how can the principal design a contract that encourages the agent to explore only worthwhile ideas?

To study this problem, we augment the canonical robust contracting model of Carroll (2015) to incorporate the agent's *sequential exploration* of alternatives. In our model, a principal ("she") contracts with an agent ("he") who can search over projects to generate a prize. Each project is modeled as a "Pandora's Box" (Weitzman, 1979), characterized by a prize distribution and a fixed cost required to make the prize collectible. The agent is protected by limited liability, and his search process is unobservable to the principal. Moreover, we assume that the principal knows only one of the agent's available projects at the time of contracting. For instance, this could be an author's "pitch," an inventor's prototype, or an industry-standard method for conducting research. Following Carroll (2015), we adopt a robustness approach to address this informational asymmetry. The principal designs a contract to maximize her worst-case expected profit given her initial knowledge.

A cornerstone result of our paper is that a *debt contract* is robustly optimal for the principal. In a debt contract, the principal collects the entire prize if it falls below a predetermined debt level, while the agent keeps any profits in excess of the debt level.

Roughly, debt is optimal because it reflects the option value of the agent's continued search. By promising payment to the agent only after he reaches a certain milestone, the agent is incentivized to "swing for the fences" and explore high-risk, high-reward projects. Our debt contract strictly outperforms linear contracts that are typically optimal in the robust contracting literature (e.g., Carroll, 2015; Walton & Carroll, 2022; Dai & Toikka, 2022; Liu, 2022). Linear contracts are suboptimal because they dampen the agent's search incentive, making the principal vulnerable to low-risk, low-reward projects that were unknown at the time of contracting.

Building on the optimal debt contract, Theorem 1 characterizes the set of robustly optimal contracts, which are all debt-like. It establishes that a contract is robustly optimal if and only if it involves a minimum debt level and enables the principal to capture the full social surplus in the worst-case scenario. The minimum debt level is crucial to discourage the agent from terminating search too early, while the full surplus extraction condition maximizes the principal's profit subject to the agent's participation. Our characterization encompasses several commonly observed contract formats, including pure debt, debt-plus-equity, and capped-earnout debt.

Our baseline model provides a novel rationale for the use of debt-like contracts in many environments; however, its predictions are not unique. In Section 5, we refine our model to better understand when specific contract formats will emerge. In particular, we show that pure debt arises when the agent can resample projects; debt-plus-equity emerges under principal moral hazard; and capped-earnout debt is appropriate when the agent is risk-averse. Lastly, in Section 6, we extend the model to accommodate the principal contracting with multiple agents who have different projects. We show that dynamically sponsoring agents and offering each a personalized debt contract is optimal, underscoring the robustness of this paper's insight.

The contributions of this paper are threefold. First, we demonstrate the importance of incorporating additional structure on the agent's moral hazard problem within a robust contracting framework. Our model's departure from Carroll (2015)—allowing the agent to sequentially search instead of making a one-time decision— drastically changes the shape of the optimal contract. One intuition for why the models generate different predictions relates to which project statistics guide the agent's decision-making. In Carroll (2015), the value of a project is captured by its *expected prize* net of cost, making linear contracts desirable. In our dynamic search setting, however, the relevant statistic for the agent's decision is the *Weitzman index*, which reflects the option value of continued search. The agent's payoff under a debt contract mirrors the index: both evaluate the expected prize in excess of a threshold. As a result, the agent's decision-making under a debt contract reproduces the socially optimal behavior.

Second, to our knowledge, our paper is the first to study contracting for sequential search *à la* Weitzman (1979) in a general setting with moral hazard and asymmetric information. Even without asymmetric information, we are unaware of the solution to the Bayesian version of this contracting problem outside a few special cases. This paper further underscores how a robustness approach can deliver simple insights that would be muddled in a Bayesian setting.

Finally, our results speak to the prevalence of debt-like contracts in practice and shed light on the forces that give rise to various debt-like instruments. An alternative interpretation of our model is one where the agent designs a security to sell to a cautious investor. If the agent's goal is to maximize the sale price—for example, because he faces a liquidity constraint—the solution coincides with that of our model. This interpretation aligns us with more conventional applications of debt within the finance literature. We provide several real-world examples of these contracts, as well as a more general discussion of applications, in Section 5 after presenting our main results.

Related Literature

Our paper primarily connects two strands of the literature that have remained largely separate. The first studies the financing of experimentation and innovation, highlighting the role of dynamic incentives, and the second investigates the provision of robust incentives in moral hazard environments, emphasizing mechanisms that perform well under limited knowledge of the underlying technology.

Lewis and Ottaviani (2008), Lewis (2012), and Ulbricht (2016) study delegated search in the context of R&D financing. A common feature of these papers is the stationarity of the underlying environment: the agent repeatedly samples from an identical distribution of outcomes, akin to facing an infinite sequence of homogeneous Pandora's boxes in our model. As a result, the literature on delegated search primarily focuses on the intensity of effort or the duration of search. A similar emphasis is found in the broader literature on contracting for experimentation, which studies how conflicts of interest affect the dynamics of payments and the length of experimentation (Bergemann & Hege, 1998, 2005; Hörner & Samuelson, 2013; Halac, Kartik, & Liu, 2016; Guo, 2016).

We examine a complementary question: not *how much* to search but *what* to search for. In many delegated search settings, the agent faces a project selection problem—such as choosing between safe and risky technologies—on top of deciding how intensely to search. Capturing this possibility requires departing from the stationary framework and allowing for the possibility of *heterogeneous* Pandora's boxes. Hence, our model also relates to the small literature on delegated project choice (see, e.g., Armstrong & Vickers, 2010, Nocke & Whinston, 2013). The closest paper in this literature is Guo and Shmaya (2023), which derives the robust, regret-minimizing policy. We differ by allowing for transfers and requiring the agent to exert effort to discover a prize. This leads to distinct tradeoffs because we have different instruments to separate incentive provision from project selection. For example, our principal does not gain from incentivizing the agent to disclose multiple prizes.

Introducing project selection into delegated search creates significant challenges. The Bayesian version of the problem is generally intractable and, when solvable, often yields highly specific, assumption-driven solutions. This motivates our study of the minimax optimal contract under minimal assumptions about the agent's technology, placing us within the literature on robust contracting stemming from Hurwicz and Shapiro (1978). As mentioned above, the paper closest to ours is the seminal contribution of Carroll (2015). We only minimally depart from Carroll's model to study dynamic search, allowing the agent to take multiple actions sequentially. Relatively few papers study robustness in a dynamic setting.[1] Chassang (2013) and Liu (2022) consider dynamic incentive problems that are separable across periods, and show that linear contracts perform well. Our incentive problem is not separable because the principal ultimately selects a single prize, and our optimal contract disregards the potential arrival of intermediate information. In this respect, we are closer to Libgober and Mu (2021) and Koh and Sanguanmoo (2024), who also derive mechanisms that are insensitive to the arrival of intermediate information. Libgober and Mu (2021) study a durable goods monopoly, and Koh and Sanguanmoo (2024) study a learning setting similar to ours. One key difference is that they assume learning benefits the agent but may impose a negative externality on the principal, while in our model, the prize accrues to the principal and the costs are borne by the agent.

One can alternatively view our model as a *fully static* extension of Carroll (2015), where the principal possesses additional knowledge on the set of feasible actions. Specifically, rather than considering all actions possible, the principal only entertains action sets whose payoff distributions can be generated by sequential search through some set

[1]We direct the reader to Carroll (2019) for a survey of this literature, which discusses some of the challenges in studying dynamics in robust contracting.

of projects.[2] Our approach provides a complementary perspective to the imposition of exogenous restrictions on the ambiguity set, which often requires making strong assumptions to maintain tractability (e.g., considering only MLRP ranked families as in Antic, 2021). We therefore contribute to a nascent theoretical literature that investigates the sensitivity of robust mechanisms to perturbations in the ambiguity set (Walton & Carroll, 2022; Kambhampati, 2024; Ball & Kattwinkel, 2024; Olszewski, 2025). Our results thus confirm that natural changes to the ambiguity set are likely to alter the optimal design, as emphasized by Olszewski (2025).

Finally, our findings complement a large finance-oriented literature that seeks to explain the prevalence of debt contracts. As we do, many papers have shown the efficacy of debt in mitigating moral hazard (e.g., Jensen & Meckling, 1976; Townsend, 1979; Innes, 1990; Hébert, 2018).[3] The common intuition is that debt contracts approximate complete ownership transfer to the agent, maximizing incentives while respecting limited liability. However, these results often rely on strong distributional assumptions regarding the agent's technology or restrictions on the contract space, such as double monotonicity. More broadly, while debt is valued for its ability to incentivize effort, it is also criticized for encouraging risk-shifting, an outcome typically undesirable for the principal (Jensen & Meckling, 1976). Instead, in our delegated search problem, both effects work in the principal's favor. Intuitively, debt contracts push the agent toward efficiently exploring risky alternatives and discouraging early termination.

[2]For instance, a static model in which the agent has exactly two actions available, (F_0, c_0) and (F_1, c_1), would be inconsistent with the possibility of sequential search. In addition, the agent should have actions that correspond to adaptively searching the two alternatives.

[3]Beyond moral hazard, many additional theories justify the widespread use of debt. These include lowering firms' tax burden (Miller, 1977), signaling firms' positive private information about profitability (Myers & Majluf, 1984), mitigating adverse selection (Dang, Gorton, & Holmström, 2011; DeMarzo & Duffie, 1999; Nachman & Noe, 1994; Yang, 2020), maintaining control or limiting investment (Aghion & Bolton, 1992; Hart & Moore, 1994; Jensen & Meckling, 1976), etc. See Tirole (2010) for an overview of the literature on security design and capital structure.

2 Model

The setup and notation of our baseline model mirror that of Carroll (2015) when possible.

A principal ("she") contracts with an agent ("he"), who sequentially searches through projects to generate a prize. Each project is a "Pandora's box" in the language of Weitzman (1979), and the set of projects available to the agent is denoted by $\mathcal{A} = \{a_i\}_{i=0}^n$, with $n \in \mathbb{N}$. The project a_i is described by a pair $(F_i, c_i) \in \Delta(\mathbb{R}_+) \times \mathbb{R}_+$ such that F_i has finite expectation. The parameter c_i represents the agent's private *cost* to learn project i's realized *prize* y_i, which is distributed according to F_i. Each project's prize is drawn independently, and we assume both players are risk-neutral.

When facing a set of projects $\mathcal{A}$, the agent engages in sequential search (with recall) *à la* Weitzman (1979). We can describe the agent's strategy as a function of two state variables: (1) the set of projects sampled up to date, denoted by $\tilde{\mathcal{A}} \subseteq \mathcal{A}$, and (2) the highest monetary wage from the sampled projects, denoted by $\tilde{y}$. Formally, a strategy is a function $\sigma : 2^{\mathcal{A}} \times \mathbb{R} \to 2^{\mathcal{A}} \cup \{\emptyset\}$, where $\sigma(\tilde{\mathcal{A}}, \tilde{y}) = a_i$ means that the agent will sample project $a_i \in \mathcal{A} \setminus \tilde{\mathcal{A}}$ next and $\sigma(\tilde{\mathcal{A}}, \tilde{y}) = \emptyset$ means that he will cease searching and present a prize that gives himself the monetary reward of $\tilde{y}$. The agent may always return a prize of zero without sampling any projects. We abusively write $a_i \in \sigma$ to denote the event that the project a_i is sampled according to the strategy σ.

As in a typical moral hazard environment, we assume that the agent's action—i.e., the search process—is unobservable and cannot be directly contracted on. In particular, the principal does not know which projects the agent sampled to generate a prize. However, she can contract on the final prize y the agent presents to her.[4] The principal's only incentive tool is thus a wage contract $w : \mathbb{R}_+ \to \mathbb{R}$, where $w(y)$ is the agent's monetary

[4] We assume the agent can only present one *single* prize, as is often the case in practice. In Section 4, we argue that the principal does not gain from offering general mechanisms, including allowing the agent to present more than one prize.

payment when the presented prize is y. All contracts must be measurable and satisfy limited liability, meaning $w(y) \geq 0$ for any $y \in \mathbb{R}_+$.

The timing of the game is summarized as follows.

1. The principal sets a contract w.

2. The agent sequentially searches among $\mathcal{A}$, after which he presents a prize y.

3. The agent's payoff is $w(y) - \sum_i c_i \mathbf{1}_{[a_i \in \sigma]}$ and the principal's payoff is $y - w(y)$.

Although the agent knows the true set of projects $\mathcal{A}$, we assume the principal only knows one of those projects, $a_0 = (F_0, c_0)$, at the time of contracting. So, she considers all $\mathcal{A} \supseteq \mathcal{A}_0 := \{a_0\}$ as possible.

The principal's objective is to determine a wage contract w that maximizes the worst-case expected payoff against all possible realizations of $\mathcal{A}$ consistent with $\mathcal{A}_0$. Formally, given a realized set of projects $\mathcal{A}$ and a wage contract w, we let $\Sigma(w, \mathcal{A})$ denote the set of optimal search strategies for the agent. For a given strategy σ, we write $\mathbb{E}_\sigma$ to denote the expectation with respect to the induced distribution over prizes generated by the agent's search. The principal's payoff for a given contract and set of projects is

$$V_P(w \mid \mathcal{A}) := \sup_{\sigma \in \Sigma(w, \mathcal{A})} \mathbb{E}_\sigma[y - w(y)],$$

where we assume the agent breaks ties in favor of the principal.[5]

The principal evaluates a wage contract w by its payoff guarantee $V_P(w)$, her expected payoff in the worst-case realization of $\mathcal{A}$:

$$V_P(w) := \inf_{\mathcal{A} \supseteq \mathcal{A}_0} V_P(w \mid \mathcal{A}).$$

[5]This assumption is for ease of exposition, and is not essential for the qualitative takeaways of this paper. See Section 4 for further discussion.

The principal seeks a contract that maximizes her payoff guarantee, defined as V_P below. We call a contract that achieves a guarantee of V_P a (robustly) optimal contract.

$$V_P := \sup_w V_P(w)$$

3 Optimal Contracts

Our main objective in this section is to identify the principal's robustly optimal contract(s). To do so, we begin by describing the agent's optimal search strategy. Given this strategy, we will argue that a debt contract is robustly optimal, which we then extend to a class of related contracts. After presenting the main results, we contrast our findings with those in the previous literature and explain why the differences emerge.

3.1 Agent's Response to a Wage Contract

The agent's optimal search strategy depends on the offered contract and the available set of projects. As a first step, we begin by recalling the solution to the problem of Weitzman (1979). For the project $a_i = (F_i, c_i)$, we define its *index* (or *reservation value*), r_i, as the smallest solution to[6]

$$c_i = \int [y_i - r_i]^+ dF_i(y_i). \tag{1}$$

A social planner who maximizes the joint welfare of the principal and the agent facing the same search problem will (1) sample the projects in descending order of their index values and (2) conclude the search whenever a realized prize y is larger than the index values of the remaining unsampled projects.[7]

[6]When c_i is positive, the solution is unique. However, when $c_i = 0$, there may be a continuum of solutions, including infinity. The index may also be negative, in which case it is never optimal for the agent to sample the project.

