

Regret-Minimizing Project Choice

Yingni Guo Eran Shmaya*

June 23, 2023

Abstract

An agent observes the set of available projects and proposes some, but not necessarily all, of them. A principal chooses one or none from the proposed set. We solve for a mechanism that minimizes the principal's worst-case regret. We compare the single-project environment in which the agent can propose only one project with the multiproject environment in which he can propose many. In both environments, if the agent proposes one project, it is chosen for sure if the principal's payoff is sufficiently high; otherwise, the probability that it is chosen decreases in the agent's payoff. In the multiproject environment, the agent's payoff from proposing multiple projects equals his maximal payoff from proposing each project alone. The multiproject environment outperforms the single-project one by providing better fallback options than rejection and by delivering this payoff to the agent more efficiently.

JEL: D81, D82, D86

Keywords: project choice, minimax regret, verifiable proposal, two-tier mechanism

*Guo: Department of Economics, Northwestern University; email: yingni.guo@northwestern.edu. Shmaya: Department of Economics, Stony Brook University; email: eran.shmaya@stonybrook.edu. We thank seminar audiences at the One World Mathematical Game Theory Seminar, the Toulouse School of Economics, the University of Bonn, University of Pennsylvania, Brown University, Northwestern University Theory Lunch Seminar, Kellogg Strategy Department Lunch Seminar, the University of Pittsburgh, Carnegie Mellon University, Stony Brook International Conference on Game Theory, Workshop on Strategic Communication and Learning, TE Session at ASSA, Arizona State University, Stanford University, Cornell University, Paris Dauphine University, MIT/Harvard, NYU, and Erice Workshop on Stochastic Methods in Game Theory for valuable feedback.

1 Introduction

Project choice is an important function of many organizations.¹ The process often involves two parties: (i) a party at a lower hierarchical level who has expertise and proposes projects, and (ii) a party at a higher hierarchical level who evaluates the proposed projects and makes the choice. This describes, for example, the relationship between a firm’s division and the firm’s headquarters when the division has a chance to choose an investment project, a factory location, or an office building.² It also applies to the relationship between a department and the university administration when the department has a hiring slot open.

This process of project choice is naturally a principal-agent problem. The agent privately observes which projects are available and proposes a subset of the available projects. The principal either chooses one from the projects proposed or rejects them all. If the two parties had identical preferences over projects, the agent would propose their jointly favorite project from the available ones, and the principal would simply rubber-stamp the agent’s proposal. In reality, however, the principal does exert power over project choice, because the agent’s preferences are not necessarily aligned with the principal’s. For example, division managers may fail to internalize projects’ externalities on other divisions; they may prefer projects that require large investments due to empire-building incentives (Harris and Raviv (1996), Berkovitch and Israel (2004)); or the department and the university may put different weights on candidates’ research and nonresearch attributes. Armed with the proposal-setting power, the agent has a tendency to propose his favorite project and hide his less preferred ones, even if those projects are “superstars” for the principal.

This paper explores how the principal counteracts the agent’s proposing bias. We address

¹For instance, Ross (1986) surveyed twelve large manufacturers in 1981-82. In the area of energy preservation alone, he collected data on 400 discretionary projects. Roughly 100 projects were completed in 1980-81, and 300 projects were either underway or prospective, intended for the period 1982-85.

²Stanley and Block (1984) surveyed 121 multinational firms; 88 percent of the respondents reported that projects were initiated from bottom up, and 70 percent reported that decision-making was centralized.

this question using a robust-design, non-Bayesian approach. Such an approach is useful when the principal faces a certain project-choice problem for the first time. It is also applicable to organizations in which the principal oversees many project-choice problems, but each problem offers little guidance on how to formulate a prior belief for another. For instance, the headquarters might oversee project choice in various divisions which operate in quite different markets. Similarly, the university oversees hiring in a range of departments; moreover, the academic labor market varies widely not only across departments but also across years.

Due to the agent’s private information, no mechanism can guarantee that the principal’s favorite project among the available ones will be chosen. We define the principal’s *regret* as the difference between his payoff from his favorite project and his expected payoff from the project chosen under the mechanism. Thus, regret can be interpreted as “money left on the table” due to the agent’s private information. We characterize a mechanism that minimizes the principal’s worst-case regret, i.e., a mechanism that leaves as little money on the table as possible in all circumstances. As we will show, this minimax regret approach allows us (i) to explain why certain incentive schemes are common, (ii) to propose new incentive schemes, and (iii) to explore questions which are intractable under the Bayesian approach.

Main results. The principal chooses at most one project. Yet, one often observes that the agent proposes multiple projects.³ The strategic role that a multiproject proposal plays in project-choice problems has not yet been explored. Do several projects within the proposal have a chance of being chosen? How does the principal benefit from allowing the agent to propose multiple projects rather than to propose only a single project?

To answer these questions, we distinguish between two environments. In the *multiproject* environment, the agent can propose any subset of the available projects. In the *single-project* environment, the agent can propose only one available project. Such a single-project restric-

³For instance, it is usual practice for a search committee to propose multiple finalists for an appointment as a dean at a business school (Byrne (2018, 2021)).

tion may emerge in organizations due to the principal's limited attention. It is also applicable to antitrust regulation, where a firm proposes one merger and the regulator decides whether to approve or reject the firm's proposal (Lyons (2003), Neven and Röller (2005), Armstrong and Vickers (2010), Ottaviani and Wickelgren (2011), Nocke and Whinston (2013)). We take the environment as exogenous and derive the optimal mechanism in each environment.

In the single-project environment, a mechanism specifies for each single-project proposal the probability that it will be approved. We show that the optimal mechanism has a two-tier structure. If the proposed project gives the principal a sufficiently high payoff, it is a top-tier project for the principal and is approved for sure. Otherwise, it is a bottom-tier project and is approved only with some probability. The probability that a bottom-tier project is approved decreases in its payoff to the agent, in order to deter the agent from hiding projects that are more valuable for the principal.

The optimal two-tier mechanism in the single-project environment aligns the agent's proposing incentives with the principal's preferences in two ways. First, if the agent has at least one top-tier project, then he will propose a top-tier project. Second, if all his projects are bottom-tier ones, then he will propose a project that maximizes the principal's expected payoff. For the principal to reject a bottom-tier project is suboptimal ex post, but is indispensable for aligning the agent's proposing incentives ex ante.

In the multiproject environment, a mechanism specifies for each proposed set of projects if and how a project will be chosen, namely, a randomization over the proposed projects and "no project." We first show that the revelation principle holds in this environment, so it is without loss of generality to focus on mechanisms under which the agent optimally proposes all the available projects.

If the agent in the multiproject environment actually proposes a single project, the optimal mechanism again has a two-tier structure. In particular, if this project's payoff to the principal is sufficiently high, he views it as a top-tier project and chooses it for sure. Other-

wise, the project is a bottom-tier one and is chosen with some probability that decreases in its payoff to the agent.

If the agent proposes multiple projects, his expected payoff must be at least his maximal expected payoff from proposing each project alone. This is because the agent had the option to propose just one of the available projects. The optimal mechanism randomizes over the proposed projects and “no project” to maximize the principal’s expected payoff, subject to the constraint of promising this maximal expected payoff to the agent.⁴ Since the agent gets his maximal expected payoff from proposing each project alone in the proposal, we call this mechanism a *proposal-wide maximal-payoff mechanism* (PMP mechanism).

The optimal PMP mechanism identifies the strategic role that multiproject proposals play in project-choice problems. Among the available projects, there is one project that gives the agent his maximal expected payoff from proposing each project alone. This project is what the agent himself wants to propose and pins down his expected payoff. The multiproject environment allows the agent to also propose other projects, so it can deliver this expected payoff to the agent more efficiently than the single-project environment does.




Figure 1: Strategic role of multiproject proposals

To illustrate this strategic role, consider the example in Figure 1. The agent has two projects, ★ and ▲. According to the optimal mechanism in subsection 4.3, if the agent

⁴This maximization problem is a linear programming problem. There exists an optimal solution whose support has at most two projects.

proposed \blacktriangle alone, it would be chosen with probability $1/2$. The agent’s and the principal’s expected payoffs would be x and y , respectively. Since x is higher than the agent’s payoff from \blackstar , proposing \blacktriangle gives the agent his maximal payoff from proposing each project alone. The multiproject environment allows the agent to also propose \blackstar , a better fallback option than rejection. The mechanism randomizes between \blacktriangle and \blackstar while promising the agent an expected payoff of x . This lifts the principal’s expected payoff from y to z .

Our result also explains when more than one proposed project have a chance of being chosen. If the agent’s favorite project is already a top-tier project for the principal, the two parties’ preferences are somewhat aligned. We show that this project will be chosen for sure. In contrast, if the agent’s favorite project is a bottom-tier project for the principal, their preferences are less aligned. The mechanism may randomize between multiple proposed projects. Figure 1 gives such an example: the mechanism “finds some middle ground” between the two parties—with some probability the choice favors the agent and with some probability it favors the principal.

We further analyze the environment in which the agent can propose up to K projects for some fixed $K \geq 2$. We show that the aforementioned PMP mechanism admits an implementation in which (i) the agent proposes up to two projects, and (ii) the principal achieves the same worst-case regret as he does in the multiproject environment. The first project is what gives the agent his maximal expected payoff from proposing each project alone; this is what the agent himself wants to propose. The second is the principal’s favorite project; this can be the fallback option when the first project is not chosen. Hence, in terms of minimizing the principal’s worst-case regret, allowing for up to two projects is enough.

Our results offer a number of insights and implications for the practical design of project-choice rules.

(1) The optimal mechanism in the single-project environment is comparable to how projects are chosen in the firms surveyed by Ross (1986). We discuss this connection in

detail in subsection 5.3.1. In the single-project environment, the threshold for being a top-tier project is higher than a zero payoff for the principal; such a project is approved for sure and the agent never proposes a bottom-tier project when he has a top-tier one. This high threshold for top-tier projects suggests one possible explanation for the finding in Poterba and Summers (1995). They show that US firms, when evaluating investment projects, use hurdle rates which are considerably higher than the cost of capital. In academic departments, top-tier projects are sometimes called slam-dunk cases or star candidates. In our experience, it is a no-brainer to propose such a candidate, who is typically approved with no questions asked. In contrast to the sure approval of top-tier projects, bottom-tier ones are rejected with positive probabilities. This result provides one possible explanation for why, in the case of faculty hiring, no slot is ever guaranteed and random rejection occurs.

(2) We show that the principal's worst-case regret in the multiproject environment is significantly lower than that in the single-project one. Our result suggests one possible explanation for the practice of proposing more projects than are needed to the principal. In particular, senior positions like a deanship are vital to the success of an organization. It is reasonable to expect that for such positions the benefit from the multiproject environment outweighs the cost of reviewing multiple candidates, so the multiproject environment is more likely to be adopted for such positions than for lower-level ones. We show that in the multiproject environment, the principal does not always choose his own favorite one among those proposed. One testable implication is that, if an outsider collects all the interactions where multiple projects are proposed to the principal, the choice is not always the principal's favorite one, but sometimes reflects the proposing party's preferences.

