Project Choice from a Verifiable Proposal

Yingni Guo Eran Shmaya*

February 27, 2021

Abstract

An agent, observing the set of available projects, can propose only available projects, 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. If the agent can propose only one project, it is chosen for sure if the principal's payoff exceeds a threshold; otherwise, the probability that it is chosen decreases in the agent's payoff. If the agent can propose multiple projects, his payoff from a multiproject proposal equals the maximal payoff from proposing each project alone. Our results highlight the benefits from randomization and the ability to propose multiple projects.

JEL: D81, D82, D86

Keywords: verifiable disclosure, evidence, project choice, regret minimization

1 Introduction

Project choice is one of the most important functions of an organization. The process often involves two parties: (i) a party at a lower hierarchical level who has expertise and proposes

*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, and the University of Bonn for valuable feedback.

projects, and (ii) a part at a higher hierarchical level who evaluates the proposed projects and makes the choice. This describes the relationship between a division and the headquarters when the division has a chance to build a factory or to choose an office building. It also applies to the relationship between a department and the university 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 chooses one from the proposed projects or rejects them all. If the two parties had identical preferences over projects, the agent would propose the project that is their shared favorite among the available ones, and the principal would always automatically approve the agent's proposal. In many applications, however, the two parties do not share the same preferences. For instance, the division may fail to internalize each project's externalities on other divisions; the department and the university may put different weights on candidates' research and nonresearch abilities. 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. How shall the principal encourage the agent to propose the principal's preferred projects? What is the principal's optimal mechanism for choosing a project?

It is easy to see that no mechanism can guarantee that the principal's favorite project among the available ones will always 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. We look for a mechanism that works fairly well for the principal in all circumstances, i.e., a mechanism that minimizes the principal's worst-case regret.

Depending on the principal's verification capacity, we distinguish 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. Besides project choice within organizations, the single-project environment also applies to antitrust regulation: a firm proposes a merger or acquisition transaction and the regulator decides whether to approve or reject it (e.g., Armstrong and Vickers (2010), Nocke and Whinston (2013)). We take the environment as exogenous and derive the optimal mechanisms in both environments. Comparing the two environments not only delineates how the principal benefits from higher verification capacity but also quantifies his gain.

We begin with the single-project environment. A mechanism specifies for each proposed single project the probability that it will be approved. In the optimal mechanism, if the proposed project gives the principal a sufficiently high payoff, it is approved for sure. If the proposed project is mediocre for the principal, it is approved only with some probability. The probability that a mediocre 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.

In the multiproject environment, a mechanism specifies for each proposed set of projects a randomization over the proposed projects and "no project." If the agent proposes only one project, the optimal mechanism takes a form similar to the one in the single-project environment. In particular, if the proposed project is sufficiently good for the principal, it is chosen for sure. Otherwise, the project is chosen with some probability that decreases in its payoff to the agent. If the agent proposes more than one project, the randomization maximizes the principal's expected payoff, subject to the constraint that the agent's expected payoff is the maximal expected payoff he would get from proposing each project alone. Under this mechanism, the more projects the agent proposes, the weakly higher his expected payoff is, so the agent is willing to propose all available projects.

Since the agent gets the maximal expected payoff from proposing each project alone, we call this mechanism the *project-wide maximal-payoff mechanism*. This mechanism implements a compromise between the two parties in the multiproject environment: with some probability the choice favors the agent and with some probability it favors the principal. Randomization is crucial for the principal's minimal worst-case regret to be lower in the multiproject environment than in the single-project one. We show in section 4.1 that if the principal is restricted to deterministic mechanisms, his minimal worst-case regret is the same in both environments.

Related literature. Our paper is closely related to Armstrong and Vickers (2010) and Nocke and Whinston (2013), which examine the same project choice problem. They assume that the principal has a prior belief over the number and the characteristics of the available projects, and mainly focus on the optimal deterministic mechanism in the single-project environment. Goel and Hann-Caruthers (2020) considers the project choice problem where the number of available projects is public information. The projects are only partially verifiable, since the agent's only constraint is not to overreport projects' payoffs to the principal. Because their agent cannot hide projects like our agent does, the resulting incentive schemes are quite different.

Our result relates to a theme in Aghion and Tirole (1997), namely, that the principal has formal authority, but the agent shares real authority due to his private information. We take this one step further. Our agent's real authority has two sources: he knows which projects are available, and he determines the proposal from which the principal chooses a project. The idea of striking a compromise is related to Bonatti and Rantakari (2016). They examine the compromise between two symmetric, competing agents whose efforts are crucial for discovering projects. We instead focus on the compromise between an agent who proposes projects and a principal who chooses among proposed projects.

