Project Choice from a Verifiable Proposal

Yingni Guo Eran Shmaya*

September 9, 2021

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. 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 proposing multiple projects equals the maximal payoff he would get from proposing each project alone. Our results highlight the benefits from the ability to propose multiple projects and from randomization.

JEL: D81, D82, D86

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

*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, Northwestern University Theory Lunch Seminar, Kellogg Strategy Department Lunch Seminar, the University of Pittsburgh, Carnegie Mellon University, Stony Brook International Conference on Game Theory, and Workshop on Strategic Communication and Learning for valuable feedback.

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 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 a factory location or an office building. 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 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 preferences: the division may fail to internalize projects' externalities on other divisions; 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 this question using a robust-design, non-Bayesian approach. Such an approach is useful when the principal faces 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, the 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 worst-case 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. Even though at most one project will be chosen, one often observes that the agent proposes multiple projects. The strategic role that a multiproject proposal plays in project-choice problems has thus far been unexplored. 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 restriction might emerge in some organizations due to the principal's limited attention. It is also applicable to antitrust regulation, for which a firm proposes one merger and the regulator decides whether to approve or reject the firm's proposal (e.g., 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 familiar 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.[1] 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.

This two-tier mechanism 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, he will propose a top-tier project for sure. Second, if all his projects are bottom-tier ones, then he will propose the principal's favorite project among the available ones. 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. As a result, 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 has a similar two-tier structure as in the single-project environment. 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. Otherwise, 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

[1]Top-tier projects are often called slam-dunk cases in many academic departments.

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.[2] Since the agent gets his maximal expected payoff from proposing each project alone, we call this mechanism a *project-wide maximal-payoff mechanism* (PMP mechanism).

This optimal 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 the mechanism can optimize over a larger set of projects and potentially improve the principal's expected payoff. Such an improvement would not be possible in the single-project environment. We show that it significantly reduces the principal's worst-case regret in the multiproject environment compared to the single-project one.

Our result intuitively explains when several proposed projects 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 several proposed projects. In such cases, 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.

[2]This maximization problem is a linear programming problem. Only projects on the Pareto frontier will be chosen. Moreover, there always exists an optimal solution whose support has at most two projects.

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.[3] 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 worst-case 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.

Goel and Hann-Caruthers (2020) consider the project-choice problem in which 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 as our agent can, he loses the proposal-setting power. The resulting incentive schemes are thus quite different.

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. Hence, our paper is closely 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)

[3]The 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.

for a survey). We will discuss the relationship of our work to this literature in more detail after we introduce the model.

Our result relates to a theme in Aghion and Tirole (1997), namely, that although the principal has formal authority, the agent shares real authority due to his private information. We take this theme 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 "finding some middle ground" is related to Bonatti and Rantakari (2016). They study how to strike a compromise between two symmetric agents whose efforts are crucial for discovering projects. In contrast, we focus on the relationship between an agent who proposes projects and a principal who chooses one or none from the projects proposed.

The worst-case 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 worst-case 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. 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. For a survey on robustness in mechanism design, see Carroll (2019).

2 Model and mechanism

Let D be the domain of all possible *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, the agent can propose only available projects. Hence, the agent's proposal must be a subset of his type, $P \subseteq A$. This is what we meant earlier when we said that projects are verifiable. 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 many organizations, the principal has limited attention, so the agent can propose at most a certain number of projects.

We begin with two environments which are natural first steps: *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. In subsection 6.1, we discuss the intermediate environments in which the agent can propose up to K projects for some fixed number $K \geqslant 2$.

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 implication of this evidence role in section 3. Second, the proposal determines the set of projects from which the principal can choose. This second role is a key difference between our paper and the evidence literature. In particular, once the agent puts his proposal on the table, there is no relevant information asymmetry left. This implies that cheap-talk communication does not help. We elaborate on this point in subsection 6.2.

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

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

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

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, on the contrary, the principal must choose a project as long as the agent has proposed some, then the agent effectively has all the bargaining power. He will propose his favorite project which will be chosen for sure.

A *mechanism* ρ attaches to each proposal $P \in \mathcal{E}$ a subprobability measure $\rho(\cdot|P)$ such that 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 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 $f(a|A) = \sum_P \mu(P)\rho(a|P)$. The interpretation is twofold. First, among the proposals that the agent can make, he selects only proposals that give him the highest expected payoff. Second, if the agent has multiple optimal proposals, then he can randomize among them.

3 The evidence structure

When the agent proposes a set P of projects, he provides evidence that his type A satisfies $P \subseteq A$. In this section, we discuss the implication of this evidence role of the agent's proposal as well as the relation to the evidence literature.

3.1 Normality in the multiproject environment

In our multiproject environment, in which $\mathcal{E} = 2^D$, the agent has the ability to provide the maximal evidence for his type. This property is called *normality* in the literature (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, in which $\mathcal{E} = \{P \subseteq D : |P| \leqslant 1\}$, normality does not hold. The single-project environment is the main focus in Armstrong and Vickers (2010) and Nocke and Whinston (2013). It is also 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.