Related literature. Our paper is closely related to Armstrong and Vickers (2010) and Nocke and Whinston (2013), who study the project-choice problem using the Bayesian approach. Armstrong and Vickers (2010) characterize the optimal deterministic mechanism in the single-project environment and show through examples that the principal does strictly

better if either randomization or multiproject proposals are allowed. Their discussion highlights the informational role of multiproject proposals.⁵ We complement their discussion by identifying the strategic role of multiproject proposals. Nocke and Whinston (2013) focus on mergers (i.e., projects) that are ex ante different and further incorporate the bargaining process among firms. They show that a tougher standard is imposed on mergers involving larger partners. We depart from these papers by taking the minimax regret approach to this multidimensional screening problem. This more tractable approach allows us to explore questions which are intractable under the Bayesian approach, including how much the principal benefits from the multiproject environment, from randomization, and from a smaller project domain.

If the principal is restricted to the single-project environment and deterministic mechanisms, a mechanism is effectively a set of permitted projects that the agent can choose with no questions asked. This connects project-choice problems to the delegation literature initiated by Holmström (1984).⁶ If either the multiproject environment or randomization is allowed, more mechanisms can be implemented than by simply specifying a set of permitted projects (as shown, respectively, in subsection 3.3 and footnote 8 of Armstrong and Vickers (2010)). Our contribution is to characterize the optimal randomized mechanisms in both the single-project and the multiproject environments.

Berkovitch and Israel (2004) study a project-choice model with binary project types (high versus low capital-intensive). The agent has an incentive to hide the principal's preferred project (which is low capital-intensive). They solve for the optimal randomized mechanism in which the principal can compensate the agent for certain proposals. They characterize

⁵The informational role of multiproject proposals can be explained by the following example. If the principal is quite certain that there are three available projects, he can get the agent's information almost for free in the multiproject environment. The mechanism chooses the principal's favorite project if three projects are proposed and chooses no project otherwise.

⁶Some later contributions include, for example, Melumad and Shibano (1991), Martimort and Semenov (2006), Alonso and Matouschek (2008), and Amador and Bagwell (2013).

conditions under which monetary incentives are not used. The principal rejects the high capital-intensive project with a positive probability, so the agent always reveals the low capital-intensive project whenever it is available. Our two-tier mechanism in the single-project environment has a similar feature: bottom-tier projects are rejected with positive probabilities so the agent proposes a top-tier project if he has at least one such project.

Che, Dessen and Kartik (2013) and Goel and Hann-Caruthers (2020) consider a project-choice problem in which the number of available projects is public information. In Che, Dessen and Kartik (2013), the agent privately learns the projects' payoffs and sends a cheap-talk message. They show that the principal benefits from randomizing between a default project and the agent's recommendation. In Goel and Hann-Caruthers (2020), the agent's only constraint is not to overreport projects' payoffs to the principal. Because the agent in these papers cannot hide projects as our agent can, he loses the proposal-setting power. The resulting incentive schemes are thus quite different.

The minimax regret approach to uncertainty dates back to Wald (1950) and Savage (1951). It has since been used broadly in game theory, mechanism design, and machine learning. A decision-theoretical axiomatization for the minimax regret criterion can be found in Milnor (1954) and Stoye (2011). Our paper contributes especially to the literature on mechanism design with the minimax regret approach. To start with, Hurwicz and Shapiro (1978) examine a moral hazard problem. Bergemann and Schlag (2008, 2011) examine monopoly pricing. Renou and Schlag (2011) apply the solution concept of ε -minimax-regret to the problem of implementing social choice correspondences. Beviá and Corchón (2019) examine the contest which minimizes the designer's worst-case regret. Malladi (2022) studies the optimal approval rules for innovation, and Guo and Shmaya (2023) study the optimal mechanism for monopoly regulation. More broadly, we contribute to the growing literature of mechanism design with worst-case objectives. For a survey on robustness in mechanism design, see Carroll (2019).

2 Model and mechanism

Let D be the domain of all possible projects. The agent's private *type* $A \subseteq D$ is a finite set of available projects. The agent proposes a set P of projects, and the principal can choose one project from this set. The set P is called the agent's *proposal*. It must satisfy two conditions. First, the agent's proposal is *verifiable* in the sense that he can propose only available projects. Hence, the agent's proposal must be a subset of his type, $P \subseteq A$. Second, $P \in \mathcal{E}$ for some fixed set \mathcal{E} of subsets of D . The set \mathcal{E} captures all the exogenous restrictions on the proposal. For instance, in the setting of antitrust regulation, the agent is restricted to proposing at most one project. In some organizations, the principal has limited attention, so the agent can propose at most a certain number of projects.

Let $u : D \rightarrow \mathbf{R}_+$ be the agent's payoff function and $v : D \rightarrow \mathbf{R}_+$ the principal's payoff function. If project a is chosen, the agent's payoff is $u(a)$ and the principal's payoff is $v(a)$. If no project is chosen, both parties get zero.

We begin with two environments which are natural first steps: *single-project* and *multi-project*. In the single-project environment, the agent can propose at most one available project, so $\mathcal{E} = \{P \subseteq D : |P| \leq 1\}$. In the multiproject environment, the agent can propose any set of available projects so $\mathcal{E} = 2^D$, the power set of D . In subsection 5.1, we discuss the intermediate environments in which the agent can propose up to K projects for some fixed number $K \geq 2$.

A *subprobability measure over D with a finite support* is given by $\pi : D \rightarrow [0, 1]$ such that

$$\text{support}(\pi) = \{a \in D : \pi(a) > 0\}$$

is finite, and $\sum_a \pi(a) \leq 1$. When we say that a project *is chosen from* a subprobability measure π with finite support, we mean that project a is chosen with probability $\pi(a)$, and

that with probability $1 - \sum_a \pi(a)$ no project is chosen.

The principal's ability to reject all proposed projects (i.e., to choose no project) is crucial for him to retain some bargaining power. If the principal were required to choose a project as long as the agent has proposed some, then the agent would propose only his own favorite project, which would be chosen for sure.

A *mechanism* ρ attaches to each proposal $P \in \mathcal{E}$ a subprobability measure $\rho(\cdot|P)$ such that $\text{support}(\rho(\cdot|P)) \subseteq P$. The interpretation is that, if the agent proposes P , then a project is chosen from the subprobability measure $\rho(\cdot|P)$. Thus, the agent's expected payoff under the mechanism ρ if he proposes P is $U(\rho, P) = \sum_{a \in P} u(a)\rho(a|P)$.

A *choice function* f attaches to each type A of the agent a subprobability measure $f(\cdot|A)$ such that $\text{support}(f(\cdot|A)) \subseteq A$. The interpretation is that, if the set of available projects is A , then a project is chosen from the subprobability measure $f(\cdot|A)$.

A choice function f is *implemented* by a mechanism ρ if, for every type A of the agent, there exists a probability measure μ with support over $\text{argmax}_{P \subseteq A, P \in \mathcal{E}} U(\rho, P)$ such that $f(a|A) = \sum_P \mu(P)\rho(a|P)$ for every $a \in A$. The agent chooses only proposals that give him the highest expected payoff and can randomize among them.

2.1 Revelation principle in the multiproject environment

Consider the multiproject environment $\mathcal{E} = 2^D$. A mechanism ρ is *incentive-compatible (IC)* if the agent finds it optimal to propose all the available projects. That is, $U(\rho, A) \geq U(\rho, P)$ for every finite set $A \subseteq D$ and every subset $P \subseteq A$. Equivalently, a mechanism ρ is IC if and only if the agent's payoff, $U(\rho, P)$, weakly increases in P with respect to set inclusion. The following proposition states the revelation principle in the multiproject environment.

Proposition 2.1. *Assume $\mathcal{E} = 2^D$. If a choice function f is implemented by some mechanism, then the mechanism f is IC and implements the choice function f .*

When the agent proposes a set P of projects, he provides evidence that his type A satisfies

$P \subseteq A$. In the multiproject environment, the agent can provide the maximal evidence for his type. This property is called *normality* by Bull and Watson (2007), which is equivalent to the *full reports condition* of Lipman and Seppi (1995). Another interpretation of the multiproject environment is to view an agent who proposes a set P as an agent who declares that his type is P . The relation that “type A can declare to be type B ” between types is reflexive and transitive, by the corresponding properties of the inclusion relation between sets. Transitivity is called the *nested-range condition* in Green and Laffont (1986). Previous papers (e.g., Green and Laffont (1986), Bull and Watson (2007)) have shown that the revelation principle holds under these conditions. We repeat the argument for completeness.

Proof of Proposition 2.1. Assume that the mechanism ρ implements the choice function f . Then for every finite set $A \subseteq D$ and every subset $P \subseteq A$, we have:

$$U(f, A) = \max_{Q \subseteq A} U(\rho, Q) \geq \max_{Q \subseteq P} U(\rho, Q) = U(f, P),$$

where the inequality follows from the fact that $Q \subseteq P$ implies $Q \subseteq A$, and the two equalities follow from the fact that ρ implements f . Hence, the mechanism f is IC. Also, by definition, if the mechanism f is IC, then it implements the choice function f . \square

Since an implementable choice function is itself an IC mechanism and vice versa, we will use both terms interchangeably when we discuss the multiproject environment.

3 The principal’s problem

The principal’s *regret* from a choice function f when the set of available projects is A is:

$$\text{RGRT}(f, A) = \max_{a \in A} v(a) - \sum_{a \in A} v(a) f(a|A).$$

The regret is the difference between what the principal could have achieved if he had known the set A of available projects and what he actually achieves. The *worst-case regret* (WCR) from a choice function f is:

$$\text{WCR}(f) = \sup_{A \subseteq D, |A| < \infty} \text{RGRT}(f, A),$$

where the supremum ranges over all possible types of the agent (i.e., all possible finite sets of available projects). The principal’s problem is to minimize $\text{WCR}(f)$ over all implementable choice functions f . This step is our only departure from the Bayesian approach. The Bayesian approach instead assigns a prior belief over the number and the payoff characteristics of the available projects. The principal’s problem under the Bayesian approach is to minimize the *expected* regret instead of the *worst-case* regret.

Wald (1950) and Savage (1972) propose considering only *admissible* choice functions (i.e., choice functions that are not weakly dominated). An implementable choice function f is *admissible* if there exists no other implementable f' such that the principal’s regret is weakly higher under f than under f' for every type of the agent and strictly higher for some type. For the rest of the paper, we focus mainly on admissible choice functions.