Since in our model the agent can propose only those projects that are available, the agent's proposal is some evidence about his private information. From this aspect, our paper is related to research on verifiable disclosure (e.g., Grossman and Hart (1980), Grossman

(1981), Milgrom (1981), Dye (1985)) and, more broadly, the evidence literature (see Dekel (2016) for a survey). We discuss the relation to the evidence literature in section 2.1.

Our paper also contributes to the literature on mechanism design in which the designer minimizes his worst-case regret. 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. Guo and Shmaya (2019) study the optimal mechanism for monopoly regulation and Malladi (2020) studies the optimal approval rules for innovation. More broadly, we contribute to the growing literature of mechanism design with worst-case objectives. See, for instance, Chassang (2013) and Carroll (2015); for a survey, see Carroll (2019).

2 Model and mechanism

Let D be a set of *verifiable projects*. Let $u : D \to \mathbf{R}_+$ be the agent's payoff function, so his payoff is $u(a)$ if project a is chosen. If no project is chosen, the agent's payoff is zero.

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, $P \subseteq A$ so the agent can propose only available projects. This is what we meant earlier when we said that projects are verifiable. Second, $P \in \mathcal{E}$ for some fixed collection $\mathcal{E}$ of subsets of D. The set $\mathcal{E}$ captures exogenous restrictions on the proposal. For instance, the principal may have limited verification capacity or there may be some institutional restrictions on the mechanism. We are interested in two environments: *single-project* and *multiproject*. In the single-project environment, the agent can propose at most one available project, so

$\mathcal{E} = \{P \subseteq D : |P| \leqslant 1\}$. In the multiproject environment, the agent can propose any set of available projects so $\mathcal{E} = 2^D$, the power set of D.

The agent's proposal P serves two roles. First, if we view a proposal as a message, then different types have access to different messages. Hence, the agent's proposal is some evidence about his type, as in Green and Laffont (1986). We explore the implications of this role of the proposal in section 2.1. Second, the proposal determines the set of projects from which the principal can choose. This second role of the proposal is a key difference between our paper and the evidence literature.

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

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

is finite, and $\sum_a \pi(a) \leqslant 1$. 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 with probability $1 - \sum_a \pi(a)$ no project is chosen.

A *mechanism* 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 $\rho(\cdot|P)$. Thus, the agent's expected payoff if he proposes P under the mechanism ρ is $U(P, \rho) = \sum_{a \in P} u(a)\rho(a|P)$.

A *choice function* attaches to each type A of available projects 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 $f(\cdot|A)$.

A choice function f is *implemented* by a mechanism ρ if, for every type A of available projects, there exists a probability measure μ with support over

$$\operatorname*{argmax}_{P \subseteq A, P \in \mathcal{E}} U(P, \rho)$$

such that $f(a|A) = \sum_P \mu(P)\rho(a|P)$. The interpretation is that the agent selects only proposals that give him the highest expected payoff among the proposals that he can choose, and that the agent can randomize if he has multiple optimal proposals.

2.1 The evidence structure

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

In our multiproject environment, where $\mathcal{E} = 2^D$, the agent has the ability to provide the maximal evidence for his type. This property is also called *normality* (Lipman and Seppi (1995), Bull and Watson (2007), Ben-Porath, Dekel and Lipman (2019)). Another interpretation of the multiproject environment is to view an agent who proposes a set P as an agent who claims that his type is P. The relation that "type A can claim 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) and is also assumed in Hart, Kremer and Perry (2017).

In our single-project environment, where $\mathcal{E} = \{P \subseteq D : |P| \leqslant 1\}$, normality does not hold. The single-project environment is the main focus in Armstrong and Vickers (2010) and Nocke and Whinston (2013), and is similar to the assumption in Glazer and Rubinstein (2006) and Sher (2014) that the speaker can make one and only one of the statements he has access to.

2.2 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 his type, that is, $A \in \text{argmax}_{P \subseteq A, P \in \mathcal{E}} U(P, \rho)$ for any A. The following proposition states the revelation principle in the multiproject environment.

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

As we explained in section 2.1, the multiproject environment satisfies normality and the nested range condition. Previous papers (e.g., Green and Laffont (1986), Bull and Watson (2007)) showed that the revelation principle holds under these assumptions. Our proposition 2.1 does not follow directly from their theorems, however, because the proposal P serves two roles in our model. In addition to providing evidence, the agent's proposal also determines the set of projects from which the principal can choose. Nonetheless, a similar argument about the revelation principle can be made within our model as well.

