Robust Contracting for Search*

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

October 13, 2024

Abstract

A principal contracts with an agent who can sequentially search over some projects to generate an award. Crucially, the principal knows only one of the available projects and evaluates a contract by its worst-case performance. We find that a debt contract is robustly optimal. When the principal's known project faces limited downside risk, this debt contract is also the unique contract that is robustly optimal and socially efficient. Moreover, linear contracts, although robustly optimal in various settings where the principal has partial knowledge about the agent's technology, are sub-optimal in our setting as they deter the agent's search incentive.

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

JEL Codes: D82, D83, D86.

*We thank Dan Barron, Dan Bernhardt, Matteo Camboni, Alexis Ghersengorin, Mallesh Pai, Alessandro Pavan, Mike Powell, Luis Rayo, Pëllumb Reshidi, Bruno Strulovici, João Thereze, and various audiences at the 35th Stony Brook International Conference on Game Theory, Duke, UNC, and Northwestern. All errors are our own.

[†]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

In many situations, a principal ("she") contracts with an agent ("he") who can sequentially search over some projects to generate a reward, whereas the principal knows only one of the available projects. Consider the example of a publisher sponsoring the creative work of a writer (or an artist, an online content creator, etc). The publisher may be aware of the initial "pitch," but the exact creative process or alternative ideas are only known to the writer. With this in mind, the publisher writes a contract that specifies how any future profits will be divided between the two.

How should the publisher (the principal) write a contract to encourage the writer (the agent) to search through projects while being robust to unknown alternatives that the agent may possess? This paper delivers two main findings. First, we construct a debt contract that maximizes the principal's worst-case guarantee. Moreover, when the principal's known project faces limited downside risk, this debt contract is the sole optimal and socially efficient contract. This finding provides a rationale for the prevalent use of advance-against-royalties contracts commonly adopted by publishers. Second, we find that linear contracts, although known to be robustly optimal in various settings where the principal has partial knowledge about the agent's technology (e.g., Carroll, 2015, Dai & Toikka, 2022, Liu, 2022), fail to provide the principal with her optimal guarantee in our setting. This is because linear contracts deter the agent's search incentive.

We study a moral hazard model combining robust contracting and sequential search. In the model, a principal designs a wage contract for an agent, who can sequentially search over some projects *à la* Weitzman (1979) and present a final award to the principal. Each project is a "Pandora's box" in the language of Weitzman (1979) — it specifies a distribution over the project's realized award and a fixed cost to make that award collectible. The contracting environment features two frictions. The first one is *limited* *liability*, i.e., the wage must be non-negative. The second one is *asymmetric information* in the sense that the principal only knows one of the projects that the agent can access. Crucially, we adopt a robustness approach to deal with the second friction — the principal evaluates a contract by its worst-case performance against the uncertainty about the projects endowed to the agent.

Theorem 1 constructs a debt contract that is worst-case optimal for the principal. In this contract, the principal collects the entire award if it is less than a certain amount, whereas the agent is the residual claimant after the award exceeds that amount. Intuitively, the optimality of this debt contract hinges on the following two facts. First, it preserves the agent's incentive to search by letting him be the residual claimant for high realizations of the award. Second, it ensures that the principal can extract the entire social surplus in the worst-case scenario under this contract.

Naturally, the next question concerns whether this debt contract is uniquely optimal. Although the answer is no, Theorem 2 fully characterizes the set of worst-case optimal contracts for the principal. A contract is optimal if and only if it satisfies two conditions. The first one (Minimum Debt Level condition) says that a minimum level of debt must be imposed, i.e., the principal collects the entire award up to some certain amount. This feature is crucial to preserve the agent's search incentive. The second one (Full Surplus Extraction condition) says that the principal must be able to fully extract the surplus in the worst-case scenario.

Given that the debt contract constructed in Theorem 1 is not uniquely optimal, what makes it special? Theorem 3 answers this question by identifying a unique feature of this debt contract — among all the robustly optimal contracts, it is the only one that achieves social efficiency when the known project's downside risk is not severe. Intuitively, this result is driven by the fact that debt contracts are *order-preserving* in the sense that they induce the agent to search in an order identical to a social planner, regardless of what projects are available to him.

The above results combined provide a rationale for the use of debt-like contracts in practice. In the publishing industry, for instance, it is common for writers to be paid an advance against any future profits their book would earn.[1] Like the debt contract featured in this paper, the writer starts earning royalty payments only after the book sale exceeds an agreed threshold. This kind of "earnout" contract encourages the writer to search for a better publication, as it gives the writer a high residual claim only when the book sells well.