While our principal takes the minimax regret approach towards the agent’s type, he calculates the expected payoff with respect to his own objective randomization. The same assumption is made by Savage (1972) when he discusses the use of randomized acts under the minimax regret approach (Savage, 1972, chapter 9.3). A similar assumption is made in the ambiguity-aversion literature. For example, in Gilboa and Schmeidler (1989), the decision maker calculates his expected payoff with respect to random outcomes (i.e., “roulette lotteries”) but evaluates acts using the maxmin approach with respect to non-unique priors. The alternative assumption that the principal takes the minimax regret approach even towards his own randomization effectively restricts the principal to deterministic mechanisms.

From now on, we assume that the set D of all possible projects is $[\underline{u}, 1] \times [\underline{v}, 1]$ for some parameters $\underline{u}, \underline{v} \in [0, 1]$, and that the functions $u(\cdot)$ and $v(\cdot)$ are projections over the first and second coordinates. Abusing notation, we denote a project $a \in D$ also by $a = (u, v)$, where u and v are the agent's and the principal's payoffs, respectively, if project a is chosen.

The parameters \underline{u} and \underline{v} quantify the uncertainty faced by the principal: the higher they are, the smaller the uncertainty. They also measure the players' preference intensity over projects. As \underline{u} increases, the agent's preferences over projects become less strong, so it becomes easier to align his incentives with those of the principal. As \underline{v} increases, the principal's preferences over projects become less strong, so the agent's tendency to propose his own favorite project becomes less costly for the principal.

4 Optimal mechanisms

4.1 Preliminary intuition

We now use an example to illustrate the fundamental trade-off faced by the principal: choosing the agent's preferred project with a very high probability incentivizes the agent to propose only his preferred project, but rejecting such a project may risk rejecting the only project that is available. We also use this example to illustrate the intuition behind the optimal mechanisms. We first explain how randomization helps to reduce the principal's WCR in the single-project environment. We then explain how the multiproject environment can further reduce the WCR. For this illustration, we assume that $\underline{v} = 0$, so $D = [\underline{u}, 1] \times [0, 1]$.

Consider the single-project environment and assume first that the principal is restricted to deterministic mechanisms. In this case, a mechanism is a set of projects that the principal approves for sure, and all other projects are rejected outright. For each such mechanism, the principal has two fears. First, if the agent has two or more projects which will be approved, then he will propose the one he likes the most, even if projects are available that are more

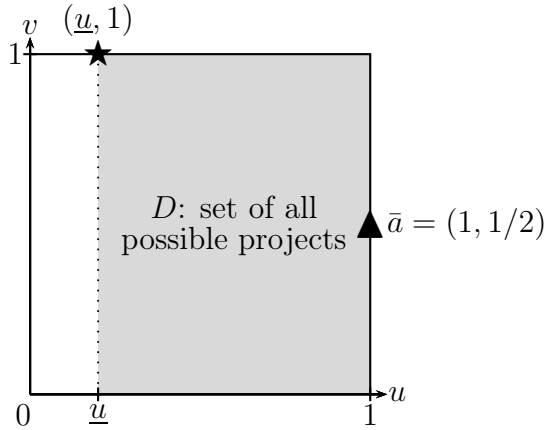


Figure 2: Preliminary intuition, $\underline{v} = 0$

valuable to the principal. Second, if the agent has only projects which will be rejected, then the principal loses the payoff from these projects. As we explain in the next paragraph, these two fears imply that the principal's WCR under any deterministic mechanism is at least $1/2$.

Consider project $\bar{a} = (1, 1/2)$ in Figure 2. It gives the agent his highest payoff 1, while giving the principal only a moderate payoff of $1/2$. If the mechanism approves \bar{a} and the set of available projects is $\{\bar{a}, (\underline{u}, 1)\}$, then the agent will propose \bar{a} rather than $(\underline{u}, 1)$, so the principal suffers regret $1/2$. If the mechanism rejects \bar{a} but \bar{a} turns out to be the only available project, then the principal also suffers regret $1/2$. Thus, the WCR under any deterministic mechanism is at least $1/2$. On the other hand, the deterministic mechanism that approves project (u, v) if and only if $v \geq 1/2$ achieves the WCR of $1/2$, so it is optimal among all the deterministic mechanisms.

We now explain how randomization can reduce the WCR in the single-project environment. We first note that, if $\underline{u} = 0$, then, even with randomized mechanisms, the principal cannot reduce his WCR below $1/2$. This is because, when the set of available projects is $\{\bar{a}, (0, 1)\}$, the only way to incentivize the agent to propose the project $(0, 1)$ is still to reject the project \bar{a} outright if \bar{a} is proposed. However, if $\underline{u} > 0$, then the principal can do better. He can approve the project \bar{a} with probability \underline{u} , while still maintaining the agent's incen-

tive to propose the principal's preferred project $(\underline{u}, 1)$ when the set of available projects is $\{\bar{a}, (\underline{u}, 1)\}$. We carry out this idea in Theorem 4.1 in subsection 4.2.

Let us now consider the multiproject environment. We again begin with deterministic mechanisms. Under deterministic mechanisms, more choice functions can be implemented in the multiproject environment than in the single-project one.⁷ However, when he is restricted to deterministic mechanisms, the principal has the same minimal WCR in the multiproject environment as in the single-project one. This is because, if the principal wants to choose $(\underline{u}, 1)$ when the set of available projects is $\{\bar{a}, (\underline{u}, 1)\}$, then the only way to incentivize the agent to include $(\underline{u}, 1)$ in his proposal is to reject the project \bar{a} when \bar{a} is proposed alone.

We now explain how randomization can help in the multiproject environment, even when $\underline{u} = 0$. While a deterministic mechanism must pick either \bar{a} or $(0, 1)$ or nothing when the agent proposes $\{\bar{a}, (0, 1)\}$, a randomized mechanism can find some middle ground by choosing each project with probability $1/2$. On the other hand, if the agent proposes only \bar{a} , the principal chooses \bar{a} with probability $1/2$, so the agent of type $\{\bar{a}, (0, 1)\}$ is willing to propose $\{\bar{a}, (0, 1)\}$ instead of just \bar{a} . The principal's regret is $1/4$ both when the agent's type is $\{\bar{a}, (0, 1)\}$ and when his type is $\{\bar{a}\}$. We carry out this idea in Theorem 4.2 in subsection 4.3. Specifically, when the agent proposes P , the principal gives the agent the maximal payoff he can offer, subject to the constraint that he can give the agent this same payoff if the agent proposes $P \cup \{(\underline{u}, 1)\}$ and can still keep his regret under control.

4.2 Optimal mechanism in the single-project environment

Since the agent in this environment can propose at most one project, a mechanism specifies the approval probability for each proposed project. Instead of using our previous notation $\rho(a|\{a\})$, we let $\alpha(u, v) \in [0, 1]$ denote the approval probability if the agent proposes the project (u, v) .

⁷For example, the principal can implement the choice function that chooses (i) the agent's favorite project, if there are at least two available projects; and (ii) nothing, if there is at most one available project.

Theorem 4.1 (Single-project environment). Assume $\mathcal{E} = \{P \subseteq D : |P| \leq 1\}$. Let

$$R^s = \max_{v \in [\underline{v}, 1]} \min \{1 - v, (1 - \underline{u})v\} = \min \left\{ \frac{1 - \underline{u}}{2 - \underline{u}}, 1 - \underline{v} \right\}. \quad (1)$$

- (i) The WCR under any mechanism is at least R^s .
- (ii) Let α^s be the mechanism given by:

$$\alpha^s(u, v) = \begin{cases} 1, & \text{if } v \geq 1 - R^s \text{ or } u = 0; \\ \underline{u}/u, & \text{if } v < 1 - R^s \text{ and } u > 0. \end{cases}$$

Under α^s , if the agent has projects with $v \geq 1 - R^s$, it is optimal for him to propose his favorite project among those with $v \geq 1 - R^s$. Otherwise, it is optimal to propose a project that maximizes the principal's expected payoff $\alpha^s(u, v)v$. The corresponding choice function has the WCR of R^s and is admissible.

- (iii) If a mechanism α implements a choice function that has the WCR of R^s , then $\alpha(u, v) \leq \alpha^s(u, v)$ for every $(u, v) \in D$.

We first explain why the WCR under any mechanism α is at least R^s . Pick any $v \in [\underline{v}, 1]$ and consider the project $(1, v)$. If $\alpha(1, v) > \underline{u}$, then the agent will not propose the principal's preferred project $(\underline{u}, 1)$ when both $(1, v)$ and $(\underline{u}, 1)$ are available. The principal's regret is then $1 - \alpha(1, v)v \geq 1 - v$. If $\alpha(1, v) \leq \underline{u}$, then when the agent has only $(1, v)$ available, the regret is $v - \alpha(1, v)v \geq (1 - \underline{u})v$. Therefore, the principal's WCR is at least $\min \{1 - v, (1 - \underline{u})v\}$. The operator $\max_{v \in [\underline{v}, 1]}$ in R^s reflects the fact that the argument holds for any $v \in [\underline{v}, 1]$.

The optimal mechanism α^s has a two-tier structure: a top tier if the principal's payoff v is above $1 - R^s$, and a bottom tier if this payoff is below $1 - R^s$. A top-tier project is approved for sure. A bottom-tier project is approved with probability \underline{u}/u , so the agent expects a constant payoff of \underline{u} from proposing a bottom-tier project. A top-tier project gives the principal a high enough payoff, so his regret is at most $1 - v \leq R^s$ if he approves such

a project. For a bottom-tier project (u, v) , the approval probability cannot exceed \underline{u}/u . Otherwise, the agent will hide $(\underline{u}, 1)$ when he has both (u, v) and $(\underline{u}, 1)$ available, causing high regret for the principal.

The agent will propose a top-tier project if he has at least one such project. If all his projects are in the bottom tier, he finds it optimal to propose a project that maximizes the principal's expected payoff. The principal still suffers regret from two sources. First, if the agent has two or more top-tier projects which will be approved for sure, he will propose what he favors instead of what the principal favors. Second, if the agent has only bottom-tier projects, his proposal is rejected with positive probability. The threshold for the top tier, $1 - R^s$, is chosen to keep the regret from both sources under control.

The approval probability $\alpha^s(u, v)$ increases in v (the principal's payoff) and decreases in u (the agent's payoff). This monotonicity in v and u is natural. In particular, the principal is less likely to approve projects that give the agent high payoffs in order to deter the agent from hiding projects that give the principal high payoffs. It is interesting to compare our optimal mechanism α^s in the single-project environment to that in Armstrong and Vickers (2010). They characterize the optimal deterministic mechanism in a Bayesian setting. Under the assumptions that (i) projects are i.i.d. and (ii) the number of available projects is independent of their payoff characteristics, they show that the optimal deterministic mechanism $\alpha(u, v)$ increases in v : a project (u, v) is approved for sure if $v \geq r(u)$ and rejected outright if $v < r(u)$. The function $r(u)$ depends on the distributions of the number and the payoffs of the available projects. Our result differs in three aspects. First, under our prior-free approach, the threshold for the top tier (i.e., where projects are approved for sure) depends only on the principal's payoff v . Second, we show how to reduce inefficient rejections in the bottom tier without jeopardizing the agent's incentive to propose top-tier projects. Third, we show that the approval probability $\alpha^s(u, v)$ is monotone decreasing in u .