Proof of Proposition 2.1. In order to show that f is IC, we need to show that $U(A, f) \geqslant U(P, f)$ for every A and every $P \subseteq A$.

By the definition that the choice function f is implemented by some mechanism ρ, it follows that $U(A, f) = \max_{P \subseteq A} U(P, \rho)$ for every P, since if the set of available projects is P, the agent can propose any $Q \subseteq P$. In particular, if $P \subseteq A$, we get that $U(A, f) \geqslant U(Q, \rho)$ for every $Q \subseteq P \subseteq A$. Therefore, $U(A, f) \geqslant \max_{Q \subseteq P} U(Q, \rho) = U(P, f)$, where the equality follows again from the fact that f is implemented by ρ. Hence, f is IC. Also, by definition, if f is IC, then f implements f. □

3 The principal's problem

Let $v : D \to \mathbf{R}_+$ be the principal's payoff, so his payoff is $v(a)$ if project a is chosen. If no project is chosen, the principal's payoff is zero.

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

$$\mathrm{RGRT}(A, f) = \max\{v(a) : a \in A\} - \sum_{a \in A} v(a) f(a|A).$$

The *worst-case regret* (WCR) from a choice function f is:

$$\mathrm{WCR}(f) = \sup\{\mathrm{RGRT}(A, f) : A\},$$

where the supremum ranges over all types (i.e., all finite sets $A \subseteq D$). The principal's problem is to minimize $\mathrm{WCR}(f)$ over all implementable choice functions f.

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

4 Optimal mechanisms

4.1 Preliminary intuition

We start with the preliminary intuition behind the optimal mechanisms. We first explain how randomization helps to reduce the WCR in the single-project environment. We then explain how the multiproject environment can further reduce the WCR. For this illustration, we assume that $\underline{v} = 0$.

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 concerns. First, if the agent has multiple projects which will be approved, then he will propose what he likes the most, even if projects are available that are more valuable to the principal. Second, if the agent has only projects which will be rejected, then the principal loses the payoff from these projects. Applied to the project $a^* = (1, 1/2)$, these two concerns imply that no matter how the principal designs the deterministic mechanism, his WCR is at least 1/2. This project a^* gives the agent his highest payoff 1, while giving the principal a moderate payoff 1/2. If the mechanism approves a^* and the set of available projects is $\{a^*, (\underline{u}, 1)\}$, then the agent will propose a^* rather than $(\underline{u}, 1)$, so the principal suffers regret 1/2. If the mechanism rejects a^* but a^* is 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 $v \geqslant 1/2$ achieves the WCR of 1/2, so it is optimal among 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 the only way to incentivize the agent to propose the project $(0, 1)$ when the set of available projects is $\{a^*, (0, 1)\}$ is still to reject the project a^* outright if a^* is proposed. However, if $\underline{u} > 0$, then the principal can do better. He can approve the project a^* with probability $\underline{u}$, while still maintaining the agent's incentive to propose the project $(\underline{u}, 1)$ when $\{a^*, (\underline{u}, 1)\}$ is the set of available projects. We carry out this idea in Theorem 4.1 below.

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 environment. (For 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 only one available project.) However, when 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 $\{a^*, (\underline{u}, 1)\}$, then the only way to incentivize the agent to include $(\underline{u}, 1)$ in his proposal is to reject the project a^* when 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 a^* or $(0,1)$ or nothing when the agent proposes $\{a^*, (0,1)\}$, a randomized mechanism can reach a compromise, in which each project is chosen with probability 1/2. Then, the principal can incentivize the agent to propose his type $\{a^*, (0,1)\}$ by choosing the project a^* with probability 1/2 if a^* is proposed alone. The regret is 1/4 both when the agent's type is $\{a^*, (0,1)\}$ and when his type is $\{a^*\}$. We carry out this idea of reaching a compromise in Theorem 4.2. When the agent proposes a set P, the principal gives the agent the highest payoff that he can give him, while maintaining the ability to offer the agent this same payoff if he proposes the set $P \cup \{(\underline{u}, 1)\}$ and keeping regret under control.