3.2 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 all the available projects. That is, $U(\rho, A) \geqslant U(\rho, P)$ 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 3.1. *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.*

As we explained in subsection 3.1, the multiproject environment satisfies normality and the nested-range condition. Previous papers (e.g., Green and Laffont (1986), Bull and Watson (2007)) have shown that the revelation principle holds under these conditions. Our proposition 3.1 does not, however, follow directly from their theorems, because the agent's proposal P serves two roles in our model. In addition to providing evidence, the proposal also determines the set of projects from which the principal can choose. Nevertheless, a similar argument for the revelation principle can be made within our model as well.

Proof of Proposition 3.1. 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:

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

where (i) the inequality follows from the fact that $Q \subseteq P$ implies $Q \subseteq A$, and (ii) 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. □

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.

4 The principal's problem

Let $v : D \to \mathbf{R}_+$ be the principal's payoff function, 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:

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

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. Savage (1951) calls this difference *loss*. We call it *regret* instead, and thereby follow the more recent game theory and computer science literature. Wald (1950) and Savage (1972) propose considering only *admissible* choice functions (i.e., choice functions that are not weakly dominated). A choice function f is *admissible* if there exists no other 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.

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

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

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 $\mathrm{WCR}(f)$ over all implementable choice functions f. This step is our only departure from the Bayesian approach. The Bayesian approach will instead assign a prior belief over the number and the payoff characteristics of the available projects. The principal's problem, then, is to minimize the *expected* regret instead of the *worst-case* regret.

While our principal takes the worst-case 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 worst-case 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. If we make the alternative assumption that the principal takes the worst-case regret approach even towards his own randomization, we effectively restrict the principal to deterministic mechanisms.

From now on, we assume that the set D of all possible verifiable 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.

5 Optimal mechanisms

5.1 Preliminary intuition

We now use an example to illustrate the fundamental trade-off faced by the principal, as well as 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

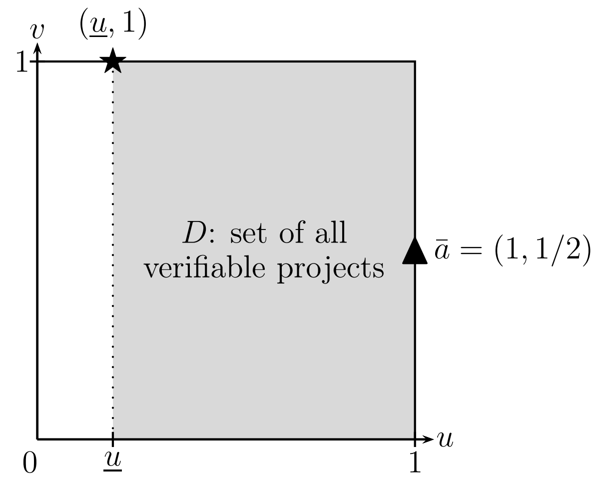

Figure 1: Preliminary intuition, $\underline{v} = 0$

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 several projects which will be approved, then he will propose the one 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 $\bar{a} = (1, 1/2)$, these two fears imply that no matter how the principal designs the deterministic mechanism, his WCR is at least $1/2$.

As shown in figure 1, the project $\bar{a}$ 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 $\bar{a}$ 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 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 incentive 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 5.1 in subsection 5.2.

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.[4] 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 $\{\bar{a}, (\underline{u}, 1)\}$, 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 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 5.2 in subsection 5.3. 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] 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 at most one available project.

5.2 Optimal mechanism in the single-project environment

Since the agent 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).

Theorem 5.1 (Single-project environment). *Assume $\mathcal{E} = \{P \subseteq D : |P| \leqslant 1\}$. Let*

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

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

(ii) Let α^s be the mechanism given by:

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

Then α^s implements a choice function that 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 $\alpha(u,v) \leqslant \alpha^s(u,v)$ for every $(u,v) \in D$.*

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 $1-R^s$. If the proposed project is a top-tier project for the principal, then it is approved for sure. If it is in the bottom tier, then the approval probability equals $\underline{u}/u$, so the agent expects a constant payoff of $\underline{u}$ from proposing a bottom-tier project.

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 will propose the principal's favorite project among the available ones. The principal still suffers regret from two sources. First, if the agent has several 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 projects in the bottom tier, 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 and only if $v \geqslant r(u)$ for some function $r(u)$. They also characterize the function $r(u)$ explicitly. Their argument can be generalized to show that the optimal randomized mechanism $\alpha(u, v)$ also increases in v, but it is unclear how to solve for the optimal $\alpha(u, v)$.

The typical situation under the worst-case regret approach to uncertainty is that multiple mechanisms can achieve the minimal WCR. Statement (iii) in Theorem 5.1 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 and a strictly higher payoff for some singleton A.