It is possible that multiple mechanisms achieve the WCR of R^s . Statement (iii) in

Theorem 4.1 says that the mechanism α^s is uniformly more generous in approving the agent's proposal than any other mechanism that has the WCR of R^s . This statement has two implications. First, among all mechanisms that have the WCR of R^s , the mechanism α^s is the agent's most preferred one for every possible type A . Second, compared to any mechanism that has the WCR of R^s , the mechanism α^s gives the principal a higher payoff (or equivalently, a lower regret) for every singleton A .

4.3 Optimal mechanism in the multiproject environment

We now present the optimal mechanism in the multiproject environment. Let $\alpha : [\underline{u}, 1] \times [\underline{v}, 1] \rightarrow [0, 1]$ be a function and consider the following *proposal-wide maximal-payoff mechanism* (PMP mechanism) induced by the function α :

- (1) If the proposal P has only one project (u, v) , it is approved with probability $\alpha(u, v)$.
- (2) If the proposal P has multiple projects, the mechanism randomizes over the proposed projects and no project to maximize the principal's expected payoff, while promising the agent an expected payoff of $\max_{(u,v) \in P} \alpha(u, v)u$. This is the maximal expected payoff the agent could get from proposing each project alone.

By the definition of a PMP mechanism, the more projects the agent proposes, the weakly higher his expected payoff will be. Therefore, PMP mechanisms are incentive-compatible. Under any incentive-compatible mechanism, the agent's payoff from proposing multiple projects must be at least his maximal payoff from proposing each project alone. A PMP mechanism has the feature that the agent receives exactly his maximal payoff from proposing each project alone, but not more.

Our next theorem shows that there exists a PMP mechanism that achieves the minimal WCR. Moreover, it is the agent's most preferred mechanism among all the mechanisms that achieve the minimal WCR and are admissible.

Theorem 4.2 (Multiproject environment). Assume $\mathcal{E} = 2^D$. For every $u \in [\underline{u}, 1]$ and $p \in [0, 1]$, let $\gamma(u, p)$ be

$$\gamma(u, p) = \min\{q \in [0, 1] : \underline{u} + q(u - \underline{u}) \geq pu\}. \quad (2)$$

Let

$$R^m = \max_{(u,v) \in D} \min_{p \in [0,1]} \max\{(1-p)v, \gamma(u,p)(1-v)\}. \quad (3)$$

(i) The WCR under any mechanism is at least R^m .

(ii) Let ρ^m be the PMP mechanism induced by

$$\alpha^m(u, v) = \max\{p \in [0, 1] : \gamma(u, p)(1 - v) \leq R^m\}. \quad (4)$$

Under ρ^m , it is optimal for the agent to propose his type A . The corresponding choice function ρ^m has the WCR of R^m . It is admissible for $\underline{v} > 0$.

(iii) If ρ is an IC and admissible mechanism which has the WCR of R^m , then $U(\rho, A) \leq U(\rho^m, A)$ for every type A .

We now explain the intuition behind $\gamma(u, p)$, R^m , and $\alpha^m(u, v)$ in (2) to (4). Pick any project (u, v) and suppose p is the approval probability if the agent proposes (u, v) alone, so he expects a payoff of pu from doing so. Then, if the agent has both (u, v) and $(\underline{u}, 1)$ available, he must expect a weakly higher payoff from proposing both projects than from proposing (u, v) alone. This implies that, when the agent proposes both projects, the probability of choosing the agent's preferred project (u, v) is at least $\gamma(u, p)$ while the probability of choosing the principal's preferred project $(\underline{u}, 1)$ is at most $1 - \gamma(u, p)$. The principal's WCR is then at least $\max\{(1-p)v, \gamma(u,p)(1-v)\}$. The first term, $(1-p)v$, is the regret when (u, v) is the only available project, while the second term, $\gamma(u,p)(1-v)$, reflects the regret when both (u, v) and $(\underline{u}, 1)$ are available. The operators $\min_{p \in [0,1]}$ and $\max_{(u,v) \in D}$ in R^m

reflect, respectively, the fact that the principal can choose p and the fact that the argument holds for any $(u, v) \in D$. Lastly, $\alpha^m(u, v)$ is the highest approval probability for (u, v) when it is proposed alone, such that (i) the agent can be incentivized to propose both available projects when his type is $A = \{(u, v), (\underline{u}, 1)\}$, and (ii) the regret from type A can be kept under R^m .

The explicit expressions for R^m and $\alpha^m(u, v)$ are presented at the end of this subsection. It follows from expressions (2) to (4) that $\alpha^m(u, v) = 1$ if $v \geq 1 - R^m$ or $u = \underline{u}$, and $\alpha^m(u, v) < 1$ otherwise. As in the case of the single-project environment, when the agent proposes only one project, it is approved for sure if its payoff to the principal is sufficiently high. For this reason, we call a project with $v \geq 1 - R^m$ a top-tier project and a project with $v < 1 - R^m$ a bottom-tier project. Figure 3 depicts these two tiers when $\underline{u} = \underline{v} = 0$.

When the agent proposes multiple projects, the principal promises the agent an expected payoff of $\max_{(u,v) \in P} \alpha^m(u, v)u$. In both panels of Figure 3, each dotted curve connects all the projects that induce the same value of $\alpha^m(u, v)u$, so it can be interpreted as an “indifference curve” for the agent. For a project in the top tier, the principal is willing to compensate the agent his full payoff from such a project. In contrast, for a project in the bottom tier, the principal is willing to compensate the agent only a discounted payoff. The lower the project’s payoff to the principal, the more severe the discounting. Hence, indifference curves are vertical in the top tier and downward sloping from left to right in the bottom tier.

The agent’s expected payoff is determined by the project (among those proposed) that is on the highest indifference curve. This is the project that the agent himself wants to propose.

So when does the principal benefit from the agent’s also proposing other projects? Figure 3 gives an intuitive answer. If the agent’s favorite project is already a top-tier project for the principal, then this project will be chosen for sure, so there is no benefit from also proposing other projects. The left panel gives such an example: \star and \blacktriangle denote the available projects and \blacktriangle will be chosen for sure. In contrast, if the agent’s favorite project is a bottom-tier




Figure 3: Optimal mechanism in the multiproject environment, $\underline{u} = \underline{v} = 0$

project for the principal, the benefit from the agent's proposing other projects as well can be significant. The right panel illustrates such an example. Among the available projects, project ■ is on the highest indifference curve, so it pins down the maximal expected payoff the agent could get from proposing each project alone. This maximal payoff is denoted by x in the figure. The mechanism randomizes between ▲ and ★ while promising the agent an expected payoff of x . In such cases, multiple projects in the proposal have a chance of being chosen: sometimes the choice favors the agent, and at other times it favors the principal.

Statement (iii) in Theorem 4.2 says that the mechanism ρ^m is the agent's most preferred mechanism among all the mechanisms that have the WCR of R^m and are admissible. This preference for the mechanism ρ^m holds for every possible type A of the agent.

Lastly, the explicit expressions for R^m and α^m are given below. The probability $\alpha^m(u, v)$ is increasing in v and decreasing in u . It is strictly so in the bottom tier.

$$R^m = \begin{cases} \frac{(1-u)(2-u-2\sqrt{1-u})}{u^2} & \text{if } \underline{v} < \frac{1-\sqrt{1-u}}{u}; \\ \frac{(1-u)(1-\underline{v})\underline{v}}{1-\underline{u}\underline{v}} & \text{otherwise;} \end{cases}$$

and

$$\alpha^m(u, v) = \begin{cases} 1, & \text{if } v \geq 1 - R^m \text{ or } u = \underline{u}; \\ (1 - \frac{R^m}{1-v}) \frac{\underline{u}}{u} + \frac{R^m}{1-v}, & \text{if } v < 1 - R^m \text{ and } u > \underline{u}. \end{cases}$$

Remark 1. The PMP mechanism ρ^m is weakly IC. We can modify this mechanism to make it strictly IC. To do this, let $h : [0, \infty) \rightarrow [0, 1]$ be a strictly increasing function and $|P|$ the number of projects in the proposal P . We modify α^m with $(1 - \varepsilon + \varepsilon h(|P|))\alpha^m$ for some small $\varepsilon > 0$ and consider the PMP mechanism induced by this modified function. If the agent has some projects with $u > 0$, then it is strictly optimal for him to propose all the available projects. The WCR under this mechanism approaches R^m as $\varepsilon \rightarrow 0$. (If the agent has only projects with $u = 0$, then there is no way to strictly incentivize him to propose.)

4.4 Comparing the WCR under the two environments

According to the definitions in (1) and (3), both R^s and R^m decrease in \underline{u} and \underline{v} . Figure 4 compares the values of R^s and R^m as \underline{u} or \underline{v} varies. The left panel depicts the WCR as a function of \underline{u} for a fixed \underline{v} . The right panel depicts the WCR as a function of \underline{v} for a fixed \underline{u} . Roughly speaking, the principal's gain from having the multiproject environment instead of the single-project environment, as measured by $R^s - R^m$, is larger when \underline{u} or \underline{v} is smaller (i.e., when the principal faces more uncertainty or when players can potentially have stronger preferences over projects).




Figure 4: WCR: single-project (dashed curve) vs. multiproject (solid curve)

5 Extensions and discussion

5.1 Intermediate environments

So far, we have focused on the single-project and the multiproject environments, which are natural first steps. Nevertheless, there are intermediate environments in which the agent can propose up to K projects for some fixed $K \geq 2$, so that $\mathcal{E} = \{P \subseteq D : |P| \leq K\}$. We call this case the K -project environment. The following proposition shows that as long as $K \geq 2$, the principal's minimal WCR is R^m .