[7]If there are multiple projects with the same index, the agent is indifferent to their ordering, and there are multiple optimal strategies.

When facing the wage contract w, the agent still uses an index strategy, except his incentives are distorted by the contract. Given a wage contract w, project a_i's *w-induced index* (or *w-induced reservation value*), r_i^w, is the smallest solution to

$$c_i = \int [w(y_i) - r_i^w]^+ dF_i(y_i). \tag{2}$$

Thus, the agent optimally samples projects in descending order of r_i^w and stops searching when his best realized monetary payoff exceeds the w-induced index values of the remaining unsampled projects.

3.2 Debt Contract Outperforms Linear Contracts

Having specified the agent's optimal response to any contract w, we now return to the principal's problem of finding the contract with the maximal payoff guarantee. We begin by identifying a simple upper bound on the principal's payoff guarantee: the surplus of the known project a_0, which we denote by s_0.

$$s_0 := \mathbb{E}_{F_0}[y] - c_0$$

Observation 1. *No contract can guarantee the principal more than the full surplus of the known project. That is, for all w, $V_P(w) \leq s_0$.*

Proof. Under the known set $\mathcal{A}_0$, $V_P(w \mid \mathcal{A}_0) \leq s_0$ for all contracts, for otherwise the agent's payoff would be strictly negative, and he would prefer to not search and return a prize of 0 instead. The observation follows because $V_P(w) \leq V_P(w \mid \mathcal{A}_0)$ by definition. □

It is not immediately clear that this upper bound is useful, as robustly extracting the entire known surplus is a demanding desideratum. Indeed, we will argue that linear contracts—a natural candidate given the preceding literature—are unable to deliver this guarantee. Nonetheless, a properly calibrated *debt contract* can extract the full surplus, even in the worst case.

To argue this, we start with a useful observation that helps identify the principal's worst-case payoff when using a contract that satisfies double monotonicity. A contract w is said to be **doubly monotone** if both $w(y)$ and $y - w(y)$ are non-decreasing in y. Although the following result only applies to doubly monotone contracts, the main results of the paper do not assume double monotonicity.

Observation 2. *Fix any doubly monotone contract w, and let $V_P(w)$ be the corresponding payoff guarantee. Either (1) $V_P(w \mid \mathcal{A}_0) = V_P(w)$ or (2) for any $\epsilon > 0$, there exists $a_1 = (\delta_x, 0)$ such that $V_P(w \mid \mathcal{A}_0 \cup \{a_1\}) < V_P(w) + \epsilon$, where δ_x is a Dirac mass for some x satisfying $w(x) > r_0^w$.*[8]

Proof. See Appendix A.1. □

One interpretation of this result is that the principal's primary concern is the presence of *cheap* and *safe* alternatives unknown at the time of contracting. The principal's worst case is either when the agent has a_0 alone or exactly one other project a_1 that crowds out the known project. The project a_1 constructed in Observation 2 delivers a constant wage that just exceeds the w-induced index of the known project a_0. Hence, the agent's optimal search strategy when facing $\mathcal{A} = \{a_0, a_1\}$ is to only sample a_1 and terminate the search afterward. Unlike in Carroll (2015), the principal's payoff is never minimized by additional projects with stochastic output.

Applying Observation 2 to linear contracts, which are doubly monotone, shows that their payoff guarantee is strictly below the surplus of the known project.

Observation 3. *If $c_0 > 0$, no linear contract robustly guarantees the full surplus,*[9] *i.e.,*

[8] The ϵ accounts for the infimum in the definition of V_P, which may not be achieved.

[9] For $c_0 = 0$, the linear contract with slope 0 coincides with the 0-debt contract, extracts the full surplus, and is robustly optimal.

$\sup_{\alpha \in [0,1]} V_P(w(y) = \alpha y) < s_0$.

Proof. See Appendix A.2 □

The main idea behind Observation 3 is seen by considering the naive linear contract $w(y) = \alpha^* y$, where $\alpha^* = \frac{c_0}{\mathbb{E}_{F_0}[y]}$. Under this contract, the principal extracts the full surplus of the known box when the agent can only sample a_0. However, in light of Observation 2, consider what happens when the agent additionally has access to some project $a_1 = (\delta_x, 0)$, where $x > 0$ is arbitrarily small. Because the induced index of the known project is zero, the agent will sample only a_1 and deterministically return a prize of x. Since x may be arbitrarily small, the payoff guarantee of the naive linear contract is 0. The proof in Appendix A.2 also shows that linear contracts with slope close to α^* cannot approximate the naive upper bound obtained in Observation 1.

Intuitively, linear contracts fail to have a high payoff guarantee because they dampen the agent's incentive to explore risky projects. When the agent encounters a safe, low-value project, he may abandon exploration prematurely, even though continued exploration benefits the principal. This occurs because linear contracts overly penalize high levels of output, which are focal in the computation of the project's index value as defined in Equation (2). This reasoning suggests that the principal's optimal contract should make low-risk projects look relatively unattractive to the agent, thereby limiting the principal's exposure to unknown alternatives.

Fortunately, *debt contracts* are immune to these issues. For any $z \in \mathbb{R}^+$, the $z-$**debt contract** is defined as the contract $w(y) := [y - z]^+$. When the principal offers the $z-$debt contract, she collects all the returns up to the debt level z, after which the agent becomes the residual claimant and collects any prize in excess of z. Conceptually, this is equivalent to offering the agent a call option on the prize with a strike price of z.

When properly calibrated to the known project, a debt contract can achieve the payoff guarantee of s_0 given in Observation 1. Let r_0 denote the index of the known project a_0

as defined using Equation (1), and let w_0 denote the r_0-debt contract $w_0(y) = [y - r_0]^+$.

Proposition 1. *The payoff guarantee of the r_0-debt contract, w_0, is the entire surplus s_0. Therefore, w_0 is robustly optimal, and $V_P = V_P(w_0) = s_0$.*

Proof. We argue that $V_P(w_0) = s_0$, which implies that w_0 is robustly optimal given the upper bound of $V_P \leq s_0$ identified in Observation 1.

First, $V_P(w_0 \mid \mathcal{A}_0) = s_0$ because the calibrated debt level r_0 leaves the agent zero expected payoff when $\mathcal{A} = \mathcal{A}_0$. To see this, notice that

$$\mathbb{E}_{F_0}[w_0(y)] - c_0 = \mathbb{E}_{F_0}[(y - r_0)^+] - c_0 = 0,$$

where the last equality follows by definition of the index r_0. In other words, when a_0 is the only available project, the contract w_0 maximizes total surplus (because the agent is still willing to sample a_0) while leaving none to the agent. This also implies that $r_0^{w_0} = 0$.

Second, we show that $V_P(w_0) = V_P(w_0 \mid \mathcal{A}_0)$, i.e., the worst-case scenario is $\mathcal{A} = \mathcal{A}_0$ *when the principal offers* w_0. The proof proceeds by contradiction. So, suppose $V_P(w_0) < V_P(w_0 \mid \mathcal{A}_0)$. Since w_0 is doubly monotone, Observation 2 implies that there exists $\mathcal{A} = \{a_0, a_1\}$, where $a_1 = (\delta_x, 0)$ is a safe project with $w_0(x) > r_0^{w_0} = 0$, such that $V_P(w_0 \mid \mathcal{A}_0) > V_P(w_0 \mid \mathcal{A})$. However, the agent's wage under w_0 is strictly positive only if the presented prize is strictly greater than r_0. This guarantees that the principal's profit is r_0 when the agent has access to the set $\mathcal{A}$, which is the best possible outcome for the principal under the contract w_0. Since $r_0 > s_0$ by Equation (1), $\mathcal{A} = \{a_0, a_1\}$ cannot be worse than $\mathcal{A}_0$.

Thus, $V_P(w_0) = V_P(w_0 \mid \mathcal{A}_0) = s_0$, and w_0 is robustly optimal. □

As the proof demonstrates, w_0 attains the payoff guarantee upper bound because it achieves two goals simultaneously. First, w_0 fully guards the principal against safe, low-value projects that may preemptively terminate the agent's search. The agent is not rewarded for presenting low prizes, and thus any prize that can terminate his search before exploring a_0 must be sufficiently large. Debt completely eliminates the risk of the agent's access to unknown projects, unlike the linear contract. Second, the debt level r_0 is chosen so that the principal extracts the entire surplus when a_0 is the only available project. Together, these two properties maximize the principal's payoff guarantee.

3.3 Characterization of Robustly Optimal Contracts

Following the construction of the optimal debt contract w_0 in Proposition 1, a natural question is whether other contracts can also attain the same payoff guarantee. The answer turns out to be yes. To generalize the previous result, notice that for any contract w, we must have

$$V_P(w) \leq V_P(w \mid \mathcal{A}_0) \leq s_0. \tag{3}$$

As outlined below Proposition 1, the structure of the debt contract w_0 causes both inequalities in Equation (3) to bind. By contrast, we have argued that linear contracts cannot simultaneously deliver both as equalities, rendering them strictly suboptimal. Since the principal's maximal payoff guarantee is exactly s_0, a necessary and sufficient condition for a contract to be robustly optimal is therefore to satisfy both sides of Equation (3) with equality. Using this logic, the following theorem characterizes the properties of robustly optimal contracts.

Theorem 1. *A contract w is robustly optimal if and only if it satisfies the following two conditions.*

1. ***Minimum Debt Level condition*** *(MDL henceforth):*

$$w(y) \leq [y - s_0]^+$$

2. ***Full Surplus Extraction condition*** *(FSE henceforth):*

$$\mathbb{E}_{F_0}[w(y)] = c_0.$$

Proof. See Appendix A.3. □

The intuition for this theorem parallels that of Proposition 1. The FSE condition ensures that the principal collects the full surplus when a_0 is the only available project.[10] The MDL condition guarantees that $\mathcal{A} = \mathcal{A}_0$ is indeed the worst case. As before, the principal collects the entire prize if it falls below s_0, whereas her payoff $y - w(y)$ is at least s_0 if the presented prize exceeds s_0. This guards the principal against safe, low-value crowding out the known project because the principal will at least collect s_0 even if the agent terminates the search before sampling a_0.

However, because we do not impose double monotonicity of the contract as in Observation 2, we must address one additional wrinkle. Previously, double monotonicity implied that any additional search by the agent only increased the prize y that was presented, and thus the principal's profit $y - w(y)$. If $w(y)$ and $y - w(y)$ are not co-monotone, the agent continuing to sample additional projects after a_0 could potentially reduce the principal's profit. Both the FSE and MDL conditions play a role in eliminating this possibility. If the agent chooses to sample some project a_1 after sampling a_0, it must have a w-induced index of exactly zero, $r_1^w = 0$. Otherwise, the agent would have sampled a_1 before a_0, which has $r_0^w = 0$ by FSE, or not at all. This implies the agent's best wage before sampling a_1 must be zero; otherwise, he would have already stopped searching. Finally, if the prize y_1 from a_1 satisfies $w(y_1) = 0$, tie-breaking in favor of the principal improves her payoff; otherwise, if $w(y_1) > 0$, then the MDL condition guarantees the principal's profit is again at least s_0.

[10] Here, we use principal-favored tie-breaking by the agent for a tight characterization. A more thorough discussion appears in Section 4.

Theorem 1 underscores the importance of debt-like instruments—such as performance targets and earnout contracts—in incentivizing exploration. For a contract to be robustly optimal, some level of debt (though not necessarily as high as r_0) is inevitable. This is reminiscent of many real-world delegated search relationships where the agent is rewarded only if the outcome reaches a certain standard. For instance, pharmaceutical startups typically earn bonuses only after passing certain clinical trials, while researchers are disproportionately rewarded for publications in top journals. Broadly, these high hurdles provide the appropriate incentives to "swing for the fences." Although such risk shifting may be detrimental in other cases, because our environment features exploration, it is a boon. The risks of failure are muted relative to the benefits of success due to the option value of exploring alternative projects.

The set of robustly optimal contracts identified in Theorem 1 includes several types of contracts that are commonly observed in practice. We depict three prominent types of contracts in Figure 1, which will be the focus of Section 5.

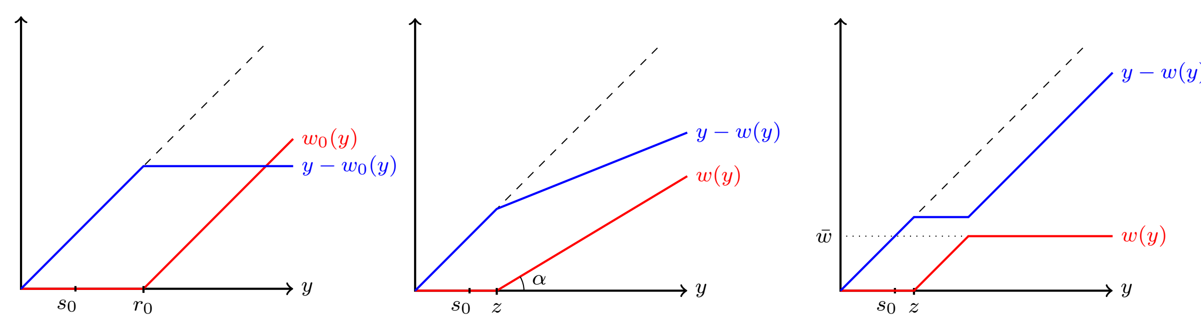

(a) Pure debt (b) Debt-plus-equity (c) Capped-earnout debt

Figure 1: Three examples of optimal contracts. Panel (a) depicts the pure debt contract w_0. Panel (b) depicts the (z, α)-debt-plus-equity contract. Panel (c) depicts the $(z, \bar{w})$-capped-earnout-debt contract. In each panel, the red and the blue curves represent the agent's and the principal's payoffs, respectively.

The first type of contract is *pure debt* (previously, just called "debt"). The only robustly optimal pure debt contract is w_0, the focus of Proposition 1. This contract has the highest debt level among all doubly-monotone contracts.

A second type of contract is *debt-plus-equity*, or convertible debt. We refer to a contract w as the (z, α)-debt-plus-equity contract if $w(y) = [\alpha(y - z)]^+$ for some $z > 0$ and $\alpha \in (0, 1)$. Under such a contract, the principal collects all the returns up to the debt level z, after which the agent collects an α share of the residual. A debt-plus-equity contract differs from a pure debt contract as the principal still shares some profit after the debt is fully repaid. Practically, this is implemented by an option in the contract to convert the debt into equity participation at a pre-specified rate. In our model, there is a continuum of debt-plus-equity contracts that are robustly optimal — specifically, a (z, α)-debt-plus-equity contract is robustly optimal if and only if $z \in [s_0, r_0)$ and $\alpha \mathbb{E}_{F_0}(y - z)^+ = c_0$.