Another main contribution of this paper is to demonstrate how the timing nature of the agent's decision affects the feature of the robustly optimal contract. From a modeling perspective, our only departure from Carroll (2015) is that the agent can sequentially search over the available projects to find his favorite instead of picking his favorite project as a static decision. It turns out that linear contracts, which are optimal in Carroll (2015), are sub-optimal in our setting (Corollary 1). We provide two angles to explain why our robustly optimal contract differs from Carroll (2015).

First, the agent cares about different statistics of a project. In Carroll (2015), the agent cares about the expected award of a project, and because of that, linear contracts are optimal since they well align the principal's and the agent's payoffs. In our setting, the right tail of the award distribution is what guides the agent's search, as indicated by the computation of the Weitzman index. Therefore, preserving the agent's incentives to search requires not significantly distorting the relative values of the high award realizations. This is achieved by debt contracts but not by linear contracts (or equivalently, equity contracts).

Second, in a one-shot moral hazard problem *à la* Carroll (2015), the principal is primarily concerned about excessively risky projects crowding out others with comparable expected values. In a sequential search problem, these risky projects are a boon to the

[1] See, e.g., https://www.booksandsuch.com/blog/sell-in-sell-through/ for the use of contracts in the publishing industry.

principal; instead, she must guard against excessively safe projects that would terminate the agent's search too early.

Related Literature

First and foremost, this paper contributes to the literature on robust contracting, particularly a strand of the literature that concerns robustness to the agent's technology (Hurwicz & Shapiro, 1978; Carroll, 2015). The closest paper to ours is Carroll (2015), with which we have intensively compared in the introduction.[2] Our paper demonstrates how the timing nature of the agent's decision affects the feature of robustly optimal contracts.[3] See Carroll (2019) for a survey on the literature.

Second, this paper contributes to the literature on delegated search (or contracting for search), such as Lewis and Ottaviani (2008), Lewis (2012), and Ulbricht (2016).[4] A common feature of those papers is that the agent can exert costly effort to repeatedly sample awards from the same distribution — to put it in our context, the agent faces infinitely many *homogeneous* Pandora's boxes. Our major departure from the literature is to consider independently *heterogeneous* Pandora's boxes.

Third, this paper is related to the literature on contracting for experimentation (Bergemann & Hege, 1998, 2005; Hörner & Samuelson, 2013; Halac, Kartik, & Liu, 2016). The main difference lies in the format of moral hazard — the agent in those papers pays effort to learn about a state of the world, while the agent in ours pays effort to learn about a

[2]Relatedly, followup works of Carroll (2015) include Dai and Toikka (2022) that extends the model to a team of agents, Liu (2022) that allows the principal to learn about the agent's technology, Kambhampati (2023) that analyzes the benefit of random contracts, etc.

[3]Chassang (2013) also studies robust contracting where the agent makes dynamic decisions. One crucial difference is that the principal in our paper only observes the final outcome (i.e., the presented award), whereas the principal in Chassang (2013) can observe the outcomes across time. Another related paper is Koh and Sanguanmoo (2024), where the agent can learn about his technology over time.

[4]Relatedly, some papers study sequential contracting with multiple agents (Kleinberg, Waggoner, & Weyl, 2016; Durandard, 2023; Ben-Porath, Dekel, & Lipman, 2021), which can be cast as sequential search problems with the principal being the searcher and agents being the boxes.

project's award and make it collectible.

Fourth, this paper is related to the literature on delegated project choice (Armstrong & Vickers, 2010; Nocke & Whinston, 2013; Guo & Shmaya, 2023). The crucial differences are that (1) we allow for transfers, and (2) the agent in our setting needs to exert costly effort to discover the award associated with each project.

Finally, this paper is related to the finance-oriented literature on the optimality of debt contracts. Some previous papers have emphasized the roles of debt contracts in mitigating the agent's moral hazard problems in the presence of limited liability (Jensen & Meckling, 1976; Townsend, 1979; Innes, 1990; Hébert, 2018; etc).[5] A common intuition behind these papers is that a debt contract is closest, among those that satisfy limited liability, to the "contract" that lets the principal sell the firm to the agent, which is optimal absent limited liability. Our paper is aligned with this literature as we delineate how debt contracts can mitigate a specific moral hazard problem where the agent can sequentially search.