Proposition 5.1 (Two is enough in minimizing WCR). *Consider the K -project environment.*

For any $K \geq 2$,

- (i) the WCR under any mechanism is at least R^m ;*
- (ii) under the PMP mechanism ρ^m induced by $\alpha^m(u, v)$, it is optimal for the agent to propose the principal's favorite project and the project that maximizes $\alpha^m(u, v)u$. The corresponding choice function has the WCR of R^m .*

Proof. Let A be the set of available projects. Let $(u_p, v_p) \in \operatorname{argmax}\{v : (u, v) \in A\}$ be a principal's favorite project. Let $(u_a, v_a) \in \operatorname{argmax}\{\alpha^m(u, v)u : (u, v) \in A\}$ be a project that gives the agent his maximal payoff from proposing each project alone. Let $P = \{(u_p, v_p), (u_a, v_a)\}$. Under the PMP mechanism ρ^m induced by $\alpha^m(u, v)$, the agent is willing to propose P since this proposal gives him $\alpha^m(u_a, v_a)u_a$, the maximal payoff he can get under ρ^m . The principal's payoff given proposal P equals his payoff if the set of available projects were actually P . By Theorem 4.2 this payoff is at least $v_p - R^m$, so the principal's regret is at most R^m . \square

Based on Theorem 4.1 and Proposition 5.1, the principal's minimal WCR drops distinctly from R^s to R^m as K goes from one to two or above. If the agent can propose at most one project, the only fallback option for the principal is to reject the proposed project. By contrast, if the agent can propose two or more projects, he can provide the principal with a

much better fallback option than rejection. This is exactly how the principal benefits from allowing the agent to propose more than one project.

Specifically, the agent can propose the following two projects: the project that gives the agent his maximal payoff from proposing each project alone, and the principal’s favorite project. The first project is what the agent himself wants to propose. The second project can be the fallback option when the first project is not chosen. According to the proof of Proposition 5.1, the randomization between these two projects alone already ensures that the principal’s WCR does not exceed R^m .

This result also suggests a parsimonious way to implement the PMP mechanism ρ^m in the multiproject environment—a way which keeps the principal’s WCR at the level of R^m . Instead of asking the agent to propose all the available projects, the principal can ask him to propose the aforementioned two projects: the one that gives the agent his maximal payoff from proposing each project alone and the one that the principal likes the most. This two-project implementation, however, leads to a choice function that is weakly dominated by the choice function if the agent proposed all the available projects, as shown by the right panel of Figure 3. If the agent proposes all three projects, the mechanism randomizes between \star and \blacktriangle while promising the agent the same payoff he would get from proposing \blacksquare alone. If the agent proposes two projects, \blacksquare and \star , the mechanism randomizes between these two while promising the agent the same payoff he would get from proposing \blacksquare alone; this leads to a strictly lower payoff for the principal than if the agent proposed all three projects.

Lastly, for any $K \geq 3$, the principal can implement the same choice function as that in statement (ii) of Theorem 4.2. This is because the linear programming problem that describes mechanism ρ^m (problem (7) in section 7) has an optimal solution whose support has at most two projects. The agent can propose these two projects and the project that gives the agent his maximal payoff from proposing each project alone.

5.2 Cheap-talk communication does not help for any \mathcal{E}

We could have started from a more general definition of a mechanism that chooses a project based on both the proposal P and a cheap-talk message m from the agent. However, cheap talk does not benefit the principal in our model, because the principal observes the payoffs that each project in P gives to both parties. There is no relevant information asymmetry left after the agent proposes P , so there is no benefit to cheap talk. (If, instead, the principal does not observe the agent's payoffs from projects in P , then cheap talk may help.)

More specifically, for any P and any cheap-talk messages m_1, m_2 , we argue that it is without loss for the principal to choose the same subprobability measure over P after (P, m_1) and after (P, m_2) . Suppose otherwise that the principal chooses a measure π_1 after (P, m_1) and chooses π_2 after (P, m_2) . If the agent strictly prefers π_1 to π_2 , then he can profitably deviate to (P, m_1) whenever he is supposed to say (P, m_2) . Hence, (P, m_2) never occurs on the equilibrium path. If the agent is indifferent between π_1 and π_2 , then the principal can pick his preferred measure between π_1 and π_2 after both (P, m_1) and (P, m_2) , without affecting the agent's incentives. This argument does not depend on the exogenous restriction \mathcal{E} on the agent's proposal P , so cheap-talk communication does not help for any \mathcal{E} .

5.3 The commitment assumption

Commitment is crucial for the principal to have some “bargaining power” in the project-choice problem. If the principal has no commitment power, sequential rationality requires that he choose his favorite project among the proposed projects. The agent now has all the bargaining power. He will propose only his favorite project, which will be chosen for sure.

In the multiproject environment, the full-commitment solution involves two types of ex post suboptimality. First, it is possible that no project is chosen even though the agent has proposed some. Second, it is possible that a worse project for the principal is chosen even though a better project for him is also proposed. Some applications may fall between the

full-commitment and the no-commitment settings: the principal can commit to choosing no project but cannot commit to choosing a worse project when a better project is also proposed. In such a partial-commitment setting, a multiproject proposal is effectively a single-project proposal which contains only the principal’s favorite project among the proposed ones. The optimal mechanism in this partial-commitment setting is then the same as that in the single-project environment characterized in Theorem 4.1. For a formal discussion of this partial-commitment environment, we refer readers to subsection 7.3 and Theorem 7.11.

5.3.1 Commitment to a randomized mechanism

We have shown that randomization plays a crucial role in counteracting the agent’s bias to propose his favorite project and hide his less preferred ones. This raises the question of how an organization implements a randomized mechanism. In addition, if we justify commitment to randomization based on repeated interaction and reputation, a tension might exist between repeated interaction and our prior-free approach.⁸ We address both points through the lens of an example from Ross (1986). Ross (1986) surveyed twelve large manufacturers and described how discretionary projects are chosen through a hypothetical example. The example shows how an organization can vary the approval probabilities for different projects using its institutional rules and the reputation of the approval authority. It also illustrates the fact that building reputation through repeated interaction and being able to form a prior belief do not always go hand in hand.

In this example, an employee has identified an approach to cutting energy costs at a heater using advanced combustion controls. He knows other approaches such as replacing the heater or adding heat exchangers. “However the projects overlap; he can advocate at most one” (Ross (1986), p. 16). Levels of approval authority vary with project sizes. The final decisions for small projects (those which cost less than \$1 million) are made at division headquarters.

⁸We thank anonymous referees for this insightful observation.

For medium projects (those between \$1 and \$10 million), they are made at the Corporate Investment Committee (CIC). The employee considers whether to propose a small project or a medium one. He has observed that going to the CIC requires “enthusiastic support...from the division, plant, and facility managers” (Ross (1986), p. 17). He has also observed that managers give higher priority to new markets or to improving the manufacturing process than to cost cutting, so enthusiastic support “would be hard to get for a project that cuts cost but has no other production benefit” (Ross (1986), p. 17). These observations imply a low approval probability for his medium-sized project. He chooses to propose the small one.

The optimal mechanism in our single-project environment provides one possible explanation for the practice in this example. Projects in high-priority areas correspond to top-tier projects, while those in low-priority areas correspond to bottom-tier projects. Among bottom-tier projects, larger projects require broader and more enthusiastic support, so they face lower approval probabilities. This counteracts employees’ tendency to propose larger projects due to empire-building incentives.

The example also suggests when reputation through repeated interaction can coexist with a prior-free approach. The principal can build a reputation for honest randomization if the agent (or an outside observer) observes previous interactions in which the principal committed to some randomized mechanism. In order for the commitment to be statistically testable, these interactions need not be with the same agent or occur in similar situations. In the example, the employee observes decisions on previous proposals by other employees, which spanned various areas and divisions. On the other hand, forming a prior belief requires some stationarity in the environment. Previous situations must be sufficiently similar to the current one, so the principal can use data on previous situations to form a prior belief about the current one. In the example, the employee is advocating a cutting-edge technology in a specific area of energy conservation, so it is unlikely that the principal has encountered many similar situations before. In fact, discretionary projects are mainly on expanded markets,

new business, or new technologies for cost cutting, so they are likely to be uncharted territory for the principal. However, being in uncharted territory does not deprive the principal of his reputation for honest randomization.

6 Conclusion

We study project-choice problems in which the agent privately observes which projects are available. How to counteract the agent's bias to propose his favorite project and hide his less preferred ones is a first-order concern. We characterize prior-free mechanisms that counteract this bias. If the agent can propose only a single project, the optimal mechanism has a two-tier structure: a top-tier project is approved with no questions asked while a bottom-tier project will be scrutinized. The more a bottom-tier project benefits the agent, the less likely it will be approved. If the agent can propose multiple projects, he can provide the principal with better fallback options than rejection. We show that the principal does not always choose his favorite project among the proposed projects.

Our analysis shows that the minimax regret approach not only is tractable but also preserves the fundamental trade-off faced by the principal. The approach might be useful for studying other variants of the project-choice problem. For instance, the principal may not observe the agent's payoffs from proposed projects; the project-choice process may involve a deeper hierarchy involving three or more parties; the principal may be able to acquire information about the agent's type at a cost; or the agent may need to exert effort to discover projects (Armstrong and Vickers (2010)). We leave these questions for future research.

7 Proofs

7.1 Proof of Theorem 4.1

We have proven statement (i) in the first paragraph after Theorem 4.1.

Claim 7.1. *The choice function under α^s in statement (ii) has the WCR of R^s .*

Proof. We call a project (u, v) *top-tier* if $v \geq 1 - R^s$ and *bottom-tier* if $v < 1 - R^s$.

Under α^s , the agent's expected payoff is \underline{u} if he proposes a bottom-tier project, and is $u \geq \underline{u}$ if he proposes a top-tier project (u, v) . Therefore, if the agent has some top-tier project, it is optimal to propose his favorite top-tier project (u, v) , in which case the regret is at most $1 - v \leq R^s$. If all the projects in A are bottom-tier projects, it is optimal to propose a project (u', v') that maximizes $\alpha^s(u, v)v$. Let $(u_p, v_p) \in \operatorname{argmax}_{(u, v) \in A} v$ be a principal's favorite project in A . Since this is a bottom-tier project, it follows from the definition of R^s in (1) that $(1 - \underline{u})v_p \leq R^s$. Thus, the regret is $v_p - \alpha^s(u', v')v' \leq v_p - \alpha^s(u_p, v_p)v_p \leq (1 - \underline{u})v_p \leq R^s$. \square

Claim 7.2. *If α implements a choice function that has the WCR of R^s , then $\alpha(u, v) \leq \alpha^s(u, v)$ for every $(u, v) \in D$.*

Proof. Fix a project (u, v) . If $v \geq 1 - R^s$ or $u = 0$, then $\alpha^s(u, v) = 1$ and therefore $\alpha(u, v) \leq \alpha^s(u, v)$. If $v < 1 - R^s$ and $u > 0$, then since the WCR under α is R^s , it must be the case that if $A = \{(u, v), (\underline{u}, 1)\}$, then the agent proposes the project $(\underline{u}, 1)$. Otherwise, the regret is at least $1 - v > R^s$. Therefore $\alpha(u, v)u \leq \alpha(\underline{u}, 1)\underline{u} \leq \underline{u}$, which implies $\alpha(u, v) \leq \underline{u}/u = \alpha^s(u, v)$, as desired. \square

Claim 7.3. *The choice function under α^s in statement (ii) is admissible.*

Proof. Suppose that α implements a choice function that has the WCR of R^s . By Claim 7.2, $\alpha(u, v) \leq \alpha^s(u, v)$ for every $(u, v) \in D$. If $\alpha(u, v) < \alpha^s(u, v)$ for some top-tier project (u, v) , then if the agent's type is $\{(u, v)\}$, α does strictly worse than α^s for the principal. Now suppose that $\alpha(u, v) = \alpha^s(u, v)$ for every top-tier project (u, v) . Then for every type A with at least one top-tier project, α and α^s implement the same project choice, which is the agent's favorite top-tier project in A . Hence, we only need to consider those A with only bottom-tier projects. Since $\alpha(u, v) \leq \alpha^s(u, v)$ and the agent proposes a project in A to

maximize the principal's expected payoff $\alpha^s(u, v)v$ under the choice function in (ii), α cannot do better than the choice function in (ii) for every type A with only bottom-tier projects.