The third type of contract we focus on is *capped-earnout debt*. We refer to a contract w as the $(z, \bar{w})$-capped-earnout-debt contract if $w(y) = \min\{\bar{w}, [y - z]^+\}$ for some $z > 0$ and $\bar{w} > 0$. A capped-earnout debt contract is similar to a debt contract except that the agent's wage (or "earnout payment") is bounded above by $\bar{w}$. The agent is fully sensitive to marginal changes in the prize between z and $\bar{w}$, but insensitive outside this region. This contract resembles a senior tranche used in structured finance. In our model, there is also a continuum of capped-earnout debt contracts that are robustly optimal — specifically, a $(z, \bar{w})$-capped-earnout-debt contract is robustly optimal if and only if $z \in [s_0, r_0)$ and $\mathbb{E}_{F_0}[\min\{\bar{w}, (y - z)^+\}] = c_0$.

Within our baseline model, there is no obvious way to select among these robustly optimal contracts. Notably, all contracts identified by Proposition 1 are admissible; each delivers the same payoff guarantee and is the principal's preferred contract for some realization of $\mathcal{A}$. An alternative refinement would be on the basis of the total surplus generated by a contract. Unfortunately, in general, there is no contract that maximizes the total surplus against any $\mathcal{A} \supseteq \mathcal{A}_0$ while remaining robustly optimal. However, under some conditions on a_0, the pure debt contract w_0 is both robustly optimal and socially efficient. This follows because pure debt contracts are *order-preserving*, i.e., the agent's search order coincides with the social planner's, despite the presence of the contract. Direct computation shows that the w_0-induced index of *any* project a_i is $r_i^{w_0} = r_i - r_0$. Because a pure debt contract uniformly reduces the indices of all projects, it does not distort the order in which the agent searches them, nor when the agent terminates search with a positive wage. However, the debt contract may prevent the agent from exploring projects with lower indices than a_0; hence, some extra conditions are required. No other type of contract besides pure debt can be socially efficient, as any alternative contract necessarily distorts the agent's search over some collection of projects. These results are presented formally in Appendix B.

Instead of direct selection based on admissibility or efficiency, we find it more instructive to enrich the model setup to identify additional forces that might select one contract or another. This analysis continues in Section 5.

3.4 Comparison with Literature

We are now in a position to explain why our main findings differ from those in the prior literature on robust contracting, where linear contracts—not debt contracts—are typically optimal. Our model departs from Carroll (2015) in one aspect only: rather than choosing a project as a one-time decision, the agent in our framework sequentially searches across alternatives. While this difference may appear minor, it leads to qualitatively different conclusions. We propose two lenses through which to understand this distinction.

First, the robustly optimal contracts in the one-shot versus sequential search models reflect different "sufficient statistics" of a project that guide the agent's choice in each environment. In Carroll (2015), the social value of a project is captured by its *expected* prize net of cost. This makes linear contracts desirable, since they align the agent's and principal's incentives over different risk distributions, thereby minimizing distortions. In our dynamic search setting, however, the relevant statistic is the Weitzman index, which captures the option value of continued search. A pure debt contract is therefore optimal because it mirrors the index: both evaluate the expected prize in excess of a threshold. As a result, the agent's decision-making under a debt contract reproduces socially optimal behavior.[11] This leaves less room for a malevolent nature to interfere, as any unknown projects can only improve social surplus and simultaneously the principal's payoff. As Theorem 1 shows, an optimal contract need not look exactly like the Weitzman index, but it must have the right shape (MDL) and level (FSE).

Second, due to the differing agent strategies, the type of unknown project that our principal guards against is thus fundamentally different than that of Carroll (2015). In Carroll (2015), the principal worries that the agent selects an excessively risky project and gambles for a favorable realization. Instead, in our setting, the principal would benefit from such gambles. So, she does not need to protect herself against excessive risk-taking, but aims to prevent the agent from settling too early on a safe, low-value project and prematurely concluding the search process (as suggested by Observation 2). Linear contracts effectively align the principal's and agent's incentives in a static moral hazard problem, but they discourage search in our dynamic problem by making safe projects relatively more attractive. Debt contracts offer a distinct advantage: they discourage early stopping by making low prizes unattractive to the agent, thus encouraging risk-taking and continued experimentation.

[11] As mentioned in Section 3.3, one caveat is that the agent will not search unknown projects whose indices are lower than the known one. But, as Observation 2 suggests, such projects will not appear in the worst case under a debt contract.

4 Discussion of Assumptions

Before turning to extensions of our baseline model, we briefly comment on the roles of several of our modeling assumptions.

Limited liability. In our setting, limited liability plays a slightly different role than in Carroll (2015) and the security design literature at large. As is typical, if we allow the agent's wage to be negative, the principal can "sell the firm" to the agent for the expected surplus, using the contract $w(y) = y - s_0$. Thus, limited liability is clearly important in the necessity direction of Theorem 1, where we characterize all optimal contracts. However, relaxing limited liability does not interfere with sufficiency because the contracts we identify give the same worst-case payoff as selling the firm, so they remain robustly optimal.

Moreover, selling the firm to the agent is *weakly dominated* by any of the contracts identified in Theorem 1. Intuitively, selling the firm must yield a deterministic payoff of s_0—the sale price—regardless of the true set of projects $\mathcal{A}$. By contrast, for any robustly optimal contract w identified in Theorem 1, the principal's expected payoff is weakly greater than s_0 across all realizations of $\mathcal{A}$, with the inequality being *strict* for some realizations of $\mathcal{A}$.[12] This provides an alternate rationale for using debt-like contracts. Relative to selling the firm, debt-like contracts not only "prepare for the worst" but also "hope for the best" (Dworczak & Pavan, 2022).

Tie-breaking rule. Assuming that the agent selects the principal-preferred search strategy when indifferent appears in the proofs of both Proposition 1 and Theorem 1, but is not essential to the qualitative takeaways of the paper. Unfavorable selection by the agent introduces two potential complications that are easily surmounted by considering contracts very close to pure debt. The first issue is that, after stopping, the agent may

[12] One such example would be if $\mathcal{A}$ included a free, full-support project.

return the lowest prize among several that equally maximize his wage. The second issue is that the agent may choose not to explore the known project if his current wage exactly equals the index of the known project. To overcome these issues, consider a debt-like contract $w(y) = \max\{\epsilon y, y - z\}$, for $\epsilon > 0$. Since this contract is strictly increasing, the agent must return the highest prize that he opens, solving the former issue. Then, by choosing ϵ and z such that the known project has a strictly positive index close enough to zero, the agent cannot stop unless he has searched the known project or produced a prize sufficiently close to paying off the entire debt. Thus, even under unfavorable selection, approximating the debt contract brings the principal arbitrarily close to full surplus extraction.

Principal's information. Our model assumes that the principal knows one of the agent's projects at the outset. A natural question is what happens if the principal knows *more than one* project available to the agent. In such an extension, we can no longer characterize the optimal contracts as in Section 3 because we cannot generally guarantee full surplus extraction.[13] However, it remains true that an optimal contract must include a minimum debt level. Any contract that does not include debt is susceptible to having some known projects crowded out by a safe, low-value alternative—the same failing as the linear contract. Adding a small amount of debt would necessarily improve on this contract because it again discourages the agent from terminating the search too early. In this case, the shape of the contract above the debt level will then depend on the specific parameters of the known projects.

Contracting space, screening, and randomization. To keep the description of the model simple, we assumed that the principal could only contract on the value of a single prize presented by the agent. Depending on the application one has in mind, more sophisticated mechanisms may be available to the principal. For instance, the principal

[13] A characterization is possible only in special cases, e.g., when the projects can be resampled repeatedly (see Section 5.1).

could use a wage contract that depends on multiple presented prizes or the identity of the project that delivered the prize. More generally, the principal may have other ways to indirectly monitor the agent's effort. Alternatively, the principal may attempt to mitigate an adversarial nature by screening the agent's private information through a menu of contracts or using randomization (Kambhampati, 2023; Kambhampati, Peng, Tang, Toikka, & Vohra, 2025). Although they may weakly dominate our debt-like contracts, none of these generalizations can improve the principal's worst-case payoff. No matter what the generalized contracting mechanism allows, the payoff guarantee is bounded above by the case when nature chooses $\mathcal{A} = \mathcal{A}_0$. Because the full surplus of s_0 is already extracted in this case, it is impossible to strictly improve on this payoff without violating the agent's outside option of not searching at all and presenting a prize of 0. So, the contracts identified in Proposition 1 remain optimal even when the principal can use more sophisticated mechanisms.

5 Contract Selection

Theorem 1 demonstrates that many contracts achieve the optimal payoff guarantee in our baseline model. To refine our model's predictions and better understand the circumstances under which different debt-like instruments may arise, we examine several extensions that uniquely select one contract among the ones identified in Theorem 1. Specifically, we argue that (i) pure debt naturally emerges when projects may be resampled; (ii) debt-plus-equity is well-suited to environments with two-sided moral hazard; and (iii) capped-earnout contracts are appropriate when it is desirable to limit the agent's risk exposure. We close each of the following subsections by illustrating some of these contracts' application domains and highlighting how their forms reflect the key forces emphasized in our framework.

To focus on the new forces introduced in these extensions, we henceforth maintain the assumption that the distribution of prizes from the known box, F_0, has full support. Full support helps establish uniqueness in these extensions, even though it does not deliver uniqueness in our baseline model.

5.1 Resampling Projects

An alternative assumption on the agent's search process, often used in the delegated search literature, is to allow the agent to resample any available projects to draw a new prize. One interpretation is that the project represents a repeatable process with the same ex-ante likelihood of success, such as exploratory drilling in an oil field or pharmaceutical experimentation with similar molecular structures. In other words, these are settings in which the principal is confident that the agent could succeed eventually, given enough effort. Extending the model to allow for resampling projects selects pure debt to be uniquely optimal.

Proposition 2. *If projects may be infinitely resampled, then any robustly optimal contract coincides with w_0 almost everywhere.*

Proof. See Appendix A.4. □

The optimality of pure debt follows from essentially the same arguments as Proposition 1. When $\mathcal{A} = \mathcal{A}_0$, the pure debt contract induces the agent to search efficiently while leaving the agent with no surplus. In this case, the agent must continue to search until he generates a prize above the debt level r_0, delivering a payoff of $V_P = r_0$ to the principal.

A rough argument for uniqueness follows from Theorem 1. Although it cannot be directly applied in this setting with repeated sampling, replacing the principal's original payoff guarantee of s_0 with the new payoff of r_0 in the MDL condition uniquely pins down the pure debt contract w_0. Intuitively, the principal must have $y - w(y) \geq r_0$ whenever the agent stops with prize y; otherwise, she is exposed to the possibility of an unknown project that generates y for certain, reducing her payoff. Combining this upper bound on w with the full surplus extraction condition leaves only the pure debt contract.

Although our model takes the perspective of a principal designing a contract for the agent, it also speaks to how an agent could design a security to sell to an investor. If the agent would like to maximize the price at which the security is sold—for example, if he is cash-constrained—but anticipates the investor's caution, the agent's optimal security would coincide with the debt-like contracts we obtain. Indeed, debt is a familiar feature of commercial finance.[14] Speaking more directly to the role of repeated sampling, oil and gas companies commonly utilize debt financing to finance capital-intensive projects, such as exploration and drilling (Clews, 2016). The motivation for using debt is simple: it protects the investor from downside risks while encouraging the borrower's effort until success.

5.2 Principal Moral Hazard

In many environments, the principal ultimately controls how the project's value is realized or reported. For instance, a publisher's negotiations with downstream sellers are critical for generating profits from a creative work. Moreover, the accounting of these profits may be somewhat opaque, combining fixed fees, per-unit prices, and other revenue-sharing agreements. This introduces a moral hazard problem, as the principal may under-monetize the project or divert part of the realized value for private benefit. To capture these forces, we extend the baseline model as follows.

After the agent presents the *actual prize* y, the principal chooses the *nominal prize* $\hat{y} \leq y$. We assume that only the nominal prize $\hat{y}$ can be contracted on. The principal's

[14]For example, government-backed R&D loans, clean energy investments, and subsidized infrastructure projects often rely on debt-like instruments. Government programs like the Small Business Administration's SBIR/STTR programs and others are designed to support a wide range of activities from early-stage research to commercialization of technologies, appealing to the efficiency properties of debt (see Appendix B).

payoff is $\hat{y} - w(\hat{y}) + k(y - \hat{y})$, where the term $k(y - \hat{y})$ captures her private benefit from diversion or shirking at a rate of k. We assume $k \leq 1 - \frac{c_0}{\mathbb{E}_{F_0}(y)}$ to rule out degenerate cases where the principal's diversion motive is too strong and deters the agent from exploring the known project a_0. This condition also implies $k \leq 1$, so any diversion of the prize is socially inefficient. Finally, to keep the analysis focused on the moral hazard dimension, we restrict attention to doubly monotone contracts.[15]

Given a contract w and a realized prize y, the principal's nominal prize $\hat{y}$ solves the following program, whose largest maximizer we denote by $\phi_w(y)$.[16]

$$\max_{\hat{y} \leq y} \ \hat{y} - w(\hat{y}) + k(y - \hat{y}).$$

When contracting, the agent anticipates the principal's possible diversion of the prize. Consequently, his search and participation decisions are shaped by the nominal prize $\hat{y} = \phi_w(y)$ rather than the actual prize y. Since the principal cannot commit to fully monetizing the prize, the agent internalizes the risk of diversion.

We now turn to the design of robustly optimal contracts. Notice that any diversion by the principal is anticipated by the agent, which further weakens the agent's incentives to search. Because diversion is also inefficient, the principal intuitively stands to benefit from writing a contract that ensures she will not divert any realized prize. We call such a contract **diversion-proof** whenever $\phi_w(y) = y$ for all $y \in \mathbb{R}_+$. The following lemma argues that there is no loss of optimality in restricting attention to diversion-proof contracts.

Lemma 1. *For any contract w, there exists a diversion-proof contract $\tilde{w}$ with a higher*

[15] This simplifies the analysis because it guarantees that maximizers will exist and allows us to skip other technical cases. Further, principal monotonicity will follow from optimality, while the agent monotonicity can be micro-founded in several ways, including free disposal.

