MANAGING PERSUASION ROBUSTLY: THE OPTIMALITY OF QUOTA RULES

By

Dirk Bergemann, Tan Gan, and Yingkai Li

October 2023

COWLES FOUNDATION DISCUSSION PAPER NO. 2372

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Managing Persuasion Robustly: The Optimality of Quota Rules*

Dirk Bergemann[†] Tan Gan[‡] Yingkai Li[§]

October 19, 2023

Abstract

We study a sender-receiver model where the receiver can commit to a decision rule before the sender determines the information policy. The decision rule can depend on the signal structure and the signal realization that the sender adopts. This framework captures applications where a decision-maker (the receiver) solicit advice from an interested party (sender). In these applications, the receiver faces uncertainty regarding the sender's preferences and the set of feasible signal structures. Consequently, we adopt a unified robust analysis framework that includes max-min utility, min-max regret, and min-max approximation ratio as special cases. We show that it is optimal for the receiver to sacrifice ex-post optimality to perfectly align the sender's incentive. The optimal decision rule is a quota rule, i.e., the decision rule maximizes the receiver's ex-ante payoff subject to the constraint that the marginal distribution over actions adheres to a consistent quota, regardless of the sender's chosen signal structure.

Keywords— communication, commitment, min-max regret, quota rules

JEL— D47, D82, D83

*Dirk Bergemann acknowledges financial support from NSF SES 1948336 and 2049754 and Yingkai Li from Sloan Research Fellowship FG-2019-12378. We thank Songzi Du, Johannes Hörner, Emir Kamenica, Xiao Lin and Eran Shmaya for helpful comments and suggestions.

[†]Department of Economics, Yale University. Email: `dirk.bergemann@yale.edu`

[‡]Department of Economics, Yale University. Email: `tan.gan@yale.edu`

[§]Cowles Foundation for Research in Economics, Yale University. Email: `yingkai.li@yale.edu`

1 Introduction

1.1 Motivation

As a central topic of economic theory, the strategic communication between a sender and a receiver has been analyzed under various assumptions of commitment power. For instance, the cheap talk model (Crawford and Sobel (1982)) considers situations where neither the sender nor the receiver have commitment power. The delegation model (Holmstrom (1980)) conveys commitment power to the receiver. Lastly, in the persuasion model (Kamenica and Gentzkow (2011)), only the sender has commitment power that allows her to credibly disclose information.

We consider a model in which both the sender and the receiver have commitment power. Specifically, in contrast to the persuasion model, before the sender adopts a signal structure for conveying information, the receiver can commit to decision rules that depend on both the chosen signal structure and the realized signal. For instance, in a firm, the CEO (receiver) could request each department manager (sender) to devise an evaluation plan for the employees (signal structure) in the context of downsizing the company. The CEO can commit to a policy for terminating employees based on their evaluations, and given this policy, the manager can commit to an evaluation plan that maximizes their department's payoff. Another example that fits into our model is civil law, where the government holds commitment power over the defendant's sentencing based on the evidence presented by the prosecutor. Then the prosecutor can propose an investigation process (signal structure) to persuade the judge to convict the defendant.

In many applications, the receiver faces uncertainty regarding the sender's utility function and the set of available signal structures. Taking the example of downsizing a firm, the CEO is uncertain about the manager's preferred department size and the set of possible evaluation plans the manager can implement. In the presence of such uncertainty, we adopt the notion of regret, which captures the difference between the optimal payoff with known utility function and set of available signal structures and the actual payoff achieved by a specific decision rule without such knowledge. Specifically, we consider the robustness objective where the receiver commits to a decision rule that minimizes the worst-case regret over all possible state-independent utility functions and sets of available signal structures for the sender. Our analyses also extend to a more general notion of regret that incorporates weighted differences between the optimal payoff without uncertainty and the actual payoff, encompassing classic max-min utility and min-max approximation ratio as special cases.

Before the robust analysis, we first consider the classic Bayesian benchmark where the receiver knows the utility function and the set of available signal structures of the sender.

We show that in the Bayesian benchmark, the receiver can achieve his first-best payoff, i.e., the payoff achieved by the optimal signal structure and optimal action after any signal realization. Therefore, the regret, which is inevitable due to the receiver's ignorance about the model primitives, measures the difference between the maximum value of information and the actual payoff of the receiver.