5.3 Optimal mechanism in the multiproject environment

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

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

2. If the proposal 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 $\max_{(u,v)\in P} \alpha(u, v)u$. 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 IC. For a mechanism to be IC, 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 is promised exactly his maximal payoff from proposing each project alone, but not more.

Our next theorem shows that there exists an optimal PMP mechanism.

Theorem 5.2 (Multiproject environment). *Assume $\mathcal{E} = 2^D$. For every $u \in [\underline{u}, 1]$ and $p \in [0, 1]$, let $\gamma(u, p)$ be*

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

Let

$$R^m = \max_{(u,v)\in D} \min_{p\in[0,1]} \max\big(v(1 - p), (1 - v)\gamma(u, p)\big). \tag{2}$$

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

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

$$\alpha^m(u,v) = \max\{p \in [0,1] : (1-v)\gamma(u,p) \leqslant R^m\}. \tag{3}$$

The mechanism ρ^m has the WCR *of R^m and is admissible.*

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

The explicit expressions for R^m and $\alpha^m(u,v)$ are presented at the end of this subsection.

It follows from expressions (1) to (3) that $\alpha^m(u,v) = 1$ if $v \geqslant 1 - R^m$ and $\alpha^m(u,v) < 1$ otherwise. As in the case of the single-project environment, when the agent proposes only one project, the project is approved for sure if its payoff to the principal is sufficiently high and is rejected with positive probability otherwise. For this reason, we still call a project with $v \geqslant 1 - R^m$ a top-tier project and a project with $v < 1 - R^m$ a bottom-tier one. Figure 2 depicts these two tiers.

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 2, 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. 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 tilt counterclockwise as the principal's payoff v further decreases.

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 2 gives an intuitive answer. If the agent's favorite project is already a top-tier project for the

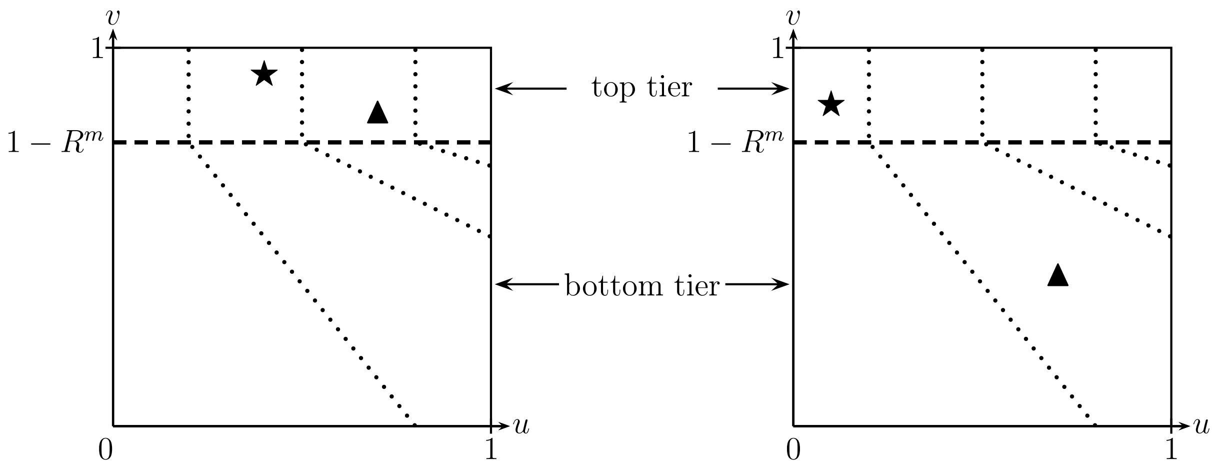

Left panel: Agent's favorite project is in the top tier.

Right panel: Agent's favorite project is in the bottom tier.

Figure 2: Optimal mechanism in the multiproject environment, $\underline{u} = \underline{v} = 0$

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: ★ and ▲ denote the available projects and ▲ will be chosen for sure. In contrast, if the agent's favorite project is a bottom-tier 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. Instead of rejecting ▲ with positive probability, the mechanism randomizes between ▲ and ★ while promising the agent the same payoff he would get from proposing ▲ alone. In such cases, several 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 5.2 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 by:

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

and

$$\alpha^m(u,v) = \begin{cases} 1, & \text{if } v \geqslant 1 - R^m \text{ or } u = 0, \\ \left(1 - \frac{R^m}{1-v}\right)\frac{\underline{u}}{u} + \frac{R^m}{1-v}, & \text{if } v < 1 - R^m \text{ and } u > 0. \end{cases}$$

Remark 1. The PMP mechanism ρ^m is weakly IC. We could modify the mechanism to make it strictly IC. To do this, let $h : [0, \infty) \to [0, 1]$ be a strictly increasing function and $|P|$ the number of projects in the proposal P. If we modify α^m with $(1 - \varepsilon + \varepsilon h(|P|))\alpha^m$ for some small $\varepsilon > 0$, the PMP mechanism induced by this modified function is strictly IC. The WCR under this modified mechanism approaches R^m as $\varepsilon \to 0$.