[16] Because w is doubly monotone, $\phi_w(y)$ is well defined. Further, given full support of F_0, it will also be optimal for the principal to select the largest maximizer in the event that there are multiple solutions, as this will relax the agent's participation constraint.

payoff guarantee, $V_P(w) \leq V_P(\tilde{w})$*. Moreover, a contract is diversion-proof if and only if*

$$\underline{D}w(y) \leq 1 - k, \ \forall y \in \mathbb{R}_+ \tag{4}$$

where $\underline{D}w(y)$ *is the left derivative of* $w(\cdot)$ *evaluated at* y*.*

Proof. See Appendix A.5. □

The proof constructively shows that any contract that permits diversion can be improved by dividing the lost social surplus between the principal and agent. This recapture of surplus also increases the agent's induced index of the known project, improving the payoff guarantee against unknown projects. Lemma 1 also highlights that the principal's marginal profit share must be everywhere above k in any diversion-proof contract. This additional constraint on the principal's contracting problem paves the way for the optimality of debt-plus-equity contracts. Recall that the (z, α)-debt-plus-equity contract has a debt level of z, above which the agent's marginal share is α.

Proposition 3. *Let* $k^* = 1 - \frac{c_0}{\mathbb{E}_{F_0}[y - s_0]^+}$*.*

(1) For any k*, there exists a* (z, α)*-debt-plus-equity contract that is optimal, where*

$$(z, \alpha) = \begin{cases} \left(s_0 \ , \ \frac{c_0}{\mathbb{E}_{F_0}(y - s_0)^+}\right) & \text{if } k < k^*, \\ \left(z_k \ , \ 1 - k\right) & \text{if } k \geq k^*, \end{cases}$$

and z_k *satisfies* $(1 - k)\mathbb{E}_{F_0}(y - z_k)^+ = c_0$*.*

(2) If $k \geq k^*$*, any optimal contract coincides with the above* $\big(z_k, 1 - k\big)$*-debt-plus-equity contract almost everywhere.*

Proof. See Appendix A.6. □

This result describes two cases, depending on whether k is above or below k^*. When $k \leq k^*$, the small diversion benefit is compatible with the debt level of s_0 that protects the principal from unknown projects, as in the baseline model. Therefore, any contract that was originally optimal according to Theorem 1 and keeps the principal's marginal profit share above k is also optimal here. Diversion-proofness is also necessary for optimality because any diversion is inefficient, meaning the principal cannot capture the total surplus s_0. The contract presented in Proposition 3 is the one that maximizes the principal's minimum profit share across prizes, and is thus most resilient to the principal's moral hazard.

When $k \geq k^*$, a debt requirement of s_0 and a large principal share are incompatible because the remaining wages are too small to incentivize the agent's participation. In this case, the principal must trade off leaving rents to the agent, satisfying diversion-proofness, and protecting herself from unknown projects. The $(z_k, 1-k)$-debt-plus-equity contract optimally balances this trade-off: any slackness in the diversion-proof constraint can be repurposed to increase the debt level by an amount that maintains the agent's participation. The constructed debt-plus-equity contract is optimal because it has the highest debt level among all diversion-proof contracts that guarantee the agent's participation.

In practice, two commonly observed contract forms that combine debt with equity participation are advance-against-royalties contracts and convertible debt contracts. Advance-against-royalties contracts are widely used in industries that rely on creative output or speculative development, including publishing, intellectual property licensing, and early-stage pharmaceutical R&D. Under this structure, the agent receives an upfront payment (the advance) but does not earn further income until the project generates sufficient revenue to "earn out" the advance. This contract effectively places the agent in a debt-like position until the threshold is reached, at which point she becomes a partial residual claimant. The structure serves three purposes: it provides liquidity to the agent up front, encourages the agent to generate successful outcomes, and incentivizes the principal to take productive actions even after recoupment. This final feature is especially important in publishing industries where the principal is responsible for downstream negotiations and distribution. A similar structure appears under convertible debt financing, in which the debt-holder can convert the debt amount into equity at a specified rate. These contracts are ubiquitous in venture capital, where investors value upside risk.[17] The combination of debt with optional equity allows investors to protect their downside through debt-like repayment expectations while also having meaningful exposure to upside outcomes. This structure is especially well-suited to environments where performance is uncertain, as it reduces the burden on the financier of assigning too much weight to their forecast of the company's valuation.[18]

5.3 Risk-Averse Agent

Risk neutrality for both the principal and agent contributes to the multiplicity of optimal contracts in Theorem 1 because every fully-extractive contract generates the same surplus and profit. If, instead, the agent is more risk-averse than the principal, we argue that a capped-earnout debt contract proves to be uniquely optimal.

To adjust the baseline setup, we now assume that the agent has a known, strictly increasing, and strictly concave utility function $u(\cdot) : Y \to \mathbb{R}_+$ over monetary rewards. We normalize $u(0) = 0$ and continue to assume that the principal is risk-neutral.

Proposition 4. *The* $(z_u, \bar{w}_u)$*-capped-earnout-debt contract,* $w_u^*(y) = \min\{\bar{w}_u, [y - z_u]^+\}$*, is robustly optimal. The parameters* $(z_u, \bar{w}_u)$ *are determined by the following two conditions:*

1. ***RA - Minimum Debt Level (RA-MDL):*** $z_u = V_{P,u}$ *where* $V_{P,u}$ *is the principal optimal guarantee;*

[17] For example, see YCombinator's introduction of the SAFE "Simple Agreement for Future Equity" contract, which allows investors the option to be repaid either in equity or a fixed nominal amount.

[18] For instance, in 2013, Tesla raised over $600 million by issuing convertible debt following the initial release of the Model S, at a time when investors were increasingly skeptical of the market for luxury electric vehicles. https://www.wsj.com/articles/SB10001424127887324767004578488891523483094

2. ***RA - Full Surplus Extraction (RA-FSE):*** $\mathbb{E}_{F_0}[u(w_u^*(y))] = c_0$.

Any robustly optimal contract coincides with w_u^* *almost everywhere.*

Proof. See Appendix A.7. □

When the agent is risk-averse, the principal has two conflicting objectives. On the one hand, fully insuring the agent (i.e., providing a constant wage contract) is socially efficient and would allow the principal to extract the most surplus. On the other hand, to preserve the agent's search incentive, the principal needs a contract that withholds wages for low prizes.

The structure of the capped-earnout contract shows how the principal balances these objectives. As in the baseline model, if the principal's payoff guarantee is V_P, the principal's robustness concern means she must take at least V_P of the prize from the agent; otherwise, a safe, low-value project would lower her payoff. Subject to this constraint, reducing the agent's exposure to risk allows the risk-neutral principal to appropriate more of the surplus. This leads to the agent's payoff for high-value prizes being constant, as the highest realizations provide little in terms of incentives. The insurance motive also causes the principal to reduce the debt level all the way to V_P because this—with a coordinated reduction in the wage cap—further smooths the agent's wage.

Capped-earnout debt contracts are widely used in performance-based compensation schemes. These contracts are popular in executive compensation packages, structured sales commissions, and acquisition "earnout" contracts, where bonuses or stock options vest only after hitting performance targets and are subject to payout ceilings.[19] Under the name of a "bet-on agreement" or a "valuation adjustment mechanism" (VAM), capped-earnout contracts are also prevalent in the booming Chinese venture capitalist market (Lin, 2020; Xue & Yun, 2022). Frequently, these contracts have clauses wherein the

[19] For example, see "first earnout payment" in the following contract: `https://www.sec.gov/Archives/edgar/data/1468328/000119312510166754/dex993.htm`

startup is required to reimburse the investor—for example, by rebuying shares at their original rate—if the startup fails to meet certain IPO price targets or other milestones. These repurchase clauses place the startup in a debt-like position and guarantee that the full earnout amount (the cap) is realized only under sufficiently strong outcomes. While these contracts are sometimes described as tranches—often used to screen or segment investors with varying degrees of risk aversion—the three regions of our contract have a different economic significance. The initial flat region mimics debt financing, discouraging low-quality effort. The intermediate increasing region encourages exploration as the agent gains from the upside. The final capped region limits windfall rents and protects the principal from overpaying in extremely successful states that provided minimal ex-ante incentives to the agent.

6 Extension to Multiple Agents

In many applications we have described, it is typical for a single principal to interact with many agents (e.g., a publisher sponsoring multiple writers). To capture the new forces at play in this setting, we consider a variation of our baseline model with *multiple agents*, each of whom has potentially *multiple projects*. Formally, we suppose there are m agents indexed by $k = 1, 2, \ldots, m$. Each agent k has access to a set of projects denoted by $\mathcal{A}^k$. As before, we assume the principal only knows one project available to each agent, which we denote by $a_0^k \in \mathcal{A}^k$.

Rather than interact with each agent simultaneously, we allow the principal to costlessly approach and sponsor agents sequentially. Upon being sponsored, an agent searches over his available projects and presents one single prize to the principal. We assume the principal can only *adopt* one presented prize, regardless of how many prizes are presented (e.g., the publisher has limited capacity to publish books presented by several writers, and would like to choose the most profitable one). The principal's realized payoff is the value of the adopted prize minus the wage payments to all agents. This setting features two layers of sequential search: the principal explores among agents, while the agents explore among their available projects. The principal needs to make two decisions: (1) the sponsoring strategy, i.e., which agents to sponsor, in what order, and when to stop; (2) the contracting strategy, i.e., what wage contract to provide each sponsored agent.

Given that the principal interacts with multiple agents, it is natural to allow an agent's wage payment to depend on all the presented prizes, including those presented by other agents. For instance, the principal could use competition between the agents to induce effort. As we will show, such contracting richness ultimately does not benefit the principal. A simple strategy in which each agent's contract only depends on his specific project can achieve the optimal payoff guarantee. Moreover, this contract need not explicitly condition on whose prize is ultimately adopted.

To state the principal's strategy, let r_0^k denote the index of agent k's project a_0^k given by Equation (1). Without loss of generality, we assume $r_0^1 \geq r_0^2 \geq ... \geq r_0^m$.

Proposition 5. *A robustly optimal strategy for the principal can be described by the following dynamic process. In round $k \in \{1, 2, ..., m\}$, the principal sponsors agent k and offers him the r_0^k-debt contract. In each round k, after seeing agent k's presented prize y^k, the principal stops and adopts the highest up-to-date presented prize,* $\max\{y^1, y^2, ..., y^k\}$*, if it is higher than r_0^{k+1}; otherwise, she continues to the next round.*

Proof. See Appendix A.8. □

The argument for optimality is conceptually similar to that of Proposition 1. Specifically, this dynamic strategy generates a principal payoff *as if she searched through all agents' projects herself.* This is clearly an upper bound on her payoff guarantee. Debt again guarantees that $\mathcal{A}^k = \{a_0^k\}$ for all k is the worst case for the principal since any premature stopping by agent k means the debt was fully repaid.

This finding relates to a small literature that studies sequential contracting with multiple agents (Kleinberg, Waggoner, & Weyl, 2016; Durandard, 2023; Ben-Porath, Dekel, & Lipman, 2025). Typically, these problems can be recast as sequential search problems with the principal being the searcher and agents being the "boxes." As in our model, the agents must be appropriately incentivized in order to produce a prize. Where we differ is through the principal's ability to offer agent-specific wage contracts, which allows her to fully separate the incentives of each agent. These contracts, when offered sequentially as in Proposition 5, obviate any strategic interaction among the agents, as each agent essentially faces a single-player decision problem. In particular, each agent's wage payment (upon being sponsored) becomes independent of other agents' strategies.

A Proofs Omitted from Main Text

A.1 Proof of Observation 2

Let w be a doubly monotone contract. Under such a contract, both players prefer a higher realized prize than a lower one, as their payoffs are non-decreasing in the presented prize. So, the agent always presents the highest discovered prize (under our tie-breaking assumption). Hence, the principal benefits from the agent sampling more projects, since it shifts the distribution of the highest realized, hence presented, prize in the sense of first-order stochastic dominance.

To prove the observation, it suffices to show that, if the principal's worst-case scenario is not $\mathcal{A} = \mathcal{A}_0$, then, for all $0 < \epsilon < V_P(w \mid \mathcal{A}_0) - V_P(w)$, there exists $\mathcal{A} = \{a_0, a_1\}$ where $a_1 = (\delta_x, 0)$ with δ_x being a Dirac mass for some x satisfying $w(x) > r_0^w$ such that $V_P(w) + \epsilon \geq V_P(w \mid \mathcal{A})$. In what follows, we construct such a set of projects.

Let $0 < \epsilon < V_P(w \mid \mathcal{A}_0) - V_P(w)$ and let $\{\mathcal{A}_n\}_n$ be a minimizing sequence in Nature's problem:

$$V_P(w) = \inf_{\mathcal{A}' \supseteq \mathcal{A}_0} V_P(w \mid \mathcal{A}') = \lim_{n \to \infty} V_P(w \mid \mathcal{A}_n).$$

By definition of $\{\mathcal{A}_n\}_n$, there exists $N \in \mathbb{N}$ such that, for all $n \geq N$, $V_P(w) \leq V_P(w \mid \mathcal{A}_n) \leq V_P(w) + \frac{\epsilon}{2}$. Observe that, under $\mathcal{A}_N$, there must be a positive probability that the known project a_0 is not sampled; for otherwise our initial observation implies that the principal would be weakly better off under $\mathcal{A}_N$ than under $\mathcal{A}_0$, and

$$V_P(w) + \frac{\epsilon}{2} \geq V_P(w \mid \mathcal{A}_N) \geq V_P(w \mid \mathcal{A}_0) > V_P(w) + \epsilon,$$

a contradiction. So, under $\mathcal{A}_N$, the known project a_0 is not sampled with positive probability.

But, if a_0 is not sampled, the agent must terminate the search for some x that satisfies $w(x) > r_0^w$ in the support of F_i, for some box $(c_i, F_i) \in \mathcal{A}_N$. Since w is non-decreasing, we can then pick x^* such that $w(x^*) > r_0^w$ and

$$x^* \leq \inf\left\{x \in \cup_{(c_i,F_i)\in\mathcal{A}_N} \text{supp}(F_i) \ : \ w(x) > r_0^w\right\} + \frac{\epsilon}{2}.$$

Consider then $\mathcal{A} = \{a_0, a_1\}$ where $a_1 = (\delta_{x^*}, 0)$. Under $\mathcal{A}$, the agent only ever opens box a_1. By definition of x^*, $V_P(w \mid \mathcal{A}) = x^* - w(x^*) \leq V_P(w \mid \mathcal{A}_N) + \frac{\epsilon}{2}$, where we used that $V_P(w \mid \mathcal{A}_N) < V_P(w \mid \mathcal{A}_0)$, hence $x^* - w(x^*) < \mathbb{E}_{F_0}[y - w(y)]$, and that w is doubly monotone. Therefore,

$$V_P(w \mid \mathcal{A}) \leq V_P(w \mid \mathcal{A}_N) + \frac{\epsilon}{2} \leq V_P(w) + \epsilon.$$

This concludes the proof.