2 Model

A principal ("she") contracts with an agent ("he"), who can sequentially search through projects to generate an award. The set of projects available to the agent is denoted by $\mathcal{A} = \{a_i\}_{i=0}^{n}$, with n being a finite number. Each project a_i is described by a pair $(F_i, c_i) \in \Delta(Y) \times \mathbb{R}^+$ where $Y := [0, \bar{y}]$. The interpretation is that the agent can exert a *cost* of c_i to learn project i's realized *award* y_i, which is a priori distributed according to F_i. We denote $\underline{y}_i := \inf y_i$ and $\bar{y}_i := \sup y_i$ according to F_i. Notice that a project is equivalent to a "Pandora's box" in the language of Weitzman (1979). Crucially, while the

[5]Besides mitigating moral hazard, debts are also justified for other reasons such as lowering firms' tax burden (Miller, 1977), signaling firms' positive private information about profitability (Myers & Majluf, 1984), etc. See Tirole (2010) for an overview of the literature on security design and capital structure.

agent knows $\mathcal{A}$, the principal only knows one of those projects, a_0 — putting this formally, the principal's only knowledge about the available projects is that $\mathcal{A} \supseteq \mathcal{A}_0 := \{a_0\}$. Also, we make a non-triviality assumption that $\mathbf{E}_{F_0}[y_0] - c_0 > 0$.

After the search process, the agent presents a single award y to the principal.[6] His presented award is set at zero if he did not sample any project. The principal observes nothing about the agent's search process except the presented award; in particular, the principal does not know the identity of the project that generates this award.[7] Therefore, her only incentive tool is a wage contract $w : Y \to \mathbb{R}^+$, where $w(y)$ is the agent's monetary payment when the presented award is y. We assume the contract w satisfies one-sided limited liability, i.e., $w(y) \geq 0$ for any $y \in Y$.[8] Both players are risk-neutral.

Given a wage contract w and knowing the set of available 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 collectible monetary award from the sampled projects, denoted by $\tilde{y}$. Formally, a strategy is a function $\sigma : 2^{\mathcal{A}} \times Y \to 2^{\mathcal{A}} \cup \{\emptyset\}$, 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 the award that gives himself the monetary award $\tilde{y}$. Let $\Sigma(w, \mathcal{A})$ denote the set of optimal search strategies for the agent. For a given strategy σ, we write $\mathbf{E}_\sigma$ to denote the expectation with respect to the induced distribution over the agent's search. We abusively write $a_i \in \sigma$ to denote the event that the project a_i is sampled according to the strategy σ.

The game proceeds as follows.

[6]We assume the agent can only present one single award, as it is common in real-world contexts including the publisher-writer example. That being said, allowing the agent to present more than one award does not change the paper's main finding, as to be discussed in Section 4.3.

[7]In Section 4.4, we discuss the situation with project-specific contracts and show that the paper's main results continue to hold.

[8]The paper's results remain true if we assume two-sided limited liability. As we will show, the optimal contract (a debt contract) yields a non-negative payoff to the principal.

1. The principal sets a contract w.

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

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

The principal's objective is to determine a wage contract w that maximizes the worst-case payoff against all possible realizations of $\mathcal{A}$. Formally, given a realized set of projects $\mathcal{A}$ and a wage contract w, the principal's payoff is

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

where we assume the following: in case the agent has multiple optimal search strategies, he adopts the one that is most preferred by the principal. The principal evaluates a wage contract by its payoff guarantee (i.e., her payoff in the worst-case realization of $\mathcal{A}$),

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

The principal seeks to solve the following problem

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

where we let V_P denote her payoff guarantee under the optimal contract.

3 Main Results

We begin by recalling the solution to the sequential search problem (Weitzman, 1979). For the project $a_i = (F_i, c_i)$, we define its *reservation value* (or *index*), r_i, as the unique solution to

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

A social planner (one 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 reservation values and (2) conclude the search whenever a realized award y is larger than the reservation values of the remaining unsampled projects.

The agent's optimal search strategy is similar except that his incentive is potentially distorted by the wage contract. Given a wage contract w, project a_i's *w-induced reservation value* (or *w-induced index*), r_i^w, is the unique solution to

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

He will sample the projects in descending order of r_i^w and cease searching when his realized monetary payoff exceeds the w-induced reservation values of the remaining unsampled projects.