5.4 Comparing the WCR under two environments

Figure 3 compares the WCR under the single-project and the multiproject environments. 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, 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).

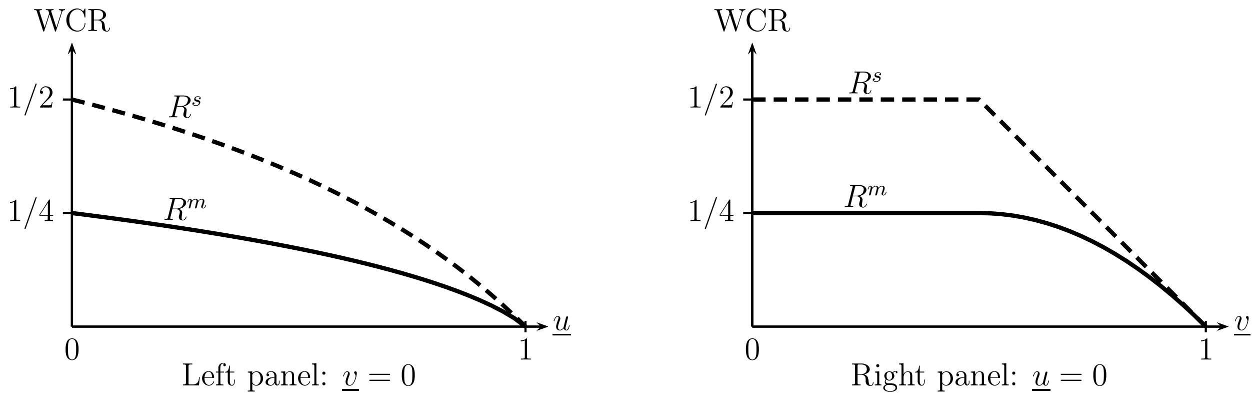

Figure 3: WCR: single-project (dashed curve) vs. multiproject (solid curve)

6 Extensions and discussion

6.1 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 $K \geqslant 2$, so that $\mathcal{E} = \{P \subseteq D : |P| \leqslant K\}$. We call this case the K-project environment. The multiproject environment is the limit case when $K \to \infty$.

The following proposition shows that as long as $K \geqslant 2$, the principal's minimal WCR is R^m.

Proposition 6.1 (Two is enough)**.** *In the K-project environment, the* WCR *under any mechanism is at least R^m. For any $K \geqslant 2$, the PMP mechanism ρ^m induced by $\alpha^m(u, v)$ 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 the 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 $P = \{(u_p, v_p), (u_a, v_a)\}$. Then under the PMP mechanism 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 the mechanism. The principal's payoff given the proposal P equals his payoff if the set of available projects were actually P. By Theorem 5.2 this payoff is at least $v_p - R^m$, so the principal's regret is at most R^m. □

Based on Theorem 5.1 and Proposition 6.1, 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 several projects rather than proposing only a single 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 serves as the fallback option when the first project is not chosen. According to the proof of Proposition 6.1, 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 mechanism ρ^m in the multiproject environment, 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 the agent to propose the aforementioned two projects: the one that the agent himself wants to propose and the one that the principal likes the most.

6.2 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, 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 over P 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. Since this argument does not depend on the exogenous restriction $\mathcal{E}$ on the agent's proposal P, cheap-talk communication does not help for any $\mathcal{E}$.

6.3 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 from the project(s) proposed. 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 with only the principal's favorite project from the project(s) proposed. The optimal mechanism in this partial-commitment setting is then the same as that in the single-project environment characterized in Theorem 5.1.

6.4 The payoff guarantee by the minimax-regret mechanism

So far we have focused on settings in which the uncertainty is nonprobabilistic. The principal evaluates possible mechanisms by the minimax-regret criterion. The resulting minimax-regret mechanism offers a reasonable contender even in settings in which the principal can formulate a prior belief over the set of available projects. We now show that this mechanism is especially appealing when the principal's expected complete-information payoff is high.

Under the minimax-regret mechanism, for every set A of available projects, the principal's payoff is at least his complete-information payoff minus the minimal worst-case regret. The complete-information payoff is his payoff from choosing his favorite project in A. The minimal WCR is R^s in the single-project environment and R^m in the multiproject one. Hence, for any prior belief over the set of available projects, the minimax-regret mechanism offers the principal a payoff guarantee which equals his expected complete-information payoff minus the minimal WCR.

This payoff guarantee increases as the principal's expected complete-information payoff increases, so it is particularly useful when this payoff is high.

To illustrate this point, we focus on the single-project environment and consider the following prior belief over the set of available projects. Suppose that the number of available projects is given by a Poisson distribution with mean λ. Projects' payoff characteristics are drawn independent of the number of available projects, and also independent of each other. For each project, the payoff u and the payoff v are drawn independently. They are uniform on $[\underline{u}, 1]$ and $[\underline{v}, 1]$, respectively.