A.2 Proof of Observation 3

Proof. Let $w(y) = \alpha y$ with $\alpha \in [0, 1]$ and let r_0^α be the associated w-induced index. By Observation 2,

$$V_P(w) = (1 - \alpha) \min\left\{r_0^\alpha, \ \mathbb{E}_{F_0}[y]\right\},$$

for $\alpha \geq \frac{c_0}{\mathbb{E}_{F_0}[y]}$, and 0 otherwise. Observe that, for $V_P(w) = s_0$, it must be that $\alpha = \frac{c_0}{\mathbb{E}_{F_0}[y]}$. But, for $\alpha = \frac{c_0}{\mathbb{E}_{F_0}[y]}$, $r_0^w = 0$ (when $c_0 > 0$). Moreover, by the implicit function theorem, r_0^α is continuous. Therefore, there exists $\delta > 0$ such that $r_0^\alpha < \epsilon < s_0 - \delta\mathbb{E}_{F_0}[y]$ for all

$\alpha < \frac{c_0}{\mathbb{E}_{F_0}[y]} + \delta$. It follows that, for all $\alpha \in [0,1]$,

$$V_P(w(y) = \alpha y) \leq \begin{cases} 0 \text{ if } \alpha < \frac{c_0}{\mathbb{E}_{F_0}[y]} \\ (1-\alpha)\epsilon \text{ if } \frac{c_0}{\mathbb{E}_{F_0}[y]} \leq \alpha < \frac{c_0}{\mathbb{E}_{F_0}[y]} + \delta \\ (1 - \frac{c_0}{\mathbb{E}_{F_0}[y]} - \delta)\mathbb{E}_{F_0}[y] \text{ otherwise} \end{cases} < s_0 - \delta \mathbb{E}_{F_0}[y],$$

which concludes the proof of Observation 3. □

A.3 Proof of Theorem 1

Part 1: MDL & FSE ⇒ w is optimal. We show that $V_P(w) = s_0$ if the contract w satisfies MDL and FSE. That is, the contract w attains the robustly optimal guarantee. Fix the agent's optimal search strategy and the realized prize $\{y_i\}_{i=0}^n$. Then we consider two cases of the realized search process: (i) a_0 is not sampled, and (ii) a_0 is sampled.

Suppose a_0 is *not sampled*. The FSE condition guarantees that the w-induced index of a_0 is exactly 0. Thus, if the agent did not sample a_0, the agent must have obtained a strictly positive wage. By the MDL condition, this implies that the agent sampled a project whose prize is higher than s_0 and stopped, and, hence, that the principal's payoff is at least s_0 after any history such that a_0 is *not sampled*.

Suppose a_0 is *sampled*. Then, it is also optimal for the agent to stop afterwards without exploring any additional projects because a_0's w-induced index is exactly zero according to the FSE condition. Because the agent breaks ties favorably for the principal, the principal's expected payoff in this case is at least s_0 after any history such that a_0 is *sampled*.

Part 2: w is optimal ⇒ MDL & FSE. We prove the contrapositive. First, suppose a contract w violates FSE. Then

$$s_0 = V_P(w_0|\mathcal{A}_0) > V_P(w|\mathcal{A}_0) \geq V_P(w),$$

where the equality is proved in Proposition 1, the first inequality holds strictly because FSE is violated, and the second inequality holds by the definition of $V_P(w)$. Hence, no contract w that violates FSE is robustly optimal as $V_P(w) < s_0$.

Second, suppose a contract w satisfies FSE but violates MDL, i.e., there exists y' such that $y' - w(y') < s_0$ and $w(y') > 0$. Consider the set of projects $\mathcal{A} = \{a_0, a_1\}$, where $a_1 = (\delta_{y'}, 0)$ is a riskless project. The agent only searches project a_1 because the w-induced index value of the project a_0 is zero by FSE. The principal's corresponding payoff, $y' - w(y')$, is strictly smaller than y', which is weakly smaller than V_P. In other words, $s_0 > V_P(w|\{a_0, a_1\}) \geq V_P(w)$, and no contract that satisfies FSE but violates MDL is robustly optimal.

Thus, no contract that violates MDL or FSE is optimal, and hence, any optimal contract satisfies MDL and FSE.

A.4 Proof of Proposition 2

We have already argued that $w_0 = [y - r_0]^+$ achieves the maximal payoff guarantee in the main text. Here, we show that this optimal contract is essentially unique.

First, recall that under w_0, the agent finds it optimal to continue sampling a_0 until he receives a draw above r_0, at which point he stops. Therefore, the principal's payoff under $w_0(y) = [y - r_0]^+$ is exactly r_0.

Next, for any other contract w' to achieve the payoff guarantee of r_0, we must have $w(y) \leq [y - r_0]^+$ for all y. If not, there exists some y such that $w(y) > [y - r_0]^+ \geq y - r_0$, so $y - w(y) < r_0$. This means adding a free, constant output project with prize y would reduce the principal's payoff below r_0, and hence w' would have a lower payoff guarantee.

Further, for any contract w' to achieve the same payoff guarantee as w_0, it also needs to extract the agent's full surplus, meaning $\mathbb{E}_{F_0}[w'(y)] = c_0$. But, by construction of the index, $\mathbb{E}_{F_0}[w_0(y)] = \mathbb{E}_{F_0}[y - r_0]^+ = c_0$. Since w' is pointwise below w_0 and it has the same expectation, they can only differ on an F_0-null set. Under the full support assumption on F_0, any such F_0-null set has measure zero.

A.5 Proof of Lemma 1

We first prove part 2 of the lemma, relating diversion-proofness to the slope of the contract. Notice that by definition, $\phi_w(y) = y$ if and only if the slope of any line connecting two points on the curve of $y \mapsto y - w(y)$ is always weakly above k. Diversion-proofness is then equivalent to $\frac{y-w(y)-(\hat{y}-w(\hat{y}))}{y-\hat{y}} \geq k$ whenever $\hat{y} \leq y$. This exactly requires the left derivative of $w(y)$ always being weakly below $1-k$, proving the second statement of the lemma.[20]

We now show that any contract that admits diversion is dominated by a diversion-proof contract. Suppose that the contract w is not diversion-proof. If w is not diversion-proof, there exists a disjoint, countable collection of open intervals $(\underline{y}_i, \overline{y}_i)_{i\in\mathbb{N}_+}$ whose union is defined $I := \cup_i(\underline{y}_i, \overline{y}_i)$, such that $\phi_w(y) < y$ for all $y \in I$ and $\phi_w(y) = y$ for all $y \notin I$.[21] In particular, this means for any i and $y \in (\underline{y}_i, \overline{y}_i)$, we have $\phi_w(y) = \underline{y}_i$ and $\underline{y}_i + k(y - \underline{y}_i) > y' + k(y - y')$ for all $y' \in (\underline{y}_i, \overline{y}_i)$. Now consider the contract $\tilde{w}$ that linearizes the wage on each interval $(\underline{y}_i, \overline{y}_i)$. By construction of the intervals and continuity of w, the slope of the wage contract on each interval will be exactly $(1-k)$. Therefore, the principal's marginal share on each $(\underline{y}_i, \overline{y}_i)$ interval will be exactly k, and $\tilde{w}$ will be diversion-proof.

[20]The left derivative is well-defined due to double-monotonicity.
[21]Slightly abusing notation, we allow $\overline{y}_i$ to be ∞.

$$\tilde{w}(y) = \begin{cases} w(\underline{y}_i) + (1-k)(y - \underline{y}_i) & \text{if } y \in (\underline{y}_i, \overline{y}_i), \\ w(y) & \text{if } y \notin I. \end{cases}$$

We now argue that this contract improves the principal's payoff guarantee, $V_P(\tilde{w}) \geq V_p(w)$. First, we claim that $\tilde{w}$ strictly increases the agent's payoff. Because $\tilde{w}$ is diversion-proof while w was not, for any $y \in (\underline{y}_i, \overline{y}_i)$, the agent now receives a wage of $w(\underline{y}_i) + (1 - k)(y - \underline{y}_i)$ which strictly exceeds their wage of $w(\underline{y}_i)$ under w due to the principal's diversion. For all $y \notin \cup_i I$, the agent's wage is unchanged. Therefore, the agent strictly benefits from $\tilde{w}$.

The improvement in the agent's payoff translates to an improved guarantee for the principal. Since $\tilde{w}$ is also double monotone, the principal's payoff guarantee under w and $\tilde{w}$ is determined by either $\mathcal{A} = \{a_0\}$ or $\mathcal{A} = \{a_0, a_1\}$. In the former case, the principal's payoff is unchanged by construction: the $\tilde{w}$ provides the diversion payoff the principal was originally receiving under w. In the latter case, the principal is strictly better off under $\tilde{w}$ because $r_0^{\tilde{w}} > r_0^w$ due to the agent's increased payoff. This means the prize necessary to crowd out the known box must be strictly larger, and $y - \tilde{w}(y)$ is also strictly increasing. Combining the two cases, the principal's payoff guarantee is larger under $\tilde{w}$ than w.

A.6 Proof of Proposition 3

As suggested by the paragraph following Proposition 3, we begin by examining whether any contract characterized in Theorem 1 satisfies the no-diversion condition (4), which states that the principal's marginal share is at least k. Among these contracts, the one least likely to violate (4) is the (z, α)-debt-plus-equity contract with $z = s_0$ and $\alpha = \frac{c_0}{\mathbb{E}_{F_0}(y-s_0)^+}$ — if this contract violates (4), any other contract characterized in Theorem 1 will violate it too. This contract satisfies the condition (4) if and only if $\frac{c_0}{\mathbb{E}_{F_0}(y-s_0)^+} \leq 1-k$, which is equivalent to $k \leq k^*$. By construction, this contract satisfies the MDL and FSE conditions of Theorem 1, so it must be optimal when $k \leq k^*$. Additionally, note that diversion-proofness is a necessary condition for optimality when $k \leq k^*$. Any contract that is not diversion-proof is also inefficient, and therefore fails to capture the entire social surplus.

Now, we turn to the case with $k > k^*$. For any $z \in \mathbb{R}_+$, let w_z denote the $(z, 1-k)$-debt-plus-equity contract. This contract has a debt level of z after which the principal's marginal share is k. Of particular importance is the debt-plus-equity contract w_{z_k}, where the debt level z_k satisfies $(1-k)\mathbb{E}_{F_0}(y-z_k)^+ = c_0$. We will argue that this contract is uniquely optimal.

Step 1: The debt-plus-equity contract w_{z_k} is optimal.

Under the contract w_{z_k}, the principal's worst-case scenario is $\mathcal{A} = \{a_0, a_1\}$ with a_1 being a project with zero exploration cost and a deterministic prize of $z_k + \epsilon$, and her payoff guarantee from w_{z_k} is thus z_k, $V_P(w_{z_k}) = z_k$.

We now argue that z_k is an upper bound on the principal's payoff guarantee. Without loss of optimality (Lemma 1), consider any diversion-proof contract w. Let r_0^w be the induced index of a_0, and let $y^w := \sup\{y \mid w(y) = r_0^w\}$ be the maximal prize that makes the agent indifferent between searching a_0 and not. The following inequalities relate the contracts w and w_{z_k}.

$$\begin{aligned}
\int (1-k)\left[y - z_k\right]^+ dF_0(y) = c_0& \\
&= \int \left[w(y) - r_0^w\right]^+ dF_0(y) \\
&\leq \int \left[w(y^w) + (1-k)(y - y^w) - r_0^w\right]^+ dF_0(y) \\
&= \int (1-k)\left[y - y^w\right]^+ dF_0(y)
\end{aligned}$$

The first two equalities follow from the definitions of z_k and r_0^w, respectively. The inequality follows because the contract is diversion-proof; therefore, the agent's marginal share above y^w can be at most $(1-k)$. The final equality is by definition of y^w.

Comparing the first and last expression above implies $y^w \leq z_k$, which further implies $y^w - w(y^w) \leq z_k$ due to limited liability. Finally, since w is doubly monotone and hence continuous, we know

$$V_P(w) \leq \inf_{\epsilon}\ y^w + \epsilon - w(y^w + \epsilon) \leq z_k.$$

Notice the middle expression is exactly the principal's payoff when the agent has access to an unknown box that freely produces prize $y^w + \epsilon$. Therefore, the payoff guarantee for any contract w is at most z_k. Since w_{z_k} achieves this guarantee, it is optimal.

Step 2: Any optimal contract must be debt-plus-equity.

By contradiction, suppose there is a contract $w \neq w_{z_k}$ that satisfies (i) the condition (4) and (ii) $\mathbb{E}_{F_0}(w(y)) \geq c_0$. We construct a debt-plus-equity contract that always outperforms it.

If $r_0^w = 0$, the debt level of w, $\inf\{y|w(y) > 0\}$, must be strictly lower than z_k; otherwise, (i) and (ii) in the last paragraph cannot both be satisfied. Hence, the principal's payoff guarantee under w must also be strictly lower than z_k, so the debt-plus-equity contract w_{z_k} strictly outperforms w.

If $r_0^w > 0$, we let $y^w := \sup\{y|w(y) = r_0^w\}$. This prize y^w is the maximal prize in hand that makes the agent indifferent between exploring a_0 and not. The principal's worst-case scenario under w is either $\mathcal{A} = \{a_0\}$ or $\mathcal{A} = \{a_0, a_1\}$ with a_1 being a project with zero exploration cost and a deterministic prize of $y^w + \epsilon$.

We now construct a contract that improves on w in two steps, as depicted in Figure 2. First, consider a debt-plus-equity contract $w_{z'}$ that satisfies $w_{z'}(y^w) = w(y^w)$. This condition ensures that if the agent were to follow the same stopping behavior (in terms of y), the principal's payoff guarantee is weakly higher. The contract $w_{z'}$ single-crosses w at

y^w: $w_{z'}(y) \le w(y)$ for $y < y^w$ and $w_{z'}(y) \ge w(y)$ for $y > y^w$. Hence, by the definition of w-induced index, we can infer that $r_0^{w_{z'}} \ge r_0^w$. Since $r_0^w = w(y^w) = w_{z'}(y^w)$, it follows that $r_0^{w_{z'}} \ge w_{z'}(y^w)$. This means that the agent finds a_0 more attractive under $w_{z'}$ than under the original w and requires a higher value of y to not search a_0.

Next, we construct another debt-plus-equity contract $w_{z''}$ such that $z'' \ge z'$ and $r_0^{w_{z''}} = w_{z''}(y^w)$. In other words, this contract (i) has a weakly larger debt level than $w_{z'}$ and (ii) faces the same project a_1 that yields the principal the lowest payoff if crowding out a_0 as w. Such a z'' exists because both $r_0^{w_z}$ and $w_z(y^w)$ are continuous in z, $w_z(y^w) \ge 0$ for any z, and $r_0^{w_z} \le 0$ when z is sufficiently large. Since $z'' \ge z'$, this also implies $w_{z''}(y) \le w_{z'}(y) \le w(y)$ for all $y \le y^w$.