4.2 Optimal mechanism in the single-project environment

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

Theorem 4.1. *Assume $\mathcal{E} = \{P \subseteq D : |P| \leqslant 1\}$. Let*

$$R^s \equiv \max_{v \in [\underline{v},1]} \min\left((1-\underline{u})v, 1-v\right) = \begin{cases} \frac{1-\underline{u}}{2-\underline{u}} & \text{if } \underline{v} \leqslant \frac{1}{2-\underline{u}}; \\ 1-\underline{v} & \text{otherwise.} \end{cases}$$

The WCR *under any mechanism is at least R^s. The mechanism*

$$\eta^s(u,v) = \begin{cases} 1, & \text{if } v \geqslant 1-R^s; \\ \underline{u}/u, & \text{if } v < 1-R^s. \end{cases}$$

has the WCR *of R^s, so it is optimal.*

The mechanism η^s consists of an *automatic-approval* region and a *chance* region. If the proposed project is sufficiently good for the principal (i.e., $v \geqslant 1 - R^s$), then it is automatically approved. If the project is mediocre for the principal (i.e., $v < 1 - R^s$), then the approval probability equals $\underline{u}/u$, so the agent expects a payoff $\underline{u}$ from proposing a mediocre project.

The agent will propose a project in the automatic-approval region if he has one such project. If all his projects are in the chance region, he will propose a project that gives the principal the highest payoff. Nonetheless, the principal still suffers regret from two sources. First, if the agent has multiple projects that will be automatically approved, he will propose what he favors instead of what the principal favors. Second, if the agent has only projects in the chance region, his proposal is rejected with positive probability. The threshold for the automatic-approval region, $1 - R^s$, is chosen to keep the regret from both sources under control.

The approval probability $\eta^s(u, v)$ increases in v (the principal's payoff) and decreases in u (the agent's payoff). This monotonicity in v and u is natural. In particularly, 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 to the one in Armstrong and Vickers (2010), which studies the single-project environment in a Bayesian setting. They focus on deterministic mechanisms, so $\eta(u, v) \in \{0, 1\}$ for every project (u, v). Under the assumptions that (i) projects are i.i.d. and (ii) the number of available projects is independent of their characteristics, they show that the optimal deterministic mechanism $\eta(u, v)$ is monotone-increasing in v: a project (u, v) is approved if and only if $v \geqslant r(u)$ for some function r. One important question to ask, which we leave for future research, is under which assumptions on the prior the optimal randomized mechanism in the Bayesian setting is monotone-increasing in v and monotone-decreasing in u.

4.3 Optimal mechanism in the multiproject environment

In the single-project environment, the only way for the principal to incentivize the agent is to reject his proposal with positive probability. The multiproject environment, however, allows the principal to "spend" this rejection probability on projects that he likes. We now characterize how to optimally reduce the amount of inefficient rejection and thus lower the principal's WCR.

Let $\eta : [\underline{u}, 1]\times[\underline{v}, 1] \to [0, 1]$ be a function and consider the following *project-wide maximal-payoff mechanism* (PMP mechanism) induced by η:

1. If the proposal P includes only one project (u, v), it is approved with probability $\eta(u, v)$.

2. If P includes 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 $\max_{(u,v)\in P} \eta(u, v)u$. This is the maximal expected payoff the agent could get from proposing each project alone.

By the definition of the PMP mechanism, the more projects the agent proposes, the weakly higher his expected payoff will be. Hence, the PMP mechanism is IC. Our next theorem shows that there exists an optimal PMP mechanism.

Theorem 4.2. *Assume $\mathcal{E} = 2^D$. For every $p \in [\underline{u}/u, 1]$, let $\alpha(p, u) = \frac{pu-\underline{u}}{u-\underline{u}}$. Let*

$$R^m \equiv \max_{(u,v)\in[\underline{u},1]\times[\underline{v},1]} \min_{p\in[\underline{u}/u,1]} \max\big(v(1-p), (1-v)\alpha(p, u)\big) = \begin{cases} \frac{(1-\underline{u})\left(2-\underline{u}-2\sqrt{1-\underline{u}}\right)}{\underline{u}^2} & \underline{v} < \frac{1-\sqrt{1-\underline{u}}}{\underline{u}}; \\ \frac{(1-\underline{u})(1-\underline{v})\underline{v}}{1-\underline{u}\,\underline{v}} & \underline{v} \geqslant \frac{1-\sqrt{1-\underline{u}}}{\underline{u}}. \end{cases}$$