We now compare the payoff guarantee and the performance of two different mechanisms: our mechanism α^s which minimizes the principal's maximal regret and a mechanism which maximizes his minimal payoff.[5]

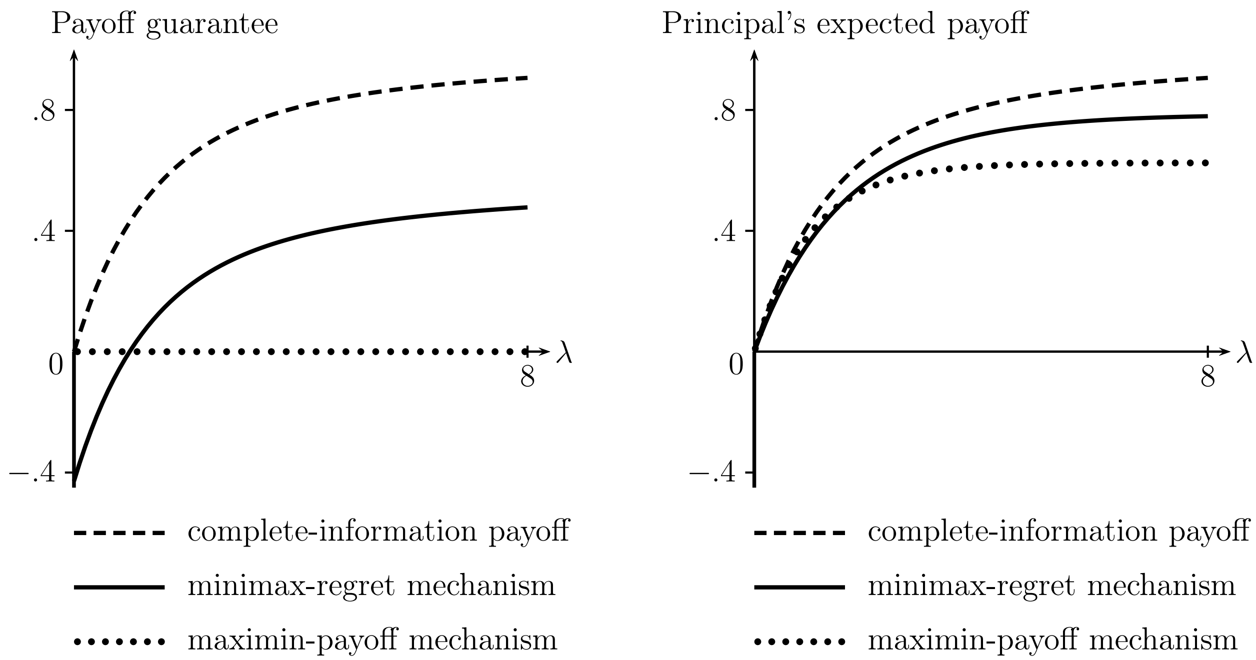

Figure 4: Comparing the minimax-regret versus maximin-payoff mechanisms, where $\underline{u} = \underline{v} = \frac{1}{4}$.

As the mean λ of the Poisson distribution increases, the principal expects more available projects, so a higher expected complete-information payoff results. In the left panel of figure 4, the dashed curve shows the principal's expected complete-information payoff as a function of λ. The solid curve shows the payoff guarantee under our minimax-regret mechanism α^s. This payoff guarantee differs from the expected complete-information payoff by the minimal WCR R^s. The dotted line shows the payoff guarantee under a maximin-payoff mechanism.

[5]The principal's minimal payoff under any mechanism is zero, so all mechanisms are optimal according to the maximin-payoff criterion. For the comparison in this section, we use the maximin-payoff mechanism that chooses the principal's favorite project among those proposed.

The minimax-regret mechanism offers a higher payoff guarantee than the maximin-payoff mechanism does when the expected complete-information payoff is sufficiently high.

The right panel of figure 4 shows the principal's expected payoff under these two mechanisms. This comparison is similar to the comparison of the payoff guarantee in the left panel. The minimax-regret mechanism outperforms the maximin-payoff mechanism if the principal's expected complete-information payoff is sufficiently high. Otherwise, the maximin-payoff mechanism does better.

7 Conclusion

A tremendous share of decision-making in organizations and in antitrust regulation can be viewed as project-choice problems. How to counteract the agent's proposing bias is a first-order concern. In this paper, we characterize prior-free mechanisms that counteract this bias, both when the agent can propose only a single project and when he can propose multiple ones. Along the way, we show that multiproject proposals play an important strategic role in providing the principal with better fallback options than rejection.

Our analysis shows that the worst-case regret approach not only is tractable but also preserves the fundamental trade-off faced by the principal. This approach thus has the potential to shed light on other variants of the project-choice problem. For instance, the project-choice process may involve a deeper hierarchy when there are three or more players; 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.