To understand the trade-offs of the receiver in robust settings, we decompose the regret into efficiency loss and agency loss. The efficiency loss captures the loss stemming from sub-optimal actions committed by the receiver, while the agency loss captures the loss stemming from the sender's strategic choice of a sub-optimal signal structure.

In the robustness setting where the receiver is uncertain about the primitives, one naive decision rule the receiver can commit to is to take the myopic optimal action for all signal realizations, as what the receiver would do if he had no commitment power. Such decision rule completely eliminates the efficiency loss. However, the decision rule is independent of the sender's choice of signal structure and is extremely vulnerable to the sender's strategic persuasion. In fact, we show that the worst-case regret for the myopic optimal decision rule is so large that it equals the entire information value of the full revealing signal structures. In this setting, a more nature conjecture for the optimal decision rule is that it should maintain a balance between efficiency loss and agency loss based on the receiver's utility function and the prior belief.

Perhaps surprisingly, we show that it is optimal for the receiver to adopt another extreme decision rule that completely avoids agency loss: quota rules. A quota rule is associated with a quota, which is a distribution over the action space. For any chosen signal structure, the quota rule specifies a decision rule that maximizes the receiver's expected payoff subject to the constraint that the marginal distribution of actions must coincide with the quota. The quota rule sacrifices some efficiency loss to create an aligned incentive between the receiver and the sender. Under the quota rule, the sender is indifferent among all signal structures as they all lead to the same marginal distribution over actions and the sender has state-independent utility. Therefore, it is incentive-compatible to ask the sender to pick the receiver's favorite signal structure. Note that quota rules also have natural interpretations in applications. In the example of firm downsizing, quota rules correspond to the policies where the CEO commits to fire, for example, ten percent of the employees in each division regardless of the managers' evaluation plans.

We further characterize the optimal quota rules in the special case of binary actions. In this case, for any given quota, the signal structure that maximizes the regret corresponds to a binary monotone partition of the state space. This partition allows us to divide the analysis of regret into two subcases: left-biased error and right-biased error. These subcases capture the worst-case regret, where the quota assigns more ex ante probability to action 0 and 1, respectively, compared to optimal decision given the signal structure. We show that the left-biased error decreases as the quota increases, while the right-biased error increases with the quota. Consequently, the optimal quota rule strikes a balance between these two errors, achieving an interior solution that equalizes them.

The explicit characterization with binary actions helps us to perform several comparative analyses to better understand how the optimal quota is determined. We first show that the optimal quota on one action is increasing if the random reward of that action increases in the first-order stochastic order. This is intuitive since the receiver will always become more favorable to the action with higher rewards even with uncertainty over the primitives. In addition, we show that the optimal quota also exhibits responsiveness to changes in second-order stochastic dominance. In particular, we show that if the prior has a mean-preserving spread over states under which one action is optimal, the optimal quota for that action would decrease. Intuitively, this implies that the receiver should try to avoid an action when its relative advantage over other actions becomes more obscure. We highlight that the statement does not imply that the receiver should optimally avoid one action because it is more riskier.

1.2 Related Work

Our paper contributes to the literature on principal-expert models, see Zermeno (2011, 2012); Deb et al. (2018); Clark and Reggiani (2021); Li et al. (2022). In those models, the receiver, i.e., the principal, has commitment power on making decisions contingent on the information received from the sender. In contrast to our model, the literature on principal-expert models assumes that both the chosen information structure and the signals are private information of the sender, and the sender cannot commit to truthfully reveal the signals to the receiver. The question of identifying the robustly optimal mechanism has also been explored for the principal-expert models, see Carroll (2019); Chen and Yu (2021). In those papers, similar to our model, the principal is also ignorant of the set of information structures available to the sender. The objective of the principal is to design a mechanism that maximizes the worst case utility, which can be viewed as a special case of minimizing the generalized regret.