The WCR *under any mechanism is at least R^m. The PMP mechanism induced by*

$$\eta^m(u, v) = \max\{p \in [\underline{u}/u, 1] : (1-v)\alpha(p, u) \leqslant R^m\} = \begin{cases} 1, & \text{if } v \geqslant 1 - R^m; \\ \left(1 - \frac{R^m}{1-v}\right)\frac{\underline{u}}{u} + \frac{R^m}{1-v}, & \text{if } v < 1 - R^m. \end{cases}$$

has the WCR *of* R^m*, so it is optimal.*

Like the function $\eta^s(u, v)$ in the single-project environment, the function $\eta^m(u, v)$ equals 1 if the payoff to the principal is sufficiently high (i.e., $v \geqslant 1 - R^m$), and is less than 1 otherwise (i.e., $v < 1 - R^m$). Hence, if the agent has only one project, the project is chosen for sure if $v \geqslant 1 - R^m$, and is chosen with partial probability otherwise. For this reason, we still call the region with $v \geqslant 1 - R^m$ the *automatic-approval* region, and we call the region with $v < 1 - R^m$ the *chance* region.

When the agent proposes more than one project, the principal promises the agent an expected payoff of $\max_{(u,v)\in P} \eta^m(u, v)u$. In both panels of figure 1, each dotted curve connects all the projects that induce the same value of $\eta^m(u, v)u$, so it can be interpreted as an "indifference curve." For a project in the automatic-approval region, the principal is willing to compensate the agent his full payoff. In contrast, for a project in the chance region, the principal is willing to compensate 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 automatic-approval region and tilt counterclockwise as v further decreases. The agent's expected payoff is determined by the project (among those proposed) that is on the highest indifference curve.

Under the optimal mechanism, if the agent's favorite project is in the automatic-approval region, then this project will be chosen for sure. In this case, there is no benefit to either party from proposing other available projects. The left panel of figure 1 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 in the chance region, the benefit to the principal from the agent's proposing multiple projects can be significant. The right panel of figure 1 illustrates such an example. Instead of rejecting $\blacktriangle$ with positive probability, the mechanism randomizes between $\blacktriangle$ and $\star$ while promising the agent the same payoff he would get from proposing $\blacktriangle$ alone. In such cases, the optimal mechanism imposes a compromise between the two parties:

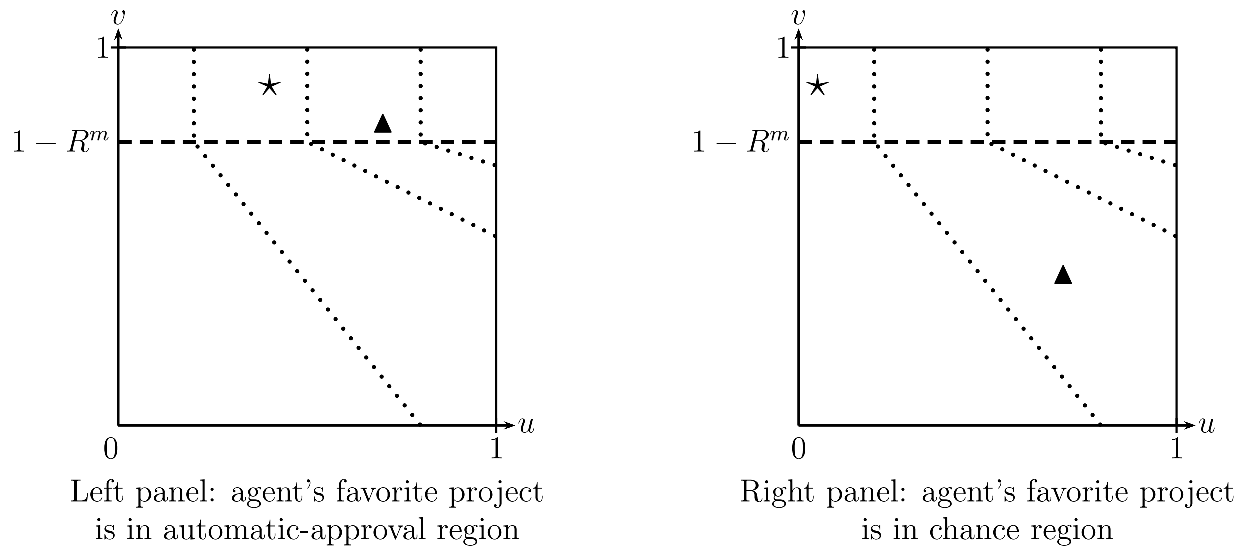

Figure 1: Reaching a compromise when agent's favorite project is in chance region, $\underline{u} = \underline{v} = 0$

sometimes the choice favors the agent, and at other times it favors the principal.