8 Proofs

8.1 Proof of Theorem 5.1

Claim 8.1. *The* WCR *from any mechanism is at least* R^s*.*

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

Claim 8.2. *The* WCR *from* α^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 α^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 $(1 - \alpha^s(u, v))v = (1 - \underline{u}/u)v \leqslant (1 - \underline{u})v \leqslant R^s$. □

Claim 8.3. *If* α *has the* WCR *of* R^s*, then* $\alpha(u, v) \leqslant \alpha^s(u, v)$ *for every* $(u, v) \in D$*. Hence,* α^s *is admissible.*

Proof. Fix a project (u, v). If $v \geqslant 1 - R^s$ or $u = 0$, then $\alpha^s(u, v) = 1$ and therefore $\alpha(u, v) \leqslant \alpha^s(u, v)$. 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 $\alpha(u, v)u \leqslant \alpha(\underline{u}, 1)\underline{u} \leqslant \underline{u}$, which implies $\alpha(u, v) \leqslant \underline{u}/u = \alpha^s(u, v)$, as desired.

Finally, if α has the WCR of R^s and $\alpha \neq \alpha^s$, then there exists $(u,v) \in D$ such that $\alpha(u,v) < \alpha^s(u,v)$. The regret is strictly higher under α than under α^s if $A = \{(u,v)\}$, so α^s is admissible. ☐

8.2 Proof of Theorem 5.2

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

$$\overline{U}(P) = \max_{\pi} \quad \underline{u} + \sum_{(u,v)\in P} \pi(u,v)(u - \underline{u}) \tag{4a}$$

$$\text{s.t.} \quad \pi(u,v) \geqslant 0, \ \forall (u,v) \in P, \tag{4b}$$

$$\sum_{(u,v)\in P} \pi(u,v) \leqslant 1, \tag{4c}$$

$$\sum_{(u,v)\in P} \pi(u,v)(1-v) \leqslant R^m. \tag{4d}$$

The following claim explains the role of $\overline{U}(P)$ in our argument: $\overline{U}(P)$ is the maximal payoff that the principal can give the agent for the proposal 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 8.4. *If ρ is an IC mechanism which has the* WCR *of at most R^m, then $U(\rho, P) \leqslant \overline{U}(P)$ 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 $\sum_{(u,v)\in P} \pi(u,v)(1-v) \leqslant R^m$. Therefore the restriction of π to the set P is a feasible point in problem (4). Moreover

$$U(\rho, \tilde{P}) = \pi(a^*)\underline{u} + \sum_{(u,v)\in P} \pi(u,v)u \leqslant \underline{u} + \sum_{(u,v)\in P} \pi(u,v)(u - \underline{u}), \tag{5}$$

where the inequality follows from $\pi(a^*) + \sum_{(u,v)\in P} \pi(u,v) \leqslant 1$. The right hand side of (5) is the objective function of (4) at π. Therefore, $U(\rho, \tilde{P}) \leqslant \overline{U}(P)$. Finally, since the mechanism ρ is IC, it follows that $U(\rho, P) \leqslant U(\rho, \tilde{P})$. Therefore, $U(\rho, P) \leqslant \overline{U}(P)$, as desired. □

When P is a singleton $\{(u,v)\}$, we also denote $\overline{U}(\{(u,v)\})$ by $\overline{U}(u,v)$. The following claim, which follows immediately from (1) and (3), explains the role of the function $\alpha^m(u,v)$ in our argument.

Claim 8.5. *When P is a singleton $\{(u,v)\}$, $\overline{U}(u,v) = \alpha^m(u,v)u$.*

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

Proof. Since ρ is IC, we have that $U(\rho, P) \geqslant U(\rho, \{(u,v)\}) = \alpha^m(u,v)u$ for every $(u,v) \in P$. □

Claim 8.4 bounds from above the agent's expected payoff in an IC mechanism which has the WCR of at most R^m. Claim 8.6 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 8.7. *$\underline{U}(P) \leqslant \overline{U}(P)$ for every P.*

Proof. The function $\overline{U}(P)$ defined in (4) is increasing in P. Therefore, from Claim 8.5 we have:

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

It follows that:

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

□

By definition, the mechanism ρ^m solves the following linear programming:

$$\rho^m(\cdot|P) \in \arg\max_{\pi} \sum_{(u,v)\in P} \pi(u,v)v \tag{6a}$$

$$\text{s.t.} \quad \pi(u,v) \geqslant 0,\ \forall (u,v) \in P, \tag{6b}$$

$$\sum_{(u,v)\in P} \pi(u,v) \leqslant 1, \tag{6c}$$

$$\sum_{(u,v)\in P} \pi(u,v)u = \underline{U}(P). \tag{6d}$$

It is possible that (6) 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 (6).

The following lemma is the core of the argument. It gives an equivalent characterization of the mechanism ρ^m.