□

7.2 Proof of Theorem 4.2

Let $a^* = (\underline{u}, 1)$. Let $\bar{U}(P)$ be the optimal value of the following linear programming with variables $\pi(u, v)$ for every $(u, v) \in P$:

$$\bar{U}(P) = \max_{\pi} \quad \underline{u} + \sum_{(u,v) \in P} \pi(u, v)(u - \underline{u}) \quad (5a)$$

$$\text{s.t.} \quad \pi(u, v) \geq 0, \quad \forall (u, v) \in P, \quad (5b)$$

$$\sum_{(u,v) \in P} \pi(u, v) \leq 1, \quad (5c)$$

$$\sum_{(u,v) \in P} \pi(u, v)(1 - v) \leq R^m. \quad (5d)$$

The following claim explains the role of $\bar{U}(P)$ in our argument: $\bar{U}(P)$ is the maximal payoff that the principal can give the agent for proposing P such that the principal can give the agent this same payoff if the agent proposed $P \cup \{a^*\}$, while still keeping regret below R^m .

Claim 7.4. *If ρ is an IC mechanism which has the WCR of at most R^m , then $U(\rho, P) \leq \bar{U}(P)$ for every proposal P .*

Proof. Let $\tilde{P} = P \cup \{a^*\}$. Let $\pi = \rho(\cdot | \tilde{P})$. Since the regret under the mechanism ρ when the set of available projects is \tilde{P} is at most R^m , it follows that $\sum_{(u,v) \in P} \pi(u, v)(1 - v) \leq R^m$. Therefore the restriction of π to the set P is a feasible point in problem (5). Moreover

$$U(\rho, \tilde{P}) = \pi(a^*)\underline{u} + \sum_{(u,v) \in P} \pi(u, v)u \leq \underline{u} + \sum_{(u,v) \in P} \pi(u, v)(u - \underline{u}), \quad (6)$$

where the inequality follows from $\pi(a^*) + \sum_{(u,v) \in P} \pi(u, v) \leq 1$. The right hand side of (6) is

the objective function of (5) at π . Therefore, $U(\rho, \tilde{P}) \leq \bar{U}(P)$. Finally, since the mechanism ρ is IC, it follows that $U(\rho, P) \leq U(\rho, \tilde{P})$. Therefore, $U(\rho, P) \leq \bar{U}(P)$, as desired. \square

When P is a singleton $\{(u, v)\}$, we also denote $\bar{U}(\{(u, v)\})$ by $\bar{U}(u, v)$. The following claim explains the role of the function $\alpha^m(u, v)$ in our argument: This is the highest possible approval probability for (u, v) when it is proposed alone, such that the principal can incentivize the agent to propose both available projects when his type is $A = \{(u, v), (\underline{u}, 1)\}$, while keeping the regret from type A below R^m .

Claim 7.5. *When $P = \{(u, v)\}$ is a singleton, $\bar{U}(u, v) = \alpha^m(u, v)u$.*

Proof. We first prove that $\bar{U}(u, v) \leq \alpha^m(u, v)u$. Let $\pi = \pi(u, v)$ be a feasible point in problem (5) with $P = \{(u, v)\}$, so that $\pi(1 - v) \leq R^m$. Let $p \in [0, 1]$ be such that $pu = \underline{u} + \pi(u - \underline{u})$. Then $\gamma(u, p) \leq \pi$ by the definition of γ in (2). Therefore $\gamma(u, p)(1 - v) \leq \pi(1 - v) \leq R^m$. Therefore $\alpha^m(u, v) \geq p$ by the definition of α^m in (4). Therefore, $\underline{u} + \pi(u - \underline{u}) = pu \leq \alpha^m(u, v)u$. Since this holds for every feasible π , it follows that $\bar{U}(u, v) \leq \alpha^m(u, v)u$, as desired.

We now prove that $\bar{U}(u, v) \geq \alpha^m(u, v)u$. Let p be such that $\gamma(u, p)(1 - v) \leq R^m$ and let $\pi = \gamma(u, p)$. Then π is feasible in problem (5) and therefore $\bar{U}(u, v) \geq \underline{u} + \pi(u - \underline{u}) \geq pu$ by the definition of γ in (2). Since this holds for every p such that $\gamma(u, p)(1 - v) \leq R^m$, it follows by the definition of α^m in (4) that $\bar{U}(u, v) \geq \alpha^m(u, v)u$, as desired. \square

For a proposal P , let $\underline{U}(P) = \max_{(u, v) \in P} \alpha^m(u, v)u$. The following claim explains the role of $\underline{U}(P)$ in our argument.

Claim 7.6. *If ρ is an IC mechanism that approves the singleton proposal $\{(u, v)\}$ with probability $\alpha^m(u, v)$, then $U(\rho, P) \geq \underline{U}(P)$.*

Proof. Since ρ is IC, $U(\rho, P) \geq U(\rho, \{(u, v)\}) = \alpha^m(u, v)u$ for every $(u, v) \in P$. \square

Claim 7.4 bounds from above the agent's expected payoff in an IC mechanism which has the WCR of at most R^m . Claim 7.6 bounds from below the agent's expected payoff in an IC mechanism which approves the singleton proposal $\{(u, v)\}$ with probability $\alpha^m(u, v)$. The following claim shows that the definition of R^m is such that both bounds can be satisfied.

Claim 7.7. $\underline{U}(P) \leq \overline{U}(P)$ for every P .

Proof. The function $\overline{U}(P)$ defined in (5) is increasing in P . Therefore, from Claim 7.5 we have:

$$\alpha^m(u, v)u = \overline{U}(u, v) \leq \overline{U}(P), \quad \forall (u, v) \in P.$$

It follows that:

$$\underline{U}(P) = \max_{(u, v) \in P} \alpha^m(u, v)u \leq \overline{U}(P).$$

□

By definition, $\rho^m(\cdot|P)$ solves the following linear programming:

$$\max_{\pi} \quad \sum_{(u, v) \in P} \pi(u, v)v \tag{7a}$$

$$\text{s.t.} \quad \pi(u, v) \geq 0, \quad \forall (u, v) \in P, \tag{7b}$$

$$\sum_{(u, v) \in P} \pi(u, v) \leq 1, \tag{7c}$$

$$\sum_{(u, v) \in P} \pi(u, v)u = \underline{U}(P). \tag{7d}$$

It is possible that (7) has multiple optimal solutions. Since all the optimal solutions are payoff-equivalent for both the principal and the agent, we do not distinguish among them. From now on, the notation $\rho(\cdot|P) \neq \rho^m(\cdot|P)$ means that $\rho(\cdot|P)$ is not among the optimal solutions to (7). We now show that, when there is only one available project, the regret under ρ^m is at most R^m .

Claim 7.8. For every singleton $A = \{(u, v)\}$, the regret under ρ^m is at most R^m .

Proof. Since ρ^m approves $\{(u, v)\}$ with probability $\alpha^m(u, v)$, the regret is $v(1 - \alpha^m(u, v))$. By the definition of R^m , there exists some $\bar{p} \in [0, 1]$ such that $\max\{v(1 - \bar{p}), \gamma(u, \bar{p})(1 - v)\} \leq R^m$. By (4), $\bar{p} \leq \alpha^m(u, v)$. Therefore, it also follows that $v(1 - \alpha^m(u, v)) \leq v(1 - \bar{p}) \leq R^m$. \square

Claim 7.9. The optimal value in problem (7) is at least $\max_{(u,v) \in P} v - R^m$.

Proof. It is sufficient to prove that the optimal values in each of the two problems derived from (7) by replacing the equality constraint (7d) with inequalities \leq and \geq are at least $\max_{(u,v) \in P} v - R^m$.

We first consider problem (7) with (7d) being replaced by:

$$\sum_{(u,v) \in P} \pi(u, v)u \leq \underline{U}(P). \quad (8)$$

Let $(u_p, v_p) \in P$ be a principal's favorite project. Let π be the subprobability over P that is given by $\pi(u, v) = \alpha^m(u_p, v_p)\delta_{(u_p, v_p), (u, v)}$ where δ is Kronecker notation so

$$\delta_{(u_p, v_p), (u, v)} = \begin{cases} 1, & \text{if } (u, v) = (u_p, v_p), \\ 0, & \text{otherwise.} \end{cases}$$

Then $\sum_{(u,v) \in P} \pi(u, v)u = \alpha^m(u_p, v_p)u_p \leq \underline{U}(P)$ by the definition of $\underline{U}(P)$, so π satisfies (8). Also $v_p(1 - \alpha^m(u_p, v_p)) \leq R^m$ by Claim 7.8. Therefore, the value of the objective function in (7) at π is at least $v_p - R^m$, as desired.

We then consider problem (7) with (7d) being replaced by:

$$\sum_{(u,v) \in P} \pi(u, v)u \geq \underline{U}(P). \quad (9)$$

Let $(u_a, v_a) \in P$ be such that $(u_a, v_a) \in \operatorname{argmax}_{(u,v) \in P} \alpha^m(u, v)u$. Let π be the probability

distribution over P such that

$$\pi(u, v) = \gamma(u_a, \alpha^m(u_a, v_a))\delta_{(u_a, v_a), (u, v)} + (1 - \gamma(u_a, \alpha^m(u_a, v_a)))\delta_{(u_p, v_p), (u, v)}.$$

Then

$$\sum_{(u, v) \in P} \pi(u, v)u = \gamma(u_a, \alpha^m(u_a, v_a))u_a + (1 - \gamma(u_a, \alpha^m(u_a, v_a)))u_p \geq \alpha^m(u_a, v_a)u_a = \underline{U}(P),$$

where the inequality follows from $u_p \geq \underline{u}$ and (2), and the last equality follows from the definitions of $\underline{U}(P)$ and (u_a, v_a) . Hence, π satisfies (9). Also

$$v_p - \sum_{(u, v) \in P} \pi(u, v)v = \gamma(u_a, \alpha^m(u_a, v_a))(v_p - v_a) \leq \gamma(u_a, \alpha^m(u_a, v_a))(1 - v_a) \leq R^m,$$

where the last inequality follows from the definition of α^m . Therefore, the value of the objective function in (7) at π is at least $v_p - R^m$, as desired. \square

The following lemma gives an equivalent characterization of the mechanism ρ^m .