Figure 2 compares the WCR under the single-project and the multiproject environments. Roughly, the principal's gain from having the multiproject environment as compared to the single-project environment, measured by $R^m - R^s$, is larger when $\underline{u}$ or $\underline{v}$ is smaller (i.e., when the principal faces more uncertainty).

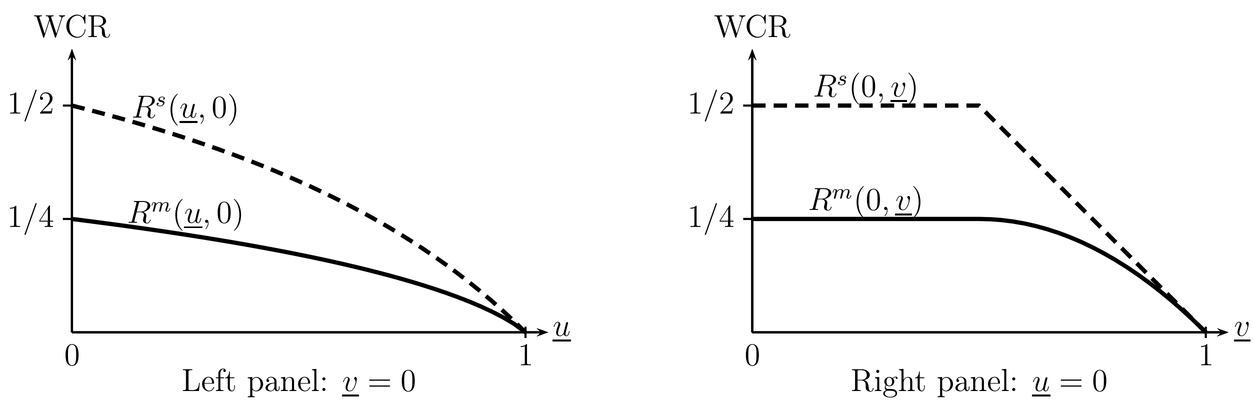

Figure 2: WCR: single-project (dashed curve) vs. multiproject (solid curve)

5 Discussion

5.1 Intermediate verification capacity

We have focused on the single-project and the multiproject environments, which are natural first steps for us to study. Nonetheless, there are intermediate environments in which the principal can verify up to K projects for some fixed $K \geqslant 2$, so $\mathcal{E} = \{P \subseteq D : |P| \leqslant K\}$. We call this the K-project environment.

Proposition 5.1 (Two are enough)**.** *For any $K \geqslant 2$, the PMP mechanism induced by $\eta^m(u,v)$ is optimal in the K-project environment. The* WCR *under this mechanism is R^m.*

Under the PMP mechanism induced by $\eta^m(u,v)$, the agent will include in his proposal a project that maximizes $\eta^m(u,v)u$, since the mechanism promises him an expected payoff of $\max_{(u,v)\in P} \eta^m(u,v)u$. The agent is indifferent among other projects, so he is willing to include the principal's favorite project (i.e., a project in $\operatorname{argmax}_{(u,v)\in A} v$). Based on the proof of Theorem 4.2, the compromise between these two projects already guarantees that the principal's regret is at most R^m. Hence, the mechanism is optimal for any $K \geqslant 2$.

Proposition 5.1 shows that having the full benefit of compromise requires not infinite or high verification capacity, but a capacity of only two projects. Furthermore, even if the principal can verify up to ten projects, it suffices to let the agent propose up to two, which provides a parsimonious way to get the full benefit of compromise.

5.2 Cheap-talk messages do 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, as in Bull and Watson (2007) and Ben-Porath, Dekel and Lipman (2019). However, in our model cheap talk does not benefit the principal. This is because the principal can choose a project only from the proposed set P and he knows the payoffs that each project in P gives to both parties. Hence, no information asymmetry remains after the agent proposes P, and so there is no benefit to cheap talk.

More specifically, for any proposal 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 after (P, m_1) and after (P, m_2). Suppose otherwise that the principal chooses a subprobability 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 messages do not help for any $\mathcal{E}$.