Having specified the agent's optimal search strategy, we return to the principal's problem. Define a contract w as a *$z-$debt contract* if $w(y) = [y - z]^+$. In such a contract, the principal collects all the returns up to the debt level z, after which the agent is the residual claimant and collects the rest. We are now ready to state the main result. Let w_0 denote the r_0-debt contract where r_0 is the index of the project a_0.

Theorem 1. *The contract w_0 is optimal for the principal. Moreover, $V_P = \mathbf{E}_{F_0}[y_0] - c_0$.*

Proof. See Appendix A.1. □

The optimality of w_0 is best explained by the following: for any contract w, we have

$$V_P(w_0) = V_P(w_0|\mathcal{A}_0) \geq V_P(w|\mathcal{A}_0) \geq V_P(w). \tag{1}$$

The equality of (1) says that $\mathcal{A} = \mathcal{A}_0$ is the worst-case scenario for the principal under the debt contract w_0, as will be shown in Appendix A.1. The first inequality of (1) holds because when $\mathcal{A} = \mathcal{A}_0$, the contract w_0 attains the highest social surplus while leaving zero payoff to the agent. The second inequality of (1) comes directly from the definition of $V_P(w)$.

Theorem 1 leads to a natural question: is w_0 uniquely worst-case optimal for the principal? Unfortunately, not quite, as the following theorem shows. Recall that V_P is the principal's payoff guarantee under the optimal contract.

Theorem 2. *A contract w is optimal if and only if the following two conditions hold.*

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

$$w(y) \begin{cases} = 0 & \text{if } y \leq V_P \\ \leq y - V_P & \text{if } y > V_P. \end{cases}$$

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

$$\mathbf{E}_{F_0}[w(y_0)] = c_0.$$

Proof. See Appendix A.2. □

Theorem 2 identifies the key properties of robustly optimal contracts. The MDL condition indicates that any robustly optimal contract w must impose a minimum debt level of V_P — specifically, the principal collects the entire return if the presented award is below V_P, whereas her payoff $y - w(y)$ is at least V_P if the presented award exceeds V_P. In other words, some level of debt (although not necessarily as high as r_0) is inevitable in any optimal contract. The FSE condition ensures that the principal fully extracts the social surplus when $\mathcal{A} = \mathcal{A}_0$.

Although the debt contract w_0 is not uniquely optimal for the principal, it strictly outperforms any other optimal contract in terms of efficiency in some situations, as the next theorem shows.

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

Theorem 3. *If the project* a_0*'s downside risk is bounded by* r_0 *(i.e.,* $\underline{y}_0 \geq r_0$*), the contract* w_0 *is the only optimal contract that is efficient.*

Proof. See Appendix A.3. $\square$

As will be clear in the proof, Theorem 3 is driven by the fact that debt contracts are the only contracts that are *order-preserving*, i.e., the agent's search order among any set of projects under such a contract is the same as the social planner's. The condition $\underline{y}_0 \geq r_0$ is satisfied as long as the known project a_0 is not exposed to severe downside risk; for example, it holds in the special case where a_0 is a riskless project (i.e., $F_0 = \delta_{c_0+r_0}$). Under this condition, the social planner would not have sampled any project whose index is lower than r_0. Hence, the social planner will sample the projects in descending order of their indexes until the project a_0, and it is not difficult to see that the agent will do the same under the debt contract w_0 because it is order-preserving.

It is also worth noticing that the limited liability assumption plays an important role in Theorem 3. Without this assumption, the principal can simply charge the agent a fixed fee and let her keep the entire award — a contract that is socially efficient and robustly optimal for the principal but violates limited liability.

3.1 Sub-optimality of Linear Contracts

Our major departure from Carroll (2015) is that we allow the agent to select his favorite project dynamically (via sequential search) instead of statically. To highlight how such dynamics of moral hazard bring a difference to the optimal contract, it is useful to study how the linear contracts, known to be optimal in Carroll (2015), perform in our setting.