Now, we are ready to show that the debt-plus-equity contract $w_{z''}$ outperforms w regardless of whether a_0 is crowded out or explored. If a_0 is not crowded out (i.e., $\mathcal{A} = \{a_0\}$), $w_{z''}$ is better than w because the expected wage paid to the agent has decreased.

$$\begin{aligned}
\int w(y)dF_0(y) &= \int_0^{y^w} w(y)dF_0(y) + \int_{y^w}^{\infty} w(y)dF_0(y) \\
&= \int_0^{y^w} w(y)dF_0(y) + c_0 + r_0^w(1 - F_0(y^w)) \\
&\ge \int_0^{y^w} w_{z''}(y)dF_0(y) + c_0 + r_0^{w_{z''}}(1 - F_0(y^w)) \\
&= \int_0^{y^w} w_{z''}(y)dF_0(y) + \int_{y^w}^{\infty} w_{z''}(y)dF_0(y) \\
&= \int w_{z''}(y)dF_0(y).
\end{aligned}$$

The second line follows by definition of y^w and the definition of r_0^w. The third line follows because $w_{z''}(y) \le w(y)$ for $y \le y^w$ and $r_0^{w_{z''}} \le r_0^w$. The inequality is strict if $w \ne w_{z''}$ almost everywhere. If a_0 is crowded out, notice that the project a_1 that yields the principal the lowest payoff if crowding out a_0 is the same for w and $w_{z''}$ — the project with no exploration cost and a deterministic prize of $y^w + \epsilon$. Since $w_{z''}(y^w) \le w(y^w)$, we know that the principal is better off under $w_{z''}$ compared to w.

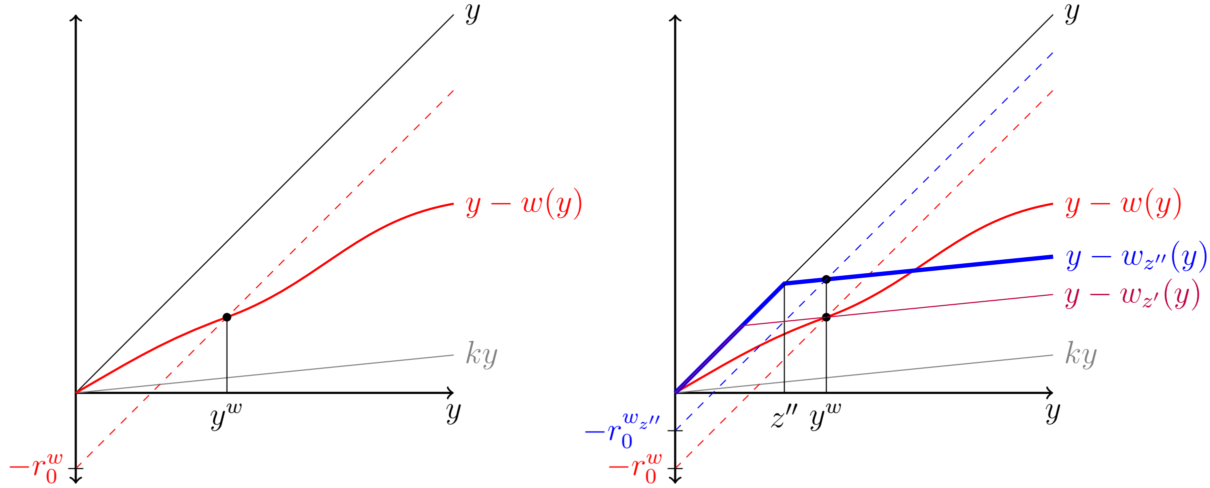

(a) Initial contract

(b) Constructed debt-plus-equity contracts

Figure 2: Illustration of the optimality of a debt-plus-equity contract. Panel (a) depicts an arbitrary initial debt contract w. Panel (b) depicts the debt-plus-equity contracts $w_{z'}$ and $w_{z''}$ that we constructed.

Step 3: w_{z_k} is best among all the debt-plus-equity contracts.

Recall that w_{z_k} is the $(z_k, 1-k)$-debt-plus-equity contract that features a debt level of z_k satisfying $(1-k)\mathbb{E}_{F_0}(y - z_k)^+ = c_0$, and a principal marginal share of k. Notice that for any debt-plus-equity contract w_z with $z < z_k$, the corresponding w_z-induced index for the known project, $r_0^{w_z}$, satisfies $\mathbb{E}_{F_0}\left[\left((1-k)(y-z)^+ - r_0^{w_z}\right)^+\right] = c_0$. Since

$$\mathbb{E}_{F_0}\left[\left((1-k)(y-z)^+ - r_0^{w_z}\right)^+\right] = \mathbb{E}_{F_0}\left[(1-k)\left(y - z - \frac{r_0^{w_z}}{1-k}\right)^+\right],$$

while $r_0^{w_{z_k}} = 0$ by definition, it follows that $r_0^{w_z} = (1-k)(z_k - z)$. Hence, the worst project to crowd out a_0 is the one with zero exploration cost and a deterministic prize of $\frac{r_0^{w_z}}{1-k} + z + \epsilon = z_k + \epsilon$, which happens to be exactly the same as the worst project to crowd out a_0 under w_{z_k}. Since $w_{z_k}(z_k) < w_z(z_k)$ for all $z < z_k$, the principal is strictly better off under w_{z_k} compared to w_z, regardless of whether a_0 is crowded out or not. Thus, the contract w_{z_k} is uniquely optimal.

A.7 Proof of Proposition 4

The overall proof proceeds in three steps. First, we prove that any robustly optimal contract *among doubly monotone contracts* coincides with the $(z_u, \bar{w}_u)$-capped-earnout-debt contract defined in the statement of Proposition 4, $w_u^*(y) = \min\{\bar{w}_u, [y - z_u]^+\}$, on the support of F_0, *when the support of* F_0 *is finite.* Next, we show that non-monotone contracts cannot improve the principal's payoff *when the support of* F_0 *is finite.* Finally, we show that w_u^* is optimal among all contracts even when the support of F_0 is infinite.

Step 1. Suppose first that $\#supp(F_0) < \infty$. The proof proceeds in five steps:

- Step 1.1 We show that any optimal contract must satisfy the modified MDL condition:

$$w(y) \leq \max\left\{u^{-1}\left(r_0^{w,u}\right), y - V_{P,u}\right\}. \tag{m-MDL}$$

- Step 1.2 We show that for any such doubly monotone contract, in the worst-case scenario, the set of available projects is $\mathcal{A}_0$.

- Step 1.3 We show that any robustly optimal contract under $\mathcal{A}_0$ takes the form

$$t^{\underline{w}}(y) = \min\left\{\underline{w}, \max\left\{u^{-1}\left(r_0^{t^{\underline{w}},u} \vee 0\right), y - V_{P,u}\right\}\right\}, \tag{5}$$

  for some constant $\underline{w}$, on the support of F_0.

- Step 1.4 We show that any robustly optimal candidate coincides with w_u^* on the support of F_0.

- Step 1.5 So far, we have shown that if a robustly optimal contract exists, it coincides with w_u^* on the support of F_0. To conclude the proof, we show that a robustly optimal contract exists.

□

Proof of Step 1.1. We prove the contrapositive. Let w be a contract that violates (m-MDL). We show that w is not optimal. Since w violates (m-MDL), there exists y^* such that $w(y^*) > u^{-1}(r_0^{w,u}) \Leftrightarrow u(w(y^*)) > r_0^{w,u}$ and $y^* - w(y^*) < V_{P,u}$. Consider the set of projects $\mathcal{A} = \{a_0, a_1\}$ where $a_1 = (\delta_{y^*}, 0)$ is a safe project that yields y^* with probability 1. Under contract w, since $u(w(y^*)) > r_0^{w,u}$, the agent faced with projects $\mathcal{A}$ only ever explore project a_1. So, $V_P(w, \mid \mathcal{A}) = y^* - w(y^*) < V_{P,u}$. But, by definition, $V_P(w) \leq V_P(w, \mid \mathcal{A})$. Thus, w is not robustly optimal. This concludes the first step. □

Proof of Step 1.2. Let w be a contract such that Eq. (m-MDL) holds, and observe that it implies that, for all y, either $y - w(y) \geq V_{P,u}$ or $u(w(y)) < r_0^{w,u}$. Let $\mathcal{A} \supseteq \mathcal{A}_0$ be a set of available projects. We distinguish two sets of histories. After all the histories such that the agent does not open a_0, then the principal gets at least $V_{P,u}$. After all the histories such that the agent opens a_0, under double monotonicity, the agent presents a prize whose distribution first-order stochastically dominates F_0. So, in both cases, the principal's payoff exceeds $V_P(w \mid \mathcal{A}_0)$. □

Proof of Step 1.3. We show the contrapositive. Let w be a (doubly monotone) contract that does not take the form of a generalized capped-earnout-debt contract, e.g., (5), on the support of F_0. We show that w is not robustly optimal. By Step 1.1, if w does not satisfy Eq. (m-MDL), it is not robustly optimal. So, assume that it does. Thus, by Step 1.2, in the worst-case scenario, $\mathcal{A} = \mathcal{A}_0$. Then, note that, if w does not incentivize search under $\mathcal{A}_0$, it is not optimal. So, assume that it does.

Since w does not take the form of a generalized capped-earnout-debt contract, e.g., (5), there exists $y_1 \in supp(F_0)$ and $\epsilon_1 > 0$ such that $w(y_1) < \max\{u^{-1}(r_0^{w,u}), y - V_{P,u}\}$. WLOG, let $y_1 = \min\{y \in supp(F_0) : w(y) + \epsilon_1 < \max\{u^{-1}(r_0^{w,u}), y - V_{P,u}\}\}$. Moreover, since w does not take the form of a generalized capped-earnout-debt contract, e.g., (5), and w is doubly monotone, there exists $y_2 \in supp(F_0)$, $y_2 > y_1$ such that

$w(y_2) > w(y_1) + \epsilon_2$ for some $\epsilon_2 > 0$. WLOG, let $y_2 = \min\{y \in supp(F_0)\ : w(y_2) > w(y_1)\}$.

Let $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, and consider the alternative contract $\tilde{w}$ defined by

$$\tilde{w}(y) = \begin{cases} w(y) \text{ if } y < y_1 \\ w(y) + k_1 \text{ if } y_1 \leq y < y_2 \\ w(y) - k_2 \text{ if } y_2 \leq y, \end{cases}$$

where $0 < k_1, k_2 < \epsilon$ are chosen so that

1. $\mathbb{E}_{F_0}[w(y)] = \mathbb{E}_{F_0}[\tilde{w}(y)]$, and

2. $\mathbb{E}_{F_0}[u(w(y))] < \mathbb{E}_{F_0}[u(\tilde{w}(y))]$.

That such k_1, k_2 exist follows from the strict concavity of the agent's utility since the distribution of $w(y)$ dominates the distribution of $\tilde{w}(y)$ in the convex order: $w_{\#}F_0 \succ_{cvx} \tilde{w}_{\#}F_0$.[22] Increasing k_2 slightly (such that, still, $k_2 < \epsilon$), with a small abuse of notation, we can assume that

1. $\mathbb{E}_{F_0}[\tilde{w}(y)] < \mathbb{E}_{F_0}[\tilde{w}(y)]$, and

2. $\mathbb{E}_{F_0}[u(w(y))] < \mathbb{E}_{F_0}[u(\tilde{w}(y))]$.

Moreover, since the support of F_0 is finite and $k_1, k_2 \leq \epsilon$, there exists a doubly monotone contract satisfying (m-MDL) that coincides with $\tilde{w}$ on the support of F_0. With a second small abuse of notation, let $\tilde{w}$ denote this doubly monotone contract. Intuitively, $\tilde{w}$ flattens w to insure the agent, increasing both the total surplus and the principal's share.

Note that, if the agent was willing to explore the known box under w, he is still willing

[22] We use $w_{\#}F_0$ and $\tilde{w}_{\#}F_0$ to denote the push-forward distributions of F_0 under the contracts w and $\tilde{w}$.

to do so under the alternative contract $\tilde{w}$. Therefore,

$$V_P(w \mid \mathcal{A}_0) < V_P(\tilde{w} \mid \mathcal{A}_0).$$

But, $\tilde{w}$ is doubly monotone and satisfies Eq. (m-MDL). Hence, $V_P(\tilde{w} \mid \mathcal{A}_0) = V_P(\tilde{w})$ by Step 1.2, and w is not robustly optimal. □

Proof of Step 1.4. Taking stock, we have that

$$V_P = \sup_{\underline{w}} V_P(t^{\underline{w}}) = \sup_{\underline{w}} V_P(t^{\underline{w}} \mid \mathcal{A}_0) > V_P(w),$$

for all w that does not take the form of a generalized capped-earnout-debt contract as defined by Eq. (5), on the support of F_0. Therefore, to prove that w_u^* is (F_0-essentially) uniquely robustly optimal, it suffices to show that w_u^* uniquely maximizes $V_P(t^{\underline{w}} \mid \mathcal{A}_0)$ among generalized capped-earnout-debt contracts, which follows from the two observations below.

First, the principal's payoff under $\mathcal{A}_0$ is maximized by a generalized capped-earnout-debt contract, $t^{\underline{w}}$, such that

$$\mathbb{E}_{F_0}[u(t^{\underline{w}}(y))] \geq c_0,$$

since $\mathbb{E}_{F_0}[u(t^{\underline{w}}(y))] < c_0 \Rightarrow V_P(t^{\underline{w}} \mid \mathcal{A}_0) = 0$.

Second, the capped-earnout-debt contract w_u^* defined in Proposition 4 is the (pointwise) smallest contract satisfying the above inequality (on the support of F_0). To see this, observe that for $\underline{w} < \bar{w}_u$, the above inequality is violated if the contract $t^{\underline{w}}$ differs from

w_u^* on the support of F_0, while for larger values of $\underline{w} > \bar{w}_u$,

$$t^{\underline{w}}(y) = \min\left\{\underline{w}, \max\left\{u^{-1}(r_0^{t^{\underline{w}},u}) \vee 0, y - V_{P,u}\right\}\right\}$$
$$\geq \min\{\bar{w}_u, \max\{0, y - V_{P,u}\}\} = w_u^*(y).$$

Moreover, the inequality is strict for $y \geq V_{P,u} + \bar{w}_u$. Thus w_u^* maximizes the principal guarantee under $\mathcal{A}_0$, uniquely on the support of F_0. □

Proof of Step 1.5. The set of doubly monotone contracts that the principal may be willing to offer is included in the set of all nondecreasing 1-Lipschitz continuous functions on $[0, \bar{y}]$ that take values in $[0, 2\max\{y \,:\, y \in supp(F_0)\}]$. So, it is (sequentially) compact by Arzelà-Ascoli (for the locally uniform semi-norms, and hence, in the product topology as well). By Weierstrass maximum theorem, it then suffices to show that the mapping

$$\mathcal{C}^{0,1} \to \mathbb{R},\ w \to \inf_{\mathcal{A} \supseteq \mathcal{A}_0} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right]$$

is upper semi-continuous to ensure that a robustly optimal contract exists. This follows from Lemma 3 below. Hence, there exists a robustly optimal contract (among doubly monotone contracts). □

This concludes Step 1 of the proof: we showed that w_u^* is (F_0-essentially) uniquely optimal *among doubly monotone contracts*.