6 Proofs

6.1 Proof of Theorem 4.1

Claim 1. *The* WCR *from any mechanism is at least R^s.*

Proof. Let $v \in [\underline{v}, 1]$. If $\eta(1, v) > \underline{u}$, then, if the agent has two projects $(1, v)$ and $(\underline{u}, 1)$, the agent will propose $(1, v)$ and the regret will be $1 - \eta(1, v)v \geqslant 1 - v$. If $\eta(1, v) \leqslant \underline{u}$, then, if the agent has only the project $(1, v)$, the regret is $v - \eta(1, v)v \geqslant v(1 - \underline{u})$. Therefore, WCR $\geqslant \min\left((1 - \underline{u})v, 1 - v\right)$ for every $v \in [\underline{v}, 1]$. □

Claim 2. *The* WCR *from η^s is R^s.*

Proof. We call a project (u, v) *good* if $v \geqslant 1 - R^s$ and *mediocre* if $v < 1 - R^s$. From the definition of R^s it follows that $(1 - \underline{u})v \leqslant R^s$ for every mediocre project.

According to η^s, if the agent proposes a mediocre project, then his expected payoff is $\underline{u}$; if the agent proposes a good project (u, v), then his expected payoff is $u \geqslant \underline{u}$. Therefore, if the agent has some good project, he will propose a good project (u, v) and the regret is at most $1 - v \leqslant R^s$. If all projects are mediocre, then the agent will propose the project (u, v) with the highest v, so the regret is at most $(1 - \eta^s(u, v))v = (1 - \underline{u}/u)v \leqslant (1 - \underline{u})v \leqslant R^s$. □

6.2 Proof of Theorem 4.2

We first note the following properties of $\alpha(p, u)$ and η^m that follow from the definitions of $\alpha(p, u)$, R^m, and η^m:

$$\alpha(p, u)u + (1 - \alpha(p, u))\underline{u} = pu \text{ for every } u \in [\underline{u}, 1], p \in [\underline{u}/u, 1]. \tag{1}$$

$$v(1 - \eta^m(u, v)) \leqslant R^m, \text{ and } (1 - v)\alpha(\eta^m(u, v), u) \leqslant R^m. \tag{2}$$

To see (2), note that, by the definition of R^m, there exists some $\bar{p} \in [\underline{u}/u, 1]$ such that $\max(v(1 - \bar{p}), (1 - v)\alpha(\bar{p}, u)) \leqslant R^m$. By the definition of η^m, the expression $\eta^m(u, v)$ is the highest p such that $(1 - v)\alpha(p, u) \leqslant R^m$. Since $(1 - v)\alpha(\bar{p}, u) \leqslant R^m$, it follows that $\bar{p} \leqslant \eta^m(u, v)$. Therefore, it also follows that $v(1 - \eta^m(u, v)) \leqslant v(1 - \bar{p}) \leqslant R^m$.

Claim 3. *The* WCR *from any mechanism is at least R^m.*

Proof. Let ρ be a mechanism. By Proposition 2.1 we can assume that ρ is IC. Fix a project $a = (u, v)$ and let $p = \rho(a|\{a\})$ be the probability that the mechanism ρ approves a when the agent proposes a alone. If the set of available projects is $A = \{a\}$, then the agent will propose the project a, and it will be implemented with probability p, causing regret $(1-p)v$. If the set of available projects is $A = \{a, (\underline{u}, 1)\}$, then, since the agent has the option to propose $P = \{a\}$ and he still prefers to propose the truth $P = A$, it follows that his payoff under the mechanism is at least pu. Let $q = \rho(a|\{a, (\underline{u}, 1)\})$. Then $qu + (1 - q)\underline{u} \geqslant pu$, which, by (1), implies that $q \geqslant \alpha(p,u)$. The regret is therefore at least $\alpha(p,u)(1-v)$. It follows that

$$\mathrm{WCR}(\rho) \geqslant \max((1-p)v, \alpha(p,u)(1-v)). \tag{3}$$

We have now proved that for every u, v there exists a p such that (3) holds. By the definition of R^m, it follows that $\mathrm{WCR}(\rho) \geqslant R^m$, as desired. □

Claim 4. *The* WCR *from the PMP mechanism induced by* η^m *is at most* R^m.

Proof. Since the PMP mechanism is IC, we assume that the agent proposes truthfully. Fix a set A of available projects, and let $(u_a, v_a), (u_p, v_p) \in A$ be such that

$$(u_a, v_a) \in \operatorname{argmax}\{\eta^m(u,v)n : (u,v) \in A\} \text{ and } (u_p, v_p) \in \operatorname{argmax}\{v : (u,v) \in A\}. \tag{4}$$

We need to prove that there exists some subprobability measure over A which gives the agent an expected payoff $\eta^m(u_a, v_a)u_a$, and induces regret of at most R^m. Let I be the set of all $y \in \mathbf{R}$ such that y is the agent's expected payoff under some subprobability measure over A that induces regret of at most R^m. Since I is convex (an interval), it is sufficient to show that there exists $\underline{y}, \bar{y} \in I$ such that $\underline{y} \leqslant \eta^m(u_a, v_a)u_a \leqslant \bar{y}$.