Lemma 8.8. *The optimal solutions to* (6) *and those to the following problem coincide. Hence, $\rho^m(\cdot|P)$ is also given by the solution to the following problem:*

$$\rho(\cdot|P) \in \arg\max_{\pi} \sum_{(u,v)\in P} \pi(u,v)v \tag{7a}$$

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

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

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

$$\sum_{(u,v)\in P} \pi(u,v)u \leqslant \overline{U}(P). \tag{7e}$$

Proof of Lemma 8.8. We discuss two cases separately.

Case 1. Assume that there exists some $(u, v) \in P$ such that $v \geqslant 1 - R^m$. Consider the following linear programming which is a relaxation of both problem (6) and problem (7):

$$\rho(\cdot|P) \in \arg\max_{\pi} \sum_{(u,v)\in P} \pi(u,v)v \tag{8a}$$

$$\text{s.t.} \quad \pi(u,v) \geqslant 0,\ \forall (u,v) \in P, \tag{8b}$$

$$\sum_{(u,v)\in P} \pi(u,v) \leqslant 1, \tag{8c}$$

$$\sum_{(u,v)\in P} \pi(u,v)u \geqslant \underline{U}(P). \tag{8d}$$

We claim that the constraint (8d) holds with equality at every optimal solution. Indeed, if (8d) is not binding then an optimal solution to (8) is also an optimal solution to the following linear programming:

$$\rho(\cdot|P) \in \arg\max_{\pi} \sum_{(u,v)\in P} \pi(u,v)v \tag{9a}$$

$$\text{s.t.} \quad \pi(u,v) \geqslant 0,\ \forall (u,v) \in P, \tag{9b}$$

$$\sum_{(u,v)\in P} \pi(u,v) \leqslant 1, \tag{9c}$$

which is derived from (8) by removing (8d). Let $v_p = \max_{(u,v)\in P} v$ and $u_p = \max_{(u,v_p)\in P} u$. By the definition of α^m in (3), $\alpha^m(u_p, v_p) = 1$ given that $v_p \geqslant 1 - R^m$. Every optimal solution π^* to problem (9) satisfies $\text{support}(\pi^*) \subseteq \text{argmax}_{(u,v)\in P} v$, which implies that:

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

This implies that every optimal solution to (8) satisfies (8d) with equality, so it is a feasible point in both (6) and (7). Since problem (8) is a relaxation of both problem (6) and (7), the optimal values of (6), (7), and (8) coincide. Hence, every optimal solution to (6) or (7) is optimal in (8). This, combined with the fact that every optimal solution to (8) is optimal in (6) and (7), implies that the optimal solutions to (6) and (7) 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 (6) and (7) coincide. Given that $v < 1 - R^m$ for every $(u, v) \in P$, the constraint (4c) in problem (4) must be slack since if it is satisfied with an equality then (4d) is violated. Therefore, in this case $\overline{U}(P)$ also satisfies

$$\begin{aligned} \overline{U}(P) = \max_{\pi} \quad & \underline{u} + \sum_{(u,v)\in P} \pi(u,v)(u - \underline{u}) \\ \text{s.t.} \quad & \pi(u,v) \geqslant 0,\ \forall (u,v) \in P, \\ & \sum_{(u,v)\in P} \pi(u,v)(1-v) \leqslant R^m, \end{aligned} \tag{10}$$

which is derived from problem (4) by removing (4c). Problem (10) 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 8.5,

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

Therefore, by Claim 8.7 we get $\overline{U}(P) = \underline{U}(P)$, as desired.

□

We now show that, when the set of available projects is a singleton, the regret under the mechanism ρ^m is at most R^m.

Claim 8.9. *For every singleton $A = \{(u, v)\}$, the regret under ρ^m is at most R^m.*

Proof. In this case, ρ^m accepts with probability $\alpha^m(u, v)$ so 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 $\max(v(1-\bar{p}), (1-v)\gamma(u,\bar{p})) \leqslant R^m$. By (3), $\bar{p} \leqslant \alpha^m(u,v)$. Therefore, it also follows that $v(1-\alpha^m(u,v)) \leqslant v(1-\bar{p}) \leqslant R^m$. □

Claim 8.10. *The optimal value in problem* (7) *is at least* $\max_{(u,v)\in P} v - R^m$.

Proof. Since the constraints (7d) and (7e) cannot both be binding, it is sufficient to prove that the optimal value in the two problems derived from (7) by removing either (7d) or (7e) is at least $v_p - R^m$ where $v_p = \max_{(u,v)\in P} v$. Let $(u_p, v_p) \in P$ denote a principal's favorite project.

If we remove (7d), let π be given by $\pi(u_p, v_p) = \alpha^m(u_p, v_p)$ and $\pi(u,v) = 0$ when $(u,v) \neq (u_p, v_p)$. Then $\sum_{(u,v)\in P} \pi(u,v)u = \alpha^m(u_p,v_p)u_p \leqslant \underline{U}(P) \leqslant \overline{U}(P)$ so (7e) is satisfied. Also $v_p(1-\alpha^m(u_p,v_p)) \leqslant R^m$ by Claim 8.9, which implies that the value of the objective function in (7) at π is at least $v_p - R^m$, as desired.