Step 2. Next, we show that non-monotone contracts cannot improve the principal's payoff, and hence, that w_u^* is optimal *among all contracts*, still under the assumption that the support of F_0 is finite. Observe that

$$\sup_{w} \inf_{\mathcal{A}\supseteq\mathcal{A}_0} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right]$$

$$\geq \sup_{K\geq 0} \sup_{w\,:\,w\leq K \text{ and (m-MDL) and double monotonicity holds}} \inf_{\mathcal{A}\supseteq\mathcal{A}_0} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right]$$

$$= \mathbb{E}_{\sigma(\mathcal{A}_0,w_u^*)}\left[y - w_u^*(y)\right]$$

$$= \sup_{K\geq 0} \sup_{w\,:w\leq K \text{ and } w \text{ satisfies (m-MDL) on } supp(F_0)} \mathbb{E}_{\sigma(\mathcal{A}_0,w)}\left[y - w(y)\right]$$

$$\geq \sup_{K\geq 0} \sup_{w\,:\,w\leq K \text{ and } w \text{ satisfies (m-MDL) on } supp(F_0)} \inf_{\mathcal{A}\supseteq\mathcal{A}_0\,:\,\forall (F,c)\in\mathcal{A},\ supp(F)\subset supp(F_0)} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right]$$

$$= \sup_{K\geq 0} \sup_{w\,:\,w\leq K \text{ on } supp(F_0)} \inf_{\mathcal{A}\supseteq\mathcal{A}_0\,:\,\forall (F,c)\in\mathcal{A},\ supp(F)\subset supp(F_0)} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right]$$

$$\geq \sup_{w} \inf_{\mathcal{A}\supseteq\mathcal{A}_0} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right].$$

The three inequalities are straightforward. The first equality holds by Step 1. The second equality holds by Lemma 2 below. The third equality holds by Step 1.1 of the proof, provided that a robustly optimal contract exists in the subproblem:

$$\sup_{w\,:\,w\leq K \text{ and } w \text{ satisfies (m-MDL) on } supp(F_0)} \inf_{\mathcal{A}\supseteq\mathcal{A}_0\,:\,\forall (F,c)\in\mathcal{A},\ supp(F)\subset supp(F_0)} \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right].$$

To see this, recall that Step 1.1 shows that any optimal contract must satisfy (m-MDL). Existence follows from Lemma 3 and Weierstrass maximum theorem since the set of contracts is payoff equivalent to a finite-dimensional compact set under our finite support assumption, which proves the equality.

As a result, w_u^* is optimal among all contracts *when the support of F_0 is finite.*

Step 3. Finally, we show that w_u^* is optimal among all contracts even *when the support of F_0 is infinite.* The proof is by contradiction. Suppose that w_u^* is not optimal. Then, there exists $\bar{w}$ such that

$$V_p(\bar{w}) > V_p(\bar{w}) + \epsilon = \mathbb{E}_{F_0}\left[y - w_u^*(y)\right] + \epsilon.$$

for some $\epsilon > 0$. The equality follows from Step 1.2 since w_u^* is doubly monotone and satisfies (m-MDL). By Lusin's theorem, there exists a closed subset K of $\mathbb{R}_+$ such that the restriction of $\bar{w}$ to C is continuous and $\int_{\mathbb{R}_+ \setminus K} \max\{(y - \bar{w}(y)), 1\} dF_0 < \frac{\epsilon}{6}$. Let $r_0^{\bar{w},u}$ denote the index of the known box under $\bar{w}$, and observe that, WLOG, we can assume that $r_0^{\bar{w},u} > 0$. (Otherwise, consider $\bar{w}(y) = \bar{w}(y) + \frac{\epsilon}{2}$.)

Since the set of distributions with finite support is dense in $\Delta(K)$ when equipped with the weak* topology, there exists a sequence of distributions $(F_0^n)_n$ with finite support such that

1. $supp(F_0^n) \subset K$, and

2. $F_0^n \rightharpoonup^* F_0(\cdot \mid y \in K)$.

Consider then the sequence of known projects $\mathcal{A}_0^n = (F_0^n, c_0^n)$, where we set $c_0^n = \mathbb{E}_{F_0^n}\left[(\bar{w}(y) - r_0^{\bar{w},u})^+\right]$. By Step 2, for all $n \in \mathbb{N}$, the $(z_u^n, \bar{w}_u^n)$-capped-earnout-debt contract defined in the statement of Proposition 4, $w_u^{*,n}$, is optimal, and, hence,

$$V_P^n = V_P^n(w_u^{*,n}) = \mathbb{E}_{F_0^n}[y - w_u^{*,n}(y)] \geq V_P^n(\bar{w}).$$

Therefore, for all $n \in \mathbb{N}$, there exists $\mathcal{A}^n \supset \mathcal{A}_0^n$ such that

$$\mathbb{E}_{\sigma(\mathcal{A}^n, \bar{w})}[y - \bar{w}(y)] \leq V_P^n(w_u^{*,n}) + \frac{\epsilon}{3} = \mathbb{E}_{F_0^n}[y - w_u^{*,n}(y)] + \frac{\epsilon}{3}.$$

Since $F_0^n \rightharpoonup^* F_0(\cdot \mid K)$ and $y - w_u^{*,n}(y) \to y - w_u^*(y)$, locally uniformly, there exists $N \in \mathbb{N}$

such that, for all $n \geq N_1$,

$$\mathbb{E}_{\sigma(\mathcal{A}^n,\bar{w})}[y-\bar{w}(y)] \leq V_P(w_u^*) + \frac{2\epsilon}{3} = \mathbb{E}_{F_0}[y - w_u^*(y)] + \frac{2\epsilon}{3}.$$

Moreover, since $r_0^{\bar{w},u} > 0$, we can choose the set $\mathcal{A}^n$ such that the indices of all boxes are strictly ordered by altering the sampling costs. Therefore, replacing F_0 by F_0^n, does not change the agent's strategy, and there exists N_2 such that for all $n \geq N_2$,

$$\mathbb{E}_{\sigma(\tilde{\mathcal{A}}^n,\bar{w})}[y-\bar{w}(y)] \leq \mathbb{E}_{\sigma(\mathcal{A}^n,\bar{w})}[y-\bar{w}(y)] + \frac{\epsilon}{3},$$

where $\tilde{\mathcal{A}}^n = \mathcal{A}^n \setminus \mathcal{A}_0^n \cup \mathcal{A}_0$ and we used that $F_0^n \rightharpoonup^* F_0(\cdot \mid K)$, that the restriction of $\bar{w}$ to K is continuous, and that $\int_{\mathbb{R}_+ \setminus K} \max\{(y-\bar{w}(y)), 1\} dF_0 < \frac{\epsilon}{6}$. Thus, for $n \geq \max\{n_1, N_2\}$,

$$\mathbb{E}_{\sigma(\tilde{\mathcal{A}}^n,\bar{w})}[y-\bar{w}(y)] \leq \mathbb{E}_{\sigma(\mathcal{A}^n,\bar{w})}[y-\bar{w}(y)] + \frac{\epsilon}{3} \leq \mathbb{E}_{F_0}[y - w_u^*(y)] + \epsilon = V_P(w_u^*) + \epsilon,$$

a contradiction. Therefore, w_u^* is optimal.

It only remains to show that it is uniquely optimal under the full support assumption. This follows from Step 1.1 and Lemma 2, since we have shown that a robustly optimal contract exists, and hence, any optimal contract must satisfy (m-MDL). But, then, for any contract w satisfying (m-MDL),

$$V_P(w) \leq V_P(w \mid \mathcal{A}_0) < V_P(w_u^* \mid \mathcal{A}_0) = V_P,$$

by Lemma 2, under our full support assumption. Thus, w_u^* is the unique robustly optimal contract.

A.7.1 Supporting Lemmas for the proof of Proposition 4

Lemma 2. *w_u^* is (F_0-essentially) uniquely optimal among contracts satisfying* (m-MDL) *under $\mathcal{A}_0$:*

$$\mathbb{E}_{\sigma(\mathcal{A}_0,w_u^*)}\left[y - w_u^*(y)\right] = \sup_{w\,:\,w \text{ satisfies (m-MDL)}} \mathbb{E}_{\sigma(\mathcal{A}_0,w)}\left[y - w(y)\right].$$

Proof. By Step 1.4, we know that, if the contract takes the form of a generalized capped-earnout-debt contract, then it yields a payoff strictly lower than $\mathbb{E}_{\sigma(\mathcal{A}_0,w_u^*)}\left[y - w_u^*(y)\right]$ to the principal. So, to prove the lemma, we only need to prove that (i) no contract satisfying (m-MDL) that does not coincide with a generalized capped-earnout-debt contract, e.g., (5), on the support of F_0 can be optimal under $\mathcal{A}_0$, and (ii) that an optimal contract exists.

(i) Let w be a contract that does not coincide with a generalized capped-earnout-debt contract, e.g., (5), on the support of F_0. We show that w is not optimal under $\mathcal{A}_0$ among the contracts that satisfy Eq. (m-MDL).

Note first that, if w does not incentivize search under $\mathcal{A}_0$, it is not optimal. So, assume that it does. Since w does not take the form of a generalized capped-earnout-debt contract, e.g., (5), there exists $y_1 < y_2$ such that $w(y_1) < \max\left\{u^{-1}\left(r_0^{w,u}\right), y - V_{P,u}\right\}$ and $w(y') \neq w(y_1)$ on a set of positive measure in $[y_1, y_2]$ under F_0. Let Y^+ be the subset of $[y_1, y_2]$ such that $w(y') > w(y_1)$ and Y^- be the subset of $[y_1, y_2]$ such that $w(y') < w(y_1)$. $Y^+ \cup Y^-$ has positive measure under F_0. Consider then the alternative contract $\tilde{w}$ defined by

$$\tilde{w}(y) = \begin{cases} \bar{w} \text{ if } y \in Y^+ \cup Y^- \subset [y_1, y_2], \\ w(y) \text{ otherwise}, \end{cases}$$

where $\bar{w} \in \mathbb{R}$ is chosen so that

1. $\mathbb{E}_{F_0}\left[w(y) \mid y \in Y^+ \cup Y^-\right] > \bar{w}$, and

2. $\mathbb{E}_{F_0}[u(w(y)) \mid y \in Y^+ \cup Y^-] < u(\bar{w})$.

Such a $\bar{w}$ exists by Jensen's inequality as the agent's utility is strictly concave, and hence, the Jensen gap is strictly positive. Moreover, we can choose $\bar{w}$ so that $\tilde{w}$ satisfies Eq. (m-MDL) by considering subsets of $Y^+ \cup Y^-$ if needed. Intuitively, $\tilde{w}$ flattens w on $Y^+ \cup Y^-$ to insure the agent, increasing both the total surplus and the principal's share. Therefore, if the agent was willing to explore the known box under w, he is still willing to do so under the alternative contract $\tilde{w}$, and,

$$V_P(w \mid \mathcal{A}_0) < V_P(\tilde{w} \mid \mathcal{A}_0).$$

But, $\tilde{w}$ satisfies Eq. (m-MDL). So, w is not optimal among the contracts that satisfy Eq. (m-MDL) under $\mathcal{A}_0$.

(ii) To prove that an optimal contract exists, we first show that we can restrict attention to contracts w such that $r_0^w = 0$. To see this, let w be a contract satisfying (m-MDL) such that $r_0^w > 0$. Consider then the alternate contract $\tilde{w}$ implicitly defined by

$$u(\tilde{w}(y)) = \left(u(w(y)) - r_0^w\right)^+, \quad \forall y.$$

Then, $\tilde{w} \leq w$, hence, $\mathbb{E}_{F_0}[y - w(y)] \leq \mathbb{E}_{F_0}[y - \tilde{w}(y)]$, and, by definition, $r_0^{\tilde{w}} = 0$. So, $\tilde{w}$ weakly dominates w. The result then follows if $\tilde{w}$ is admissible, i.e., $\tilde{w}$ satisfies (m-MDL).

But,

$$\begin{aligned}\tilde{w}(y) &= u^{-1}\left((u(w(y)) - r_0^w)^+\right)\\ &\leq \max\left\{0, u^{-1}\left((u(w(y)) - r_0^w)\right)\right\}\\ &\leq \max\left\{0, w(y) - u^{-1}(r_0^w)\right\}\\ &\leq \max\left\{0, y - V_{P,u} - u^{-1}(r_0^w)\right\}\\ &\leq \max\left\{0, y - V_{P,u}\right\}.\end{aligned}$$

where the second inequality follows from the convexity of u^{-1} and the third inequality from w satisfying (m-MDL). So, $\tilde{w}$ satisfies (m-MDL), and, hence, we can look for an optimal contract in the subset

$$\mathcal{C} = \{w \geq 0 \ : \ w(y) \leq \max\{0, y - V_P\}\}.$$

Existence then follows from Lemma 10 in Durandard and Ghersengorin (2024) since $\mathcal{C}$ is convex. This concludes the proof of Lemma 2. □

Lemma 3. *For all superset $\mathbb{A}$ of set of boxes, the mapping*

$$\mathcal{C}^0([0,\bar{y}]) \subset \mathbb{R}_+^{[0,\bar{y}]} \to \mathbb{R}, \ w \to \inf_{\mathcal{A}\in\mathbb{A}} \mathbb{E}_{\sigma(\mathcal{A},w)}[y - w(y)]$$

is sequentially upper semi-continuous when $\mathbb{R}_+^{[0,\bar{y}]}$ is equipped with the product topology (i.e., for pointwise convergence).

Proof. The infimum of a family of upper semi-continuous functions is upper semi-continuous. So, the results follows if $w \to \mathbb{E}_{\sigma(\mathcal{A},w)}[y - w(y)]$ is upper semi-continuous for all $\mathcal{A}$, which we prove now.

Let $\mathcal{A} \supseteq \mathcal{A}_0$. Let $\{w_n\}_{n\in\mathbb{N}}$ be a sequence of contract that converges to some w^*. We show that

$$\limsup_{n\to\infty} \mathbb{E}_{\sigma(\mathcal{A},w_n)}\left[y - w_n(y)\right] \leq \mathbb{E}_{\sigma(\mathcal{A},w^*)}\left[y - w^*(y)\right].$$

Observe first that, if $\sigma(\mathcal{A}, w_n) \overset{weak^*}{\rightarrow} \sigma(\mathcal{A}, w^*)$, where we denote the distribution of the presented prize under $\mathcal{A}$ and w by $\sigma(\mathcal{A}, w)$ with a small abuse of notation, the above inequality holds as an equality as the bilinear form $(\sigma, w) \to \mathbb{E}_\sigma\left[y - w(y)\right]$ is continuous. So, suppose that $\sigma(\mathcal{A}, w_n) \overset{weak^*}{\not\rightarrow} \sigma(\mathcal{A}, w^*)$. Since the agent is rational, he follows the Weitzmann index strategy. From the definition of the indices and the implicit function theorem, the mapping $w \to r^w$ is continuous for each box. As a result, $\sigma(\mathcal{A}, w_n) \overset{weak^*}{\not\rightarrow} \sigma(\mathcal{A}, w^*)$ only if $r_i^{w^*} = r_j^{w^*}$ for some boxes i and j. That is, the agent is indifferent between two boxes under contract w^*. But, by assumption, when that's the case, the agent breaks the tie in favor of the principal. Therefore, the above inequality also holds in that case.