Lemma 7.10. *The optimal solutions to (7) and those to the following problem coincide.*

$$\max_{\pi} \sum_{(u, v) \in P} \pi(u, v)v \tag{10a}$$

$$\text{s.t. } \pi(u, v) \geq 0, \forall (u, v) \in P, \tag{10b}$$

$$\sum_{(u, v) \in P} \pi(u, v) \leq 1, \tag{10c}$$

$$\sum_{(u, v) \in P} \pi(u, v)u \geq \underline{U}(P), \tag{10d}$$

$$\sum_{(u, v) \in P} \pi(u, v)u \leq \bar{U}(P). \tag{10e}$$

Proof of Lemma 7.10. We discuss two cases separately.

Case 1. Assume that there exists some $(u, v) \in P$ such that $v \geq 1 - R^m$. Consider the following linear programming which is a relaxation of both problem (7) and problem (10):

$$\max_{\pi} \sum_{(u,v) \in P} \pi(u, v)v \quad (11a)$$

$$\text{s.t. } \pi(u, v) \geq 0, \forall (u, v) \in P, \quad (11b)$$

$$\sum_{(u,v) \in P} \pi(u, v) \leq 1, \quad (11c)$$

$$\sum_{(u,v) \in P} \pi(u, v)u \geq \underline{U}(P). \quad (11d)$$

We claim that the constraint (11d) holds with equality at every optimal solution. Indeed, if (11d) is not binding then an optimal solution to (11) is also an optimal solution to the following linear programming:

$$\max_{\pi} \sum_{(u,v) \in P} \pi(u, v)v \quad (12a)$$

$$\text{s.t. } \pi(u, v) \geq 0, \forall (u, v) \in P, \quad (12b)$$

$$\sum_{(u,v) \in P} \pi(u, v) \leq 1, \quad (12c)$$

which is derived from (11) by removing (11d). Let $v_p = \max_{(u,v) \in P} v$ and $u_p = \max_{(u,v_p) \in P} u$. Since $v_p \geq 1 - R^m$, it follows from the definition of α^m in (4) that $\alpha^m(u_p, v_p) = 1$. Every optimal solution π to problem (12) satisfies $\text{support}(\pi) \subseteq \text{argmax}_{(u,v) \in P} v$, which implies that:

$$\sum_{(u,v) \in P} \pi(u, v)u \leq u_p = \alpha^m(u_p, v_p)u_p \leq \underline{U}(P).$$

We proved that every optimal solution to (11) satisfies (11d) with equality, so it is a feasible

point in both (7) and (10). Since problem (11) is a relaxation of both problem (7) and (10), the optimal values of (7), (10), and (11) coincide.

Case 2. Assume now that $v < 1 - R^m$ for every $(u, v) \in P$. We claim that $\underline{U}(P) = \overline{U}(P)$ and therefore problems (7) and (10) coincide. Given that $v < 1 - R^m$ for every $(u, v) \in P$, the constraint (5c) in problem (5) must be slack since if it is satisfied with an equality then (5d) is violated. Therefore, in this case $\overline{U}(P)$ also satisfies

$$\begin{aligned} \overline{U}(P) = \max_{\pi} \quad & \underline{u} + \sum_{(u,v) \in P} \pi(u, v)(u - \underline{u}) \\ \text{s.t.} \quad & \pi(u, v) \geq 0, \forall (u, v) \in P, \\ & \sum_{(u,v) \in P} \pi(u, v)(1 - v) \leq R^m, \end{aligned} \tag{13}$$

which is derived from problem (5) by removing (5c). Problem (13) admits a solution π^* with the property that, for some $(u^*, v^*) \in P$, the only non-zero element of π^* is $\pi^*(u^*, v^*)$. Therefore, by Claim 7.5,

$$\overline{U}(P) = \overline{U}(u^*, v^*) = \alpha^m(u^*, v^*)u^* \leq \underline{U}(P).$$

Therefore, by Claim 7.7 we get $\overline{U}(P) = \underline{U}(P)$, as desired.

□

Proof of Theorem 4.2. 1. Fix $(u, v) \in D$ and let $P = \{(u, v)\}$ and $\tilde{P} = \{(u, v), (\underline{u}, 1)\}$.

Let p be the probability that ρ chooses (u, v) when the proposal is P . So, $\text{RGRT}(\rho, P) = (1 - p)v$. Since the mechanism is IC, the agent's expected payoff under \tilde{P} must be at least pu . By definition of $\gamma(u, p)$, this implies that when the proposal is \tilde{P} the mechanism chooses (u, v) with probability at least $\gamma(u, p)$. So, $\text{RGRT}(\rho, \tilde{P}) \geq \gamma(u, p)(1 - v)$. Therefore, the principal's WCR is at least $\max\{(1 - p)v, \gamma(u, p)(1 - v)\}$. Since the

principal can choose p , it follows that

$$\text{WCR}(\rho) \geq \min_{p \in [0,1]} \max \{(1-p)v, \gamma(u,p)(1-v)\}.$$

This inequality must hold for any $(u, v) \in D$. Hence, the WCR under any mechanism is at least R^m .

2. The mechanism ρ^m is IC, and it solves problem (7). By Claim 7.9, the optimal value in problem (7) is at least $\max_{(u,v) \in P} v - R^m$. Since the objective function in (7) is the principal's payoff under ρ^m , the principal's regret is at most R^m .

We next argue that ρ^m is admissible. Let ρ be an IC mechanism which has the WCR of R^m and let $\alpha(u, v)$ be the probability that ρ approves a singleton proposal $\{(u, v)\}$. By Claims 7.4 and 7.5, $\alpha(u, v) \leq \alpha^m(u, v)$. Assume first that $\alpha(u, v) < \alpha^m(u, v)$ for some project (u, v) . Then, since $v > 0$, if the agent's type is the singleton $\{(u, v)\}$ the principal's payoff is strictly higher under ρ^m than under ρ . Assume now that $\alpha(u, v) = \alpha^m(u, v)$ for every (u, v) . We claim that for every proposal P the principal's payoff is weakly higher under ρ^m than under ρ . Indeed, since ρ is IC, it follows from Claim 7.4 that $U(\rho, P) \leq \overline{U}(P)$, and, from Claim 7.6, that $U(\rho, P) \geq \underline{U}(P)$. Therefore $\rho(\cdot|P)$ is a feasible point in problem (10). Since by Lemma 7.10 $\rho^m(\cdot|P)$ is the optimal solution to (10), the principal's payoff is weakly higher under ρ^m than under ρ .

3. Let ρ be an IC, admissible mechanism which has the WCR of R^m and which differs from ρ . We want to show that $U(\rho, P) \leq U(\rho^m, P)$ for every finite $P \subseteq D$. Recall that $U(\rho^m, P) = \underline{U}(P)$ for every P .

We first construct a new mechanism $\tilde{\rho}$ based on ρ and ρ^m :

$$\tilde{\rho}(\cdot|P) = \begin{cases} \rho^m(\cdot|P), & \text{if } U(\rho, P) \geq \underline{U}(P) \\ \rho(\cdot|P), & \text{if } U(\rho, P) < \underline{U}(P). \end{cases}$$

By definition, $U(\tilde{\rho}, P) = \min \{U(\rho, P), U(\rho^m, P)\}$. The functions $U(\rho, P)$ and $U(\rho^m, P)$ are increasing in P since ρ and ρ^m are IC. Therefore $U(\tilde{\rho}, P)$ is increasing in P , so $\tilde{\rho}$ is also IC. Moreover, for every P either $\tilde{\rho}(\cdot|P) = \rho(\cdot|P)$ or $\tilde{\rho}(\cdot|P) = \rho^m(\cdot|P)$. Therefore the WCR under $\tilde{\rho}$ is also R^m .

We next argue that for every P , $\tilde{\rho}$ gives the principal a weakly higher payoff than ρ does.

- (a) Consider a set P such that $U(\rho, P) < \underline{U}(P)$. Then $\tilde{\rho}(\cdot|P) = \rho(\cdot|P)$, so $\tilde{\rho}$ gives the principal the same payoff as ρ does.
- (b) Consider a set P such that $U(\rho, P) \geq \underline{U}(P)$. From Claim 7.4 we know that $U(\rho, P) \leq \overline{U}(P)$ for every P . Therefore, $\rho(\cdot|P)$ is a feasible point in problem (10). It follows from Lemma 7.10 that ρ^m gives the principal a weakly higher payoff than ρ does. Moreover, if $\rho(\cdot|P) \neq \rho^m(\cdot|P)$, then ρ^m gives the principal a strictly higher payoff than ρ does.

Since $\tilde{\rho}(\cdot|P) = \rho^m(\cdot|P)$ for every P such that $U(\rho, P) \geq \underline{U}(P)$, $\tilde{\rho}$ gives the principal a weakly higher payoff than ρ does for every such P .

We have argued that $\tilde{\rho}$ gives the principal a weakly higher payoff than ρ does for every P . On the other hand, ρ is admissible, so there cannot be a P such that $\tilde{\rho}$ gives the principal a strictly higher payoff than ρ does. This implies that $\rho(\cdot|P) = \rho^m(\cdot|P)$ for every P such that $U(\rho, P) \geq \underline{U}(P)$, so $U(\rho, P) = U(\rho^m, P) = \underline{U}(P)$. Hence $U(\rho, P) \leq \underline{U}(P)$ for every P .

□

7.3 Partial commitment in the multiproject environment

Let $v_p(P) = \max_{(u,v) \in P} v$ be the principal's payoff from his favorite project in P . We consider an environment in which the principal cannot commit to choosing a worse project when a better project is also proposed. This means that the mechanism ρ must satisfy $\rho((u, v)|P) =$

0 for any $(u, v) \in P$ such that $v \neq v_p(P)$. We now show that this environment is essentially equivalent to the single-project environment.

Let $u_p(P) = \min_{(u,v) \in P, v=v_p(P)} u$ be the agent's lowest payoff among the principal's favorite projects in P . We now define a mechanism ρ^* which is adapted from the optimal mechanism α^s in the single-project environment. The mechanism ρ^* chooses $(u_p(P), v_p(P))$ with probability $\alpha^s(u_p(P), v_p(P))$ and chooses other projects with zero probability:

$$\rho^*((u, v)|P) = \begin{cases} \alpha^s(u_p(P), v_p(P)) & \text{if } (u, v) = (u_p(P), v_p(P)); \\ 0 & \text{if } (u, v) \neq (u_p(P), v_p(P)). \end{cases}$$

For a mechanism ρ , we let $\rho((\cdot, v_p(P))|P)$ denote the probability that a principal's favorite project in P is chosen when the agent proposes P .