1. The subprobability measure that picks (u_p, v_p) with probability $\eta^m(u_p, v_p)$ induces regret $v_p(1-\eta^m(u_p, v_p)) \leqslant R^m$, by (2). It gives the agent the expected payoff $\underline{y} = u_p\eta^m(u_p, v_p) \leqslant \eta^m(u_a, v_a)u_a$, where the inequality follows from the definition of (u_a, v_a).

2. The probability measure that picks (i) (u_a, v_a) with probability $\alpha(\eta^m(u_a, v_a), u_a)$, and (ii) (u_p, v_p) with probability $(1-\alpha(\eta^m(u_a, v_a), u_a))$ gives the following regret:

$$\alpha(\eta^m(u_a, v_a), u_a))(v_p - v_a) \leqslant \alpha(\eta^m(u_a, v_a), u_a)(1-v_a) \leqslant R^m,$$

where the second inequality follows from (2), and gives the agent payoff of

$$\bar{y} = \alpha(\eta^m(u_a, v_a), u_a)u_a + (1 - \alpha(\eta^m(u_a, v_a), u_a))u_p$$
$$\geqslant \alpha(\eta^m(u_a, v_a), u_a)u_a + (1 - \alpha(\eta^m(u_a, v_a), u_a))\underline{u} = \eta^m(u_a, v_a)u_a,$$

where the last equality follows from (1).

□

References

Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy*, 105(1): 1–29.

Armstrong, Mark, and John Vickers. 2010. "A Model of Delegated Project Choice." *Econometrica*, 78(1): 213–244.

Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman. 2019. "Mechanisms with evidence: Commitment and robustness." *Econometrica*, 87(2): 529–566.

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.

Beviá, Carmen, and Luis Corchón. 2019. "Contests with dominant strategies." *Economic Theory*.

Bonatti, Alessandro, and Heikki Rantakari. 2016. "The Politics of Compromise." *American Economic Review*, 106(2): 229–59.

Bull, Jesse, and Joel Watson. 2007. "Hard evidence and mechanism design." *Games and Economic Behavior*, 58(1): 75 – 93.

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

Carroll, Gabriel. 2019. "Robustness in Mechanism Design and Contracting." *Annual Review of Economics*, 11(1): 139–166.

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

Dekel, Eddie. 2016. "On Evidence in Games and Mechanism Design." *Econometric Society Presidential Address.*

Dye, Ronald A. 1985. "Strategic Accounting Choice and the Effects of Alternative Financial Reporting Requirements." *Journal of Accounting Research*, 23(2): 544–574.

Glazer, Jacob, and Ariel Rubinstein. 2006. "A study in the pragmatics of persuasion: a game theoretical approach." *Theoretical Economics*, 1: 395–410.

Goel, Sumit, and Wade Hann-Caruthers. 2020. "Project selection with partially verifiable information."

Green, Jerry R., and Jean-Jacques Laffont. 1986. "Partially Verifiable Information and Mechanism Design." *The Review of Economic Studies*, 53(3): 447–456.

Grossman, Sanford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *The Journal of Law and Economics*, 24(3): 461–483.

Grossman, S. J., and O. D. Hart. 1980. "Disclosure Laws and Takeover Bids." *The Journal of Finance*, 35(2): 323–334.

Guo, Yingni, and Eran Shmaya. 2019. "Robust Monopoly Regulation." *Working paper.*

Hart, Sergiu, Ilan Kremer, and Motty Perry. 2017. "Evidence Games: Truth and Commitment." *American Economic Review*, 107(3): 690–713.

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.

Malladi, Suraj. 2020. "Judged in Hindsight: Regulatory Incentives in Approving Innovations." *Available at SSRN.*

Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *The Bell Journal of Economics*, 12(2): 380–391.

Nocke, Volker, and Michael D. Whinston. 2013. "Merger Policy with Merger Choice." *American Economic Review*, 103(2): 1006–33.

Renou, Ludovic, and Karl H. Schlag. 2011. "Implementation in minimax regret equilibrium." *Games and Economic Behavior*, 71(2): 527 – 533.

Savage, Leonard J. 1972. *The foundations of statistics.* Courier Corporation.

Savage, L. J. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association*, 46(253): 55–67.

Sher, Itai. 2014. "Persuasion and dynamic communication." *Theoretical Economics*, 9(1): 99–136.