Since $\mathcal{A} \supseteq \mathcal{A}_0$ was arbitrary, we have shown that $w \to \mathbb{E}_{\sigma(\mathcal{A},w)}\left[y - w(y)\right]$ is upper semi-continuous for all $\mathcal{A} \supseteq \mathcal{A}_0$. □

A.8 Proof of Proposition 5

The proof of Proposition 5 builds on the idea of the proof of Proposition 1. First, we identify an upper bound on the principal's payoff guarantee. Second, we show that the strategy specified in the statement of the proposition achieves this upper bound.

Step 1. Identifying an upper bound for the principal's payoff guarantee. Consider the possible scenario where the only available projects are the ones known for each agent. In this scenario, the principal's payoff cannot exceed the highest possible social surplus, which can be computed by analyzing a social planner's optimal search strategy when facing m projects $\{a_0^k\}_{k=1}^m$. Following Weitzman (1979), the social planner optimally explores the projects in descending order of their indexes and stops when the highest up-to-date prize is higher than all the unexplored projects' indexes. We denote by V_S^* the social planner's optimal expected payoff. Since $\left\{a_0^k\right\}_{k=1}^m$ is one possible choice by nature, V_S^* is an upper bound for the principal's payoff guarantee.

Step 2. The startegy in Proposition 5 achieves V_S^*. To prove that the proposed strategy is optimal, it suffices to prove:

1. Under the proposed strategy, the worst-case scenario is $\mathcal{A}_0 = \{a_0^k\}_{k=1}^m$.

2. The proposed strategy guarantees the principal a payoff of V_S^* under $\mathcal{A}_0$.

First, we show that, under the proposed strategy, the principal's worst-case scenario is exactly the aforementioned one where each agent only has access to his known project, $\mathcal{A}_0 = \{a_0^k\}_{k=1}^m$. The proof of this argument is identical to that of Proposition 1 — under a debt contract, having more projects induces the agent to present a weakly better prize in the first-order stochastic dominance sense, which increases the principal's payoff.

Second, we show that, in the principal's worst-case scenario, the strategy specified in Proposition 5 achieves V_S^*, and, hence, is robustly optimal. To see this, observe first that it induces the socially efficient outcome, as the order of search among the projects $\{a_0^k\}_{k=1}^m$ and the stopping rule under this strategy is identical to that of the social planner's optimal strategy. Moreover, this strategy leaves no surplus for the agents, as they each face a debt contract that satisfies the FSE condition. Hence, the principal's expected payoff in this scenario is the entire social surplus.

B Efficiency Concerns

If the principal only evaluates contracts by their worst-case performance, a natural criterion to tie-break among them is on the basis of total surplus. This section investigates whether any optimal contract also maximizes total surplus. Given that there may be many true sets of projects $\mathcal{A}$, we follow a strong notion of efficiency. We denote the agent's expected payoff given a contract w and a set of projects $\mathcal{A}$ as $V_A(w \mid \mathcal{A})$.

Definition 1. *Let $V_S(w \mid \mathcal{A}) := \sup_{\sigma\in\Sigma(w,\mathcal{A})}[V_P(w|\mathcal{A}) + V_A(w|\mathcal{A})]$ denote the total surplus induced by the contract w when the realized set of projects is $\mathcal{A}$. The contract w is said to be **efficient** if it maximizes $V_S(w \mid \mathcal{A})$ for any $\mathcal{A}$.*

Despite the demanding criterion, efficiency refinement works for some parameters of a_0, as shown by the following proposition. Let $\underline{y}_0 := \inf(\text{supp}(F_0))$ denote the lowest possible prize that may be generated by the project a_0.

Proposition 6. *If $c_0 \geq E_{F_0}[y] - \underline{y}_0$, the pure debt contract w_0 is the only optimal contract that is efficient.*[23] *If $c_0 < E_{F_0}[y] - \underline{y}_0$, no optimal contract is efficient.*

Proof. See Appendix B.1. □

This result is driven by an important property that pure debt contracts are *order-preserving*, i.e., the agent's search order coincides with the social planner's, despite the presence of the contract. Direct computation shows that the w_0-induced index of *any* project a_i is $r_i^{w_0} = r_i - r_0$. Because a pure debt contract uniformly reduces the indices of all projects, it does not distort the order in which the agent searches them, nor when the agent terminates search with a positive wage.

[23] Unlike the argument in Section 4, limited liability is important for uniqueness in this theorem; otherwise, selling the firm to the agent would also be robustly optimal and efficient.

The only possibility that a debt contract generates inefficient search is when it discourages the agent from searching some project a_1 after the known a_0 has been explored. Without the w_0 contract in place, it may have been efficient to explore a_1 because $r_1 \geq 0$, but $r_1^{w_0} < 0$. This possibility is excluded by the assumption on a_0, which implies that even the social planner will never find it optimal to explore projects with a lower index than a_0. We thus conclude that the agent's w_0-induced search order among any $\mathcal{A}$ is the same as that of the social planner.

Proposition 6 relies on the condition $c_0 \geq E_{F_0}[y] - \underline{y}_0$, which holds in situations where either (i) the cost of exploring the known project is large or (ii) the "downside risk" of a_0, measured by $E_{F_0}[y] - \underline{y}_0$, is small. This condition holds, for instance, in the special case where a_0 is a riskless project whose potential is well-known at the time of contracting.

B.1 Proof of Proposition 6

Part 1: If $c_0 \geq E_{F_0}[y] - \underline{y}_0$, w_0 is efficient. The condition $c_0 \geq E_{F_0}[y] - \underline{y}_0$ is equivalent to $\underline{y}_0 \geq r_0$, which ensures that the social planner's search process will not continue after sampling the project a_0. This is also true for the agent when he faces the debt contract w_0. Hence, as we mentioned in the explanation following Proposition 6, it suffices to show that under w_0, the agent's search order among the projects whose indexes are weakly higher than r_0 is identical to that of the social planner's. Intuitively, this is because the agent is the full residual claimant under a debt contract. To see this formally, notice that $r_i^{w_0} = r_i - r_0$ if $r_i \geq r_0$, which is true because $[y_i - r_i]^+ = [(y_i - r_0)^+ - (r_i - r_0)]^+$ when $r_i \geq r_0$, making $\int [y_i - r_i]^+ dF_i(y_i) = \int [(y_i - r_0)^+ - (r_i - r_0)]^+ dF_i(y_i)$.

Part 2: If $c_0 < E_{F_0}[y] - \underline{y}_0$, w_0 is not efficient. The condition $c_0 < E_{F_0}[y] - \underline{y}_0$ is equivalent to $\underline{y}_0 < r_0$. Let $a_1 = (\delta_{r_0 - \epsilon}, 0)$ and $\mathcal{A} = \{a_0, a_1\}$, where ϵ is positive but arbitrarily small. In this case, the social planner will first sample project a_0 and continue to sample a_1 if $y < r_0$, whereas the agent will only sample a_0 and stop.

Part 3: No robustly optimal contract $w \neq w_0$ is efficient. To ensure that the agent's search order among projects with indexes higher than r_0 is the same as the social planner's for any realization of $\mathcal{A}$, we need $w(y') - w(y'') = y' - y''$ for any $y' > y'' \geq r_0$, i.e., the agent needs to be the full residual claimant. If this condition is violated, we can construct the following counterexamples.

Suppose $w(y') - w(y'') > y' - y''$. Let $a_1 = (\delta_{y'}, y' - y'' + \epsilon)$, $a_2 = (\delta_{y''}, 0)$, and $\mathcal{A} = \{a_0, a_1, a_2\}$, where ϵ is positive but arbitrarily small. In this case, the social planner will first sample project a_2, whereas the agent will first sample a_1.

Suppose $w(y') - w(y'') < y' - y''$. Let $a_1 = (\delta_{y'}, y' - y'')$, $a_2 = (\delta_{y''}, \epsilon)$, and $\mathcal{A} = \{a_0, a_1, a_2\}$, where ϵ is positive but arbitrarily small. In this case, the social planner will first sample project a_1, whereas the agent will first sample a_2.

Hence, an efficient contract must satisfy $w(y') - w(y'') = y' - y''$ for any $y' > y'' \geq r_0$. Combining this condition, the FSE condition specified in Theorem 1, and the limited liability assumption, we conclude that an efficient and optimal contract must satisfy $w(y) = 0$ if $y \leq r_0$.

References

Aghion, P., & Bolton, P. (1992). An incomplete contracts approach to financial contracting. *The Review of Economic Studies*, *59*(3), 473–494.

Antic, N. (2021). Contracting with unknown technologies. *Working paper*.

Armstrong, M., & Vickers, J. (2010). A model of delegated project choice. *Econometrica*, *78*(1), 213–244.

Ball, I., & Kattwinkel, D. (2024). Robust robustness. *arXiv preprint arXiv:2408.16898*.

Ben-Porath, E., Dekel, E., & Lipman, B. L. (2025). Mechanism design for acquisition of stochastic evidence. *Theoretical Economics*.

Bergemann, D., & Hege, U. (1998). Venture capital financing, moral hazard, and learning. *Journal of Banking & Finance*, *22*(6-8), 703–735.

Bergemann, D., & Hege, U. (2005). The financing of innovation: Learning and stopping. *RAND Journal of Economics*, 719–752.

Carroll, G. (2015). Robustness and linear contracts. *American Economic Review*, *105*(2), 536–63.

Carroll, G. (2019). Robustness in mechanism design and contracting. *Annual Review of Economics*, *11*(1), 139–166.

Chassang, S. (2013). Calibrated incentive contracts. *Econometrica*, *81*(5), 1935–1971.

Clews, R. (2016). *Project finance for the international petroleum industry*. Amsterdam : Boston: Elsevier.

Dai, T., & Toikka, J. (2022). Robust incentives for teams. *Econometrica*, *90*(4), 1583–1613.

Dang, T. V., Gorton, G., & Holmström, B. (2011). Ignorance and the optimality of debt for liquidity provision. *Working paper*.

DeMarzo, P., & Duffie, D. (1999). A liquidity-based model of security design. *Economet-*

rica, *67*(1), 65–99.

Durandard, T. (2023). Dynamic delegation in promotion contests. *arXiv:2308.05668*.

Durandard, T., & Ghersengorin, A. (2024). Robust regulation of labour contracts. *arXiv preprint arXiv:2411.04841*.

Dworczak, P., & Pavan, A. (2022). Preparing for the worst but hoping for the best: Robust (bayesian) persuasion. *Econometrica*, *90*(5), 2017–2051.

Guo, Y. (2016). Dynamic Delegation of Experimentation. *American Economic Review*, *106*(8), 1969–2008.

Guo, Y., & Shmaya, E. (2023). Regret-minimizing project choice. *Econometrica*, *91*(5), 1567–1593.

Halac, M., Kartik, N., & Liu, Q. (2016). Optimal contracts for experimentation. *The Review of Economic Studies*, *83*(3), 1040–1091.

Hart, O., & Moore, J. (1994). A theory of debt based on the inalienability of human capital. *The Quarterly Journal of Economics*, *109*(4), 841–879.

Hébert, B. (2018). Moral hazard and the optimality of debt. *The Review of Economic Studies*, *85*(4), 2214–2252.

Hörner, J., & Samuelson, L. (2013). Incentives for experimenting agents. *The RAND Journal of Economics*, *44*(4), 632–663.

Hurwicz, L., & Shapiro, L. (1978). Incentive structures maximizing residual gain under incomplete information. *The Bell Journal of Economics*, 180–191.

Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory*, *52*(1), 45–67.

Jensen, M., & Meckling, W. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, *3*(4), 305–360.

Kambhampati, A. (2023). Randomization is optimal in the robust principal-agent problem. *Journal of Economic Theory*, *207*, 105585.

Kambhampati, A. (2024). Proper Robustness and the Efficiency of monopoly Screening. *Working Paper*.

Kambhampati, A., Peng, B., Tang, Z., Toikka, J., & Vohra, R. (2025). Randomization and the Robustness of Linear Contracts. *Working Paper*.

Kleinberg, R., Waggoner, B., & Weyl, E. G. (2016). Descending price optimally coordinates search. *arXiv:1603.07682*.

Koh, A., & Sanguanmoo, S. (2024). Robust technology regulation. *arXiv:2408.17398*.

Lewis, T. R. (2012). A theory of delegated search for the best alternative. *The RAND Journal of Economics*, *43*(3), 391–416.

Lewis, T. R., & Ottaviani, M. (2008). Search agency. *Mimeo*.

Libgober, J., & Mu, X. (2021). Informational robustness in intertemporal pricing. *The Review of Economic Studies*, *88*(3), 1224–1252.

Lin, L. (2020). Contractual innovation in china's venture capital market. *European Business Organization Law Review*, *21*, 101-138.

Liu, C. (2022). Robust contracts with exploration. *arXiv:2212.00157*.

Miller, M. H. (1977). Debt and taxes. *Journal of Finance*, *32*(2), 261–275.

Myers, S., & Majluf, N. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, *13*, 187–221.

Nachman, D. C., & Noe, T. H. (1994). Optimal design of securities under asymmetric information. *The Review of Financial Studies*, *7*(1), 1–44.

Nocke, V., & Whinston, M. D. (2013). Merger policy with merger choice. *American Economic Review*, *103*(2), 1006–1033.

Olszewski, W. (2025). On evaluating contracts by their performance in the worst cases. *Working paper*.

Tirole, J. (2010). *The theory of corporate finance.* Princeton university press.

Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, *21*(2), 265–293.

Ulbricht, R. (2016). Optimal delegated search with adverse selection and moral hazard. *Theoretical Economics*, *11*(1), 253–278.

Walton, D., & Carroll, G. (2022). A general framework for robust contracting models. *Econometrica*, *90*(5), 2129–2159.

Weitzman, M. L. (1979). Optimal Search for the Best Alternative. *Econometrica*, *47*(3), 641.

Xue, Y., & Yun, X. (2022). The optimal performance target of valuation adjustment mechanism agreement with real options perspective. *PLOS ONE*, *17*(11), e0277509.

Yang, M. (2020). Optimality of debt under flexible information acquisition. *The Review of Economic Studies*, *87*(1), 487–536.