Theorem 7.11. *Consider the multiproject environment with partial commitment.*

- (i) *The WCR under any mechanism ρ is at least R^s .*
- (ii) *Under ρ^* , if the agent has projects with $v \geq 1 - R^s$, it is optimal for him to propose his favorite project among those with $v \geq 1 - R^s$. Otherwise, it is optimal to propose a project that maximizes the principal's expected payoff $\alpha^s(u, v)v$. The corresponding choice function has the WCR of R^s and is admissible.*
- (iii) *If a mechanism ρ implements a choice function that has the WCR of R^s , then*

$$\rho((\cdot, v_p(P))|P) \leq \rho^*((\cdot, v_p(P))|P), \text{ for every } P.$$

Proof. (i) Fix $v \in [\underline{v}, 1)$. Let $x = \rho((1, v)|\{(1, v)\})$ be the probability that ρ chooses $(1, v)$ if the agent proposes only one project $(1, v)$.

If $x \leq \underline{u}$, then, if the agent has only the project $(1, v)$, the regret is $(1 - x)v \geq (1 - \underline{u})v$. Now suppose that $x > \underline{u}$. If the agent has two projects $(1, v)$ and $(\underline{u}, 1)$, his payoff is at most \underline{u} if he proposes both projects. This is because ρ must assign zero probability to

$(1, v)$ when $(\underline{u}, 1)$ is also proposed. Moreover, his payoff is at most \underline{u} if he proposes only $(\underline{u}, 1)$. On the other hand, his payoff is strictly higher than \underline{u} if he proposes only $(1, v)$ since $x > \underline{u}$. Hence, the agent proposes only $(1, v)$ and the regret is $1 - \alpha(1, v)v \geq 1 - v$. Therefore, $\text{WCR} \geq \min \{(1 - \underline{u})v, 1 - v\}$ for every $v \in [\underline{v}, 1)$.

- (ii) We call a project (u, v) *top-tier* if $v \geq 1 - R^s$ and *bottom-tier* if $v < 1 - R^s$. From the definition of R^s it follows that $(1 - \underline{u})v \leq R^s$ for every bottom-tier project.

Under ρ^* , the agent's expected payoff is \underline{u} if all of the proposed projects are bottom-tier, and is at least \underline{u} if at least one of the proposed projects is top-tier. Therefore, if the agent has some top-tier project, it is optimal to propose the agent's favorite top-tier project (u, v) , in which case the regret is at most $1 - v \leq R^s$. If all the projects in A are bottom-tier, it is optimal to propose a project (u', v') that maximizes $\alpha^s(u, v)v$. Let $(u_p, v_p) \in \operatorname{argmax}_{(u, v) \in A} v$ be a principal's favorite project. Thus, the regret is $v_p - \alpha^s(u', v')v' \leq v_p - \alpha^s(u_p, v_p)v_p \leq (1 - \underline{u})v_p \leq R^s$.

We prove that this choice function is admissible in part (iii).

- (iii) If P includes a top-tier project or $u_p(P) = 0$, then $\rho((\cdot, v_p(P))|P) \leq 1 = \rho^*((\cdot, v_p(P))|P)$. If all projects in P are bottom-tier and $u_p(P) > 0$, we claim that $\rho((\cdot, v_p(P))|P) \leq \underline{u}/u_p(P) = \rho^*((\cdot, v_p(P))|P)$. Suppose otherwise that $\rho((\cdot, v_p(P))|P) > \underline{u}/u_p(P)$ for some P . If the agent's type is $P \cup \{(\underline{u}, 1)\}$, then he will not include $(\underline{u}, 1)$ in his proposal, since proposing just P leads to a strictly higher payoff than \underline{u} . As a result, the regret is strictly higher than R^s , contradicting the fact that the WCR under ρ is at most R^s .

Lastly, we show that the choice function in statement (ii) is admissible. Fix a mechanism ρ which has the WCR of R^s . If the agent proposes only bottom-tier projects, his expected payoff cannot exceed \underline{u} . (Otherwise, the agent has incentive to hide $(\underline{u}, 1)$, causing a higher regret than R^s .)

Consider the set of all single-project proposals with one top-tier project. If there exists one such proposal P such that ρ does not choose the proposed project with probability 1, then if the agent's type is actually P , ρ does strictly worse than ρ^* for the principal. Now suppose that, for every single-project proposal with one top-tier project, ρ chooses the proposed project with probability 1. Then for every A with at least one top-tier project, ρ and ρ^* implement the same project choice, which is the agent's favorite top-tier project in A . Hence, we only need to consider those A with only bottom-tier projects. Since $\rho((\cdot, v_p(P))|P) \leq \rho^*((\cdot, v_p(P))|P)$, the principal's payoff is higher under ρ^* than under ρ for every $P \subset A$. Moreover, under the choice function in (ii), the agent chooses a project in A to maximize the principal's expected payoff, so ρ cannot do better than the choice function in (ii) for every A with only bottom-tier projects. □

References

- Alonso, Ricardo, and Niko Matouschek.** 2008. "Optimal Delegation." *The Review of Economic Studies*, 75(1): 259–293.
- Amador, Manuel, and Kyle Bagwell.** 2013. "The Theory of Optimal Delegation With an Application to Tariff Caps." *Econometrica*, 81(4): 1541–1599.
- Armstrong, Mark, and John Vickers.** 2010. "A Model of Delegated Project Choice." *Econometrica*, 78(1): 213–244.
- Bergemann, Dirk, and Karl H. Schlag.** 2008. "Pricing without Priors." *Journal of the European Economic Association*, 6(2/3): 560–569.
- Bergemann, Dirk, and Karl Schlag.** 2011. "Robust Monopoly Pricing." *Journal of Economic Theory*, 146(6): 2527–2543.

- Berkovitch, Elazar, and Ronen Israel.** 2004. “Why the NPV Criterion does not Maximize NPV.” *The Review of Financial Studies*, 17(1): 239–255.
- Beviá, Carmen, and Luis Corchón.** 2019. “Contests with Dominant Strategies.” *Economic Theory*, 1–19.
- Bull, Jesse, and Joel Watson.** 2007. “Hard Evidence and Mechanism Design.” *Games and Economic Behavior*, 58(1): 75 – 93.
- Byrne, John A.** 2018. “Two Wharton Profs Among Finalists For Haas Deanship.” <https://poetsandquants.com/2018/04/10/two-wharton-profs-among-finalists-for-haas-deanship/>.
- Byrne, John A.** 2021. “HEC Paris: New Dean Will Put Focus On Impact.” <https://poetsandquants.com/2021/02/17/hec-paris-new-dean-will-put-focus-on-impact/?pq-category=business-school-news>.
- Carroll, Gabriel.** 2019. “Robustness in Mechanism Design and Contracting.” *Annual Review of Economics*, 11(1): 139–166.
- Che, Yeon-Koo, Wouter Dessein, and Navin Kartik.** 2013. “Pandering to Persuade.” *American Economic Review*, 103(1): 47–79.
- Gilboa, Itzhak, and David Schmeidler.** 1989. “Maxmin Expected Utility with Non-unique Prior.” *Journal of Mathematical Economics*, 18(2): 141–153.
- Goel, Sumit, and Wade Hann-Caruthers.** 2020. “Project Selection with Partially Verifiable Information.” *Working paper*.
- Green, Jerry R., and Jean-Jacques Laffont.** 1986. “Partially Verifiable Information and Mechanism Design.” *The Review of Economic Studies*, 53(3): 447–456.
- Guo, Yingni, and Eran Shmaya.** 2023. “Robust Monopoly Regulation.” *Working paper*.

- Harris, Milton, and Artur Raviv.** 1996. “The Capital Budgeting Process: Incentives and Information.” *The Journal of Finance*, 51(4): 1139–1174.
- Holmström, Bengt.** 1984. “On the Theory of Delegation.” *in: Bayesian Models in Economic Theory*. Ed. by M. Boyer, and R. Kihlstrom. North-Holland, Amsterdam, 115–141.
- Hurwicz, Leonid, and Leonard Shapiro.** 1978. “Incentive Structures Maximizing Residual Gain under Incomplete Information.” *The Bell Journal of Economics*, 9(1): 180–191.
- Lipman, Barton L, and Duane J Seppi.** 1995. “Robust Inference in Communication Games with Partial Provability.” *Journal of Economic Theory*, 66(2): 370–405.
- Lyons, Bruce R.** 2003. “Could Politicians be More Right Than Economists? A Theory of Merger Standards.” *Working paper*.
- Malladi, Suraj.** 2022. “Delegated Costly Screening.” *Working paper*.
- Martimort, David, and Aggey Semenov.** 2006. “Continuity in Mechanism Design without Transfers.” *Economics Letters*, 93(2): 182–189.
- Melumad, Nahum D., and Toshiyuki Shibano.** 1991. “Communication in Settings with No Transfers.” *The RAND Journal of Economics*, 22(2): 173–198.
- Milnor, John.** 1954. “Games against Nature.” In *Decision Processes*. , ed. Robert Thrall, Clyde Coombs and Robert Davis. Wiley.
- Neven, Damien J, and Lars-Hendrik Röller.** 2005. “Consumer Surplus vs. Welfare Standard in a Political Economy Model of Merger Control.” *International Journal of Industrial Organization*, 23(9-10): 829–848.
- Nocke, Volker, and Michael D. Whinston.** 2013. “Merger Policy with Merger Choice.” *American Economic Review*, 103(2): 1006–33.
- Ottaviani, Marco, and Abraham L Wickelgren.** 2011. “Ex Ante or Ex Post Competition Policy? A Progress Report.” *International Journal of Industrial Organization*, 29(3): 356–359.

- Poterba, James M, and Lawrence H Summers.** 1995. "A CEO Survey of US Companies' Time Horizons and Hurdle Rates." *MIT Sloan Management Review*, 37(1): 43–53.
- Renou, Ludovic, and Karl H. Schlag.** 2011. "Implementation in Minimax Regret Equilibrium." *Games and Economic Behavior*, 71(2): 527 – 533.
- Ross, Marc.** 1986. "Capital Budgeting Practices of Twelve Large Manufacturers." *Financial Management*, 15–22.
- Savage, Leonard J.** 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association*, 46(253): 55–67.
- Savage, Leonard J.** 1972. *The Foundations of Statistics*. Courier Corporation.
- Stanley, Marjorie T, and Stanley B Block.** 1984. "A Survey of Multinational Capital Budgeting." *Financial Review*, 19(1): 36–54.
- Stoye, Jörg.** 2011. "Axioms for Minimax Regret Choice Correspondences." *Journal of Economic Theory*, 146(6): 2226–2251.
- Wald, Abraham.** 1950. *Statistical Decision Functions*. Wiley.