Corollary 1. *No linear contract is optimal for the principal.*

Proof. All the linear contracts violate the MDL condition in Theorem 2, i.e., they do not contain a debt level of at least V_P. $\square$

To see why the violation of MDL makes a linear contract sub-optimal, we consider the linear contract $w_{lin}(y) = \alpha y$ with $\alpha = \frac{c_0}{\mathbf{E}_{F_0}[y_0]}$ so that the FSE condition is satisfied. The w_{lin}-induced index for a_0 solves

$$c_0 = \int [w_{lin}(y_0) - r_0^{w_{lin}}]^+ dF_0(y_0).$$

Plugging in the expression for α, we see that $r_0^{w_{lin}} = 0$. Now consider the alternative project $(\delta_x, 0)$, where δ_x is a Dirac mass on some $x > 0$. This project has a strictly positive index *and* $\alpha x > 0$. Therefore, under w_{lin}, the agent will search $(\delta_x, 0)$ and never proceed to a_0. Since this is true regardless of x, we get $V_P(w_{lin}) = \inf_{x>0}(1-\alpha)x = 0$.

We provide two angles to explain why our optimal contract differs from Carroll (2015). First, the agent cares about different statistics of a project. In Carroll (2015), the agent cares about the expected award of a project, and because of that, linear contracts are optimal since they well align the principal's and the agent's payoffs. In our setting, the right tail of the award distribution is what guides the agent's search, as indicated by the computation of the Weitzman index. Therefore, preserving the agent's incentives to search requires not significantly distorting the relative values of the high award realizations.

This is achieved by debt contracts, but not by linear contracts (or equivalently, equity contracts).

Second, in a one-shot moral hazard problem *à la* Carroll (2015), the principal is primarily concerned about excessively risky projects crowding out others with comparable expected values. In a sequential search problem, these risky projects are a boon to the principal; instead, she must guard against excessively safe projects that would terminate the agent's search too early.

4 Discussion

4.1 Multiple Agents

Consider a variation of the model with m agents indexed by $k = 1, 2, ..., m$. Each agent k has access to a set of projects denoted by $\mathcal{A}^k$, whereas the principal only knows one project available to each agent, which we denote by $a_0^k \in \mathcal{A}^k$. The principal can sponsor any agent at any time. Upon being sponsored, an agent searches over his available projects and presents one single award to the principal. The principal can only *adopt* one presented award, regardless of how many awards are presented (e.g., the publisher can only publish one book among those presented by several writers). The principal's realized payoff is the utility of the adopted award minus the wage payment to all the sponsored agents.

This setting features two layers of sequential search — the principal can explore among agents, whereas the agents can explore among 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. In principle, the wage payment to an agent can depend on all the presented awards, including those presented by others. But as we will show, such richness does not benefit the principal — she can already achieve her optimal payoff guarantee by providing each sponsored agent a simple wage contract that only depends on his own presented award (notably, regardless of whether that award is adopted or not).

Let r_0^k denote the index of agent k's project a_0^k. Without loss of generality, we assume $r_0^1 \geq r_0^2 \geq ... \geq r_0^m$.

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

Proof. We begin by identifying an upper bound for the principal's payoff guarantee. Consider the possible scenario where there are no available projects other than the known one 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 $\{a_0^k\}_{k=1}^m$. As in Weitzman (1979), the social planner will explore the projects in descending order of their indexes and stop when the highest up-to-date award is higher than all the unexplored projects' indexes. By doing so, the social planner's expected payoff is denoted by V_S^*. This is an upper bound for the principal's payoff guarantee.

Next, we show that the strategy specified in Proposition 1 achieves a payoff guarantee that attains the above upper bound. It suffices to prove the following two arguments. The first argument resembles the equality in (1), and the second one resembles the first inequality in (1).

First, under this strategy, the principal's worst-case scenario is exactly the aforementioned one where each agent only has access to his known project. The proof of this argument is the same as that in Theorem 1 — under a debt contract, having more projects induces the agent to present a weakly better award in the first-order stochastic dominance sense.

Second, in the principal's worst-case scenario, the strategy specified in Proposition 1 achieves 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. Meanwhile, this strategy leaves no surplus for all the sponsored agents since they each face a debt contract that satisfies the FSE condition. Hence, the principal's expected payoff in this scenario is also V_S^*, the entire social surplus. □

4.2 Screening

This paper features not only moral hazard but also adverse selection. Yet, the baseline model assumes that the principal offers a single contract w and does not screen the agent's private information about $\mathcal{A}$. Can the principal benefit from screening the agent by providing a more general mechanism? We find that the answer is no.

By the taxation principle, any mechanism is equivalent to a menu of contracts.