If we remove (7e), let π be the optimal solution to (4) and let π' be the probability distribution over P such that $\pi'(u,v) = \pi(u,v)$ when $(u,v) \neq (u_p,v_p)$ and $\pi'(u_p,v_p) = 1 - \sum_{(u,v)\in P\setminus\{(u_p,v_p)\}} \pi(u,v)$, so π' is derived from π by allocating the probability of choosing no project to (u_p,v_p). Then

$$\sum_{(u,v)\in P} \pi'(u,v)u = u_p + \sum_{(u,v)\in P} \pi(u,v)(u-u_p) \geqslant \underline{u} + \sum_{(u,v)\in P} \pi(u,v)(u-\underline{u}) = \overline{U}(P) \geqslant \underline{U}(P),$$

where the last equality follows from the fact that π is optimal in (4). Therefore, π' satisfies (7d). Also

$$\sum_{(u,v)\in P} \pi'(v)(v_p - v) = \sum_{(u,v)\in P} \pi(v)(v_p - v) \leqslant \sum_{(u,v)\in P} \pi(v)(1-v) \leqslant R^m$$

where the last inequality follows from (4d), as desired. □

Proof of Theorem 5.2. 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 accepts (u,v) when the proposal is P. So, $\mathrm{RGRT}(\rho,P) =$

$(1-p)v$. 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 accepts (u,v) with probability at least $\gamma(u,p)$. So, $\mathrm{RGRT}(\rho,\tilde{P}) \geqslant (1-v)\gamma(u,p)$. Therefore $\mathrm{WCR}(\rho) \geqslant \max((1-p)v,(1-v)\gamma(u,p))$.

2. The mechanism ρ^m is IC, and it solves problem (7) by Lemma 8.8. By Claim 8.10, 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 π, 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 accepts a singleton proposal $\{(u,v)\}$. Then, ρ^m is not weakly dominated by ρ based on the following two claims:

   (a) If the agent's type A is a singleton $\{(u,v)\}$, then $\alpha(u,v) \leqslant \alpha^m(u,v)$ by claims 8.4 and 8.5. Hence, the principal's payoff is weakly higher under ρ^m than under ρ for singleton A.

   (b) Suppose that $\alpha(u,v) = \alpha^m(u,v)$ for every (u,v). Fix a proposal P and let $\pi = \rho(\cdot|P)$ so $U(\rho,P) = \sum_{(u,v)\in P} \pi(u,v)u$. Then, since ρ is IC, it follows from Claim 8.6 that $U(\rho,P) \geqslant \underline{U}(P)$, and, from Claim (8.4), that $U(\rho,P) \leqslant \overline{U}(P)$. Therefore π is a feasible point in problem (7). Since $\rho^m(\cdot|P)$ is the optimal solution to (7), 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 $U(\rho,P) \leqslant U(\rho^m,P)$ 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:

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

By definition, $U(\tilde{\rho}, P) = \min(U(\rho, P), U(\rho^m, P))$. The functions $U(P,\rho)$ and $U(P,\rho^m)$ are increasing in P since ρ and ρ^m are IC. Therefore $U(P,\tilde{\rho})$ 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 $U(\rho, P) \geqslant \underline{U}(P)$. From Claim 8.4 we know that $U(P,\rho) \leqslant \overline{U}(P)$ for every P. Therefore, $\rho(\cdot|P)$ is a feasible point in problem (7). It follows from Lemma 8.8 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 $U(\rho, P) \geqslant \underline{U}(P)$, $\tilde{\rho}$ 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 for every P such that $U(\rho, P) \geqslant \underline{U}(P)$, $\rho(\cdot|P) = \rho^m(\cdot|P)$, so $U(\rho, P)$ is equal to $\underline{U}(P)$. Hence, for every P, $U(\rho, P) \leqslant \underline{U}(P)$.

□

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.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics*, 18(2): 141–153.

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.

Kasberger, Bernhard, and Karl H Schlag. 2020. "Robust bidding in first-price auctions: How to bid without knowing what others are doing." *Available at SSRN 3044438.*

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. 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.

Milnor, John. 1954. "Games against nature, in "Decision Processes"(RM Thrall, CH Coombs, and RL Davis, Eds.)."

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): 829–848. Merger Control in International Markets.

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. Special Issue: Selected Papers, European Association for Research in Industrial Economics 37th Annual Conference, Istanbul, Turkey, September 2-4, 2010.

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.

Stoye, Jörg. 2011. "Axioms for minimax regret choice correspondences." *Journal of Economic Theory*, 146(6): 2226–2251.

Wald, Abraham. 1950. "Statistical decision functions."