Proposition 2. *The singleton menu $\{w_0\}$ achieves the principal's optimal payoff guarantee among all mechanisms.*

Proof. Slightly abusing the notation, for any menu of contracts $\mathcal{M}$, we have

$$V_P(\{w_0\}) = V_P\left(\{w_0\} \mid \mathcal{A}_0\right) \geq V_P\left(\mathcal{M} \mid \mathcal{A}_0\right) \geq V_P\left(\mathcal{M}\right).$$

The equality and the second inequality remain true as in Theorem 1. The first inequality holds for the following reason: when the agent's type is $\mathcal{A}_0$, the contract w_0 already maximizes the social surplus while leaving no surplus for the agent; in other words, no menu can improve the principal's payoff when $\mathcal{A} = \mathcal{A}_0$. □

4.3 Multiple Presented Awards

We consider a variation of the model where the agent can present more than one award. We maintain the assumption that the principal can only adopt one award. Let the set of presented awards be denoted by $\mathbf{y} = \{y^{(1)}, y^{(2)}, ..., y^{(k)}\}$ with $k \leq n+1$, and the principal's selected award be denoted by $\hat{y} \in \mathbf{y}$. The principal can use a general wage contract $w(\mathbf{y}, \hat{y})$ that conditions the agent's wage on both the number of presented awards and their values. Moreover, the principal will select $\hat{y} \in \mathbf{y}$ to maximize $\hat{y} - w(\mathbf{y}, \hat{y})$. Viewed in that light, our problem resembles the project choice problem studied by Guo and Shmaya (2023), with two key differences. First, we allow for transfers. Second, the agent needs to exert effort to "discover" the award associated with each project.

These two differences turn out to drastically change the structure of the optimal mechanism. Contrary to Guo and Shmaya (2023), the principal cannot gain from incentivizing the agent to disclose multiple awards. The contract w_0 remains robustly optimal for the principal, since (1) continues to hold. In particular, the equality of (1) still holds because it only concerns the specific contract w_0. The first inequality remains true because in the scenario where $\mathcal{A} = \mathcal{A}_0$, the generalization of the contract space has no bites. The second inequality still holds by definition.

Why doesn't the principal benefit from incentivizing the agent to propose multiple projects in our setting? Intuitively, this is because the transfer allows the principal to fully separate incentive provision from project selection. Unlike in Guo and Shmaya (2023), where the principal incentivizes the disclosure of multiple projects to hedge strategic risks and improve her fallback option, in our problem, the principal can always pick her favorite award. The debt contract w_0 not only motivates search but also aligns the agent's and principal's interests in selecting the best project. As a result, the principal does not need to compensate the agent to disclose the best available project, which is then implemented.

Moreover, linear contracts as defined in Carroll (2015) remain sub-optimal in this setting, where the contract space is enlarged.

4.4 Project-Specific Contracts

Another way to generalize the contract space is to let the wage depend on the identity of the project that generates the presented award. Such a contract can be written as $w(y_i, a_i)$.

Similarly, this generalization does not affect the optimality of the debt contract w_0 since (1) remains true for the same reason as Section 4.3. Intuitively, $\mathcal{A} = \mathcal{A}_0$ is established as the worst-case scenario when the contract w_0 is used, while this generalization of the contract space is irrelevant under this worst-case scenario.

As before, linear contracts remain sub-optimal in this generalized setting.

4.5 Resampling of Projects

In our baseline model, the agent can only sample each project once. This paper's results remain valid if we allow for finite resampling, i.e., each project can be sampled for finite times. This is because multiple samplings of the same project can be regarded as one-time sampling of multiple projects with identical (F_i, c_i).

This is also true if we allow for infinite resampling. In particular, w_0 remains optimal because $\mathcal{A} = \mathcal{A}_0$ is still the worst-case scenario for the principal under the debt contract w_0. In particular, under a debt contract under which the agent is willing to sample the project a_0, the agent always samples until getting some positive payoff, which means that the principal has already collected r_0 in full.

A Proofs

A.1 Proof of Theorem 1

As we mentioned in the explanation following Theorem 1, it suffices to show the following two statements. First, $V_P(w_0|\mathcal{A}_0) \geq V_P(w|\mathcal{A}_0)$. This is because

$$\mathbf{E}_{F_0}[w_0(y_0)] - c_0 = \mathbf{E}_{F_0}[(y_0 - r_0)^+] - c_0 = 0.$$

In other words, when a_0 is the only available project, the contract w_0 maximizes social surplus (because the agent is still willing to sample the project a_0) while leaving zero expected payoff to the agent.

Second, $V_P(w_0) = V_P(w_0|\mathcal{A}_0)$; that is, when using the contract w_0, the principal's worst-case scenario is $\mathcal{A} = \mathcal{A}_0$. To show this, fix the realized awards $\{y_i\}_{i=0}^n$ and the agent's optimal search strategy. The realized search process must belong to one of the two following cases.

Suppose the project a_0 is *not sampled* by the agent under the realized search process. It must be that the agent has sampled another project a_i whose award y_i satisfies $w_0(y_i) > r_0^{w_0} = 0$. This implies that the agent has surpassed the debt level, so the principal surely collects a payoff of r_0, which is the best situation for the principal under the r_0-debt contract.

Suppose the project a_0 is *sampled* by the agent under the realized search process. Since the contract w_0 is weakly increasing, the agent always reports the maximal sampled award at the end of the process. The distribution of the maximal sampled award $\max\{w_0(y_i)\}_{i\in\sigma}$ under any $\mathcal{A} \supsetneqq \mathcal{A}_0$ must first-order stochastically dominate that under $\mathcal{A}_0$. Since $y - w(y)$ is weakly increasing, the principal's expected payoff attains the minimum under $\mathcal{A}_0$.

Having shown the optimality of w_0, it is immediate that the principal's payoff guarantee is the social surplus of the project a_0, i.e., $V_P = V_P(w_0|\mathcal{A}_0) = \mathbf{E}_{F_0}[y_0] - c_0$.

A.2 Proof of Theorem 2

Part 1: MDL & FSE $\rightarrow$ w is optimal. We directly show that $V_P(w) = V_P$ if the contract w satisfies MDL and FSE. Fix the realized award $\{y_i\}_{i=0}^n$ and the agent's optimal search strategy, and then we consider two cases of the realized search process.

Suppose a_0 is *not sampled*, then the agent must have sampled a project whose award is higher than V_P. This ensures that the principal will attain at least V_P according to the MDL condition.

Suppose a_0 is *sampled*, then it is also the last project sampled by the agent because a_0's w-induced reservation value is exactly zero according to the FSE condition. Hence, the principal's expected payoff in this case is V_P.

Part 2: w is optimal $\rightarrow$ MDL & FSE. We prove by contradiction. First, suppose a contract w violates FSE. We have the following:

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

where the equality is proved in Theorem 1, the first inequality holds strictly because FSE is violated, and the second inequality holds by the definition of $V_P(w)$. Hence, the contract w is not robustly optimal as $V_P(w) < V_P$.

Second, suppose a contract w satisfies FSE but violates MDL for some $y' \leq V_P$, i.e., $w(y') > 0$. We consider the counterexample where $\mathcal{A} = \{a_0, a_1\}$ with $a_1 = (\delta_{y'}, 0)$ being a riskless project. The agent will only search project a_1 because the w-induced reservation value of the project a_0 is zero due to FSE. The principal's corresponding payoff, $y' - w(y')$, is strictly smaller than y', which is weakly smaller than V_P. In other words, $V_P > V_P(w|\{a_0, a_1\}) \geq V_P(w)$.

Third, suppose a contract w satisfies FSE but violates MDL for some $y' > V_P$, i.e., $w(y') > y' - V_P$. We consider the counterexample where $\mathcal{A} = \{a_0, a_1\}$ with $a_1 = (\delta_{y'}, 0)$ being a riskless project. The agent will only search project a_1 because the w-induced reservation value of the project a_0 is zero due to FSE. The principal's corresponding payoff, $y' - w(y')$, is strictly smaller than V_P because MDL is violated. In other words, $V_P > V_P(w|\{a_0, a_1\}) \geq V_P(w)$.

A.3 Proof of Theorem 3

Part 1: w_0 is efficient. The condition $\underline{y}_0 \geq r_0$ ensures that the social planner's search process will not continue after sampling the project a_0, which is also true for the agent when he faces the debt contract w_0. Hence, as we mentioned in the explanation following Theorem 3, 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: Any other optimal contract is not efficient. First, 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 counter-examples.

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 2, and the limited liability assumption, we conclude that an efficient and optimal contract must satisfy $w(y) = 0$ if $y \leq r_0$.

References

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

Ben-Porath, E., Dekel, E., & Lipman, B. L. (2021). Mechanism design for acquisition of stochastic evidence.

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.

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

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

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.

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.

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, Northwestern University.*.

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

Miller, M. H. (1977). Debt and taxes. *the 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.

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

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.

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