In the robustness setting, the performance measure we take is min-max regret. The objective of max-min regret has been micro-founded in Stoye (2011), and has been widely adopted in applications such as monopoly pricing (Bergemann and Schlag (2011)), monopoly regulation (Guo and Shmaya (2019)), contest design (Beviá and Corchón (2022)), project choice (Guo and Shmaya (2022)) and bandit learning (Slivkins et al. (2019)). The objective of minimizing worst-case generalized regret was recently proposed in Anunrojwong et al. (2023) for designing robust mechanisms in auction settings, which includes the objective of max-min utility (see Bergemann and Schlag (2011); Carroll (2017); Brooks and Du (2021)) and min-max approximation ratio (see Hartline and Lucier (2015); Allouah and Besbes (2020); Hartline et al. (2023) as special cases.

Our paper also relates to the literature on robust persuasion (see Dworczak and Pavan (2022); de Clippel and Zhang (2022); Babichenko et al. (2022)). The main difference is that in these papers, the receiver does not have commitment power and can only myopically best respond to the signals received from the sender. Moreover, the sender in these models faces uncertainty over the environments and tries to design robust mechanisms to maximize her worst case performance.

By viewing the information structure of the agent as an abstract observable action, our model shares similar structure as contract design with observable actions. Laux (2001); Zhao (2008); Chen (2012) consider the setting where both observable and unobservable actions are available to the agent, and the principal seeks to minimize the cost of incentivizing a given effort choice. Guo and Shmaya (2022) consider the project choice problem where the proposal of the agent can also be viewed as an abstract observable action. Similar to our paper, they also consider the robust setting where the set of actions available to the agent is unknown to the principal.

Quota mechanisms have been discussed in other settings as well. For example, it is discussed in cheap talk Chakraborty and Harbaugh (2007), in mechanism design with large samples Jackson and Sonnenschein (2007), and in the credibility of Bayesian persuasion Lin and Liu (2022).[1] Chakraborty and Harbaugh (2007) point out that quota rules can help to elicit information when the sender has state-independent utility. They consider multi-issue cheap-talk problems, and study equilibria where the Sender assigns a ranking to each issue. In such equilibria, a message is a complete or partial ordering of all the issues, and any on-path deviation is a different ordering that maintains the same distribution of rankings. In a mechanism design setting, Jackson and Sonnenschein (2007) show that the utility costs associated with incentive constraints become negligible when the decision problem is linked with a large number of independent copies of itself. This is established by defining a mechanism in which agents must budget their representations of preferences so that the frequency of preferences across problems mirrors the underlying distribution of preferences, and then arguing that agents' incentives are to satisfy their budget by being as truthful as possible. In Lin and Liu (2022), quota mechanisms are associated with the credibility of the sender's persuasion scheme. They define the sender's disclosure policy as credible

[1] Lin and Liu (2022) have a comprehensive literature review on quota mechanisms.

if the sender cannot profit from tampering with her messages while keeping the marginal distribution of messages unchanged.

2 Model

2.1 Payoff Environment

We consider a sender-receiver game with finite states $\theta \in \Theta$ and finite actions $a \in A$. The receiver (he) has state-dependent utility

$$u(\theta, a) : \Theta \times A \to [0, 1],$$

and the sender (she), on the other hand, has state-independent utility

$$v(a) : A \to [0, 1],$$

where the payoffs of both players are normalized to $[0, 1]$ without loss of generality.

Both the sender and the receiver have a common prior $\rho \in \Delta\Theta$ over the unknown state. Additionally, the sender can commit to a signal structure for sending informative signals to the receiver. Without loss of generality, we take a belief based approach and represent a signal structure π as a distribution over posteriors that satisfies Bayesian consistency. Formally, we denote the set of feasible signal structures as

$$\Sigma = \left\{ \pi \in \Delta(\Delta\Theta) \,\middle|\, \int_{\Delta\Theta} \mu \,\mathrm{d}\pi(\mu) = \rho \right\}.$$

We assume that the sender only has limited access to the set of feasible signal structures. That is, there exists a set:

$$\Pi \subseteq \Sigma$$

such that the sender can commit to signal structure π only if $\pi \in \Pi$.

In contrast to the literature on Bayesian persuasion (Kamenica and Gentzkow, 2011), we allow the receiver to have commitment power on the decision rules before receiving a signal from the sender. That is, instead of myopically best responding to the realized signal, the receiver can ex ante commit to a decision rule:

$$\alpha(\pi, \mu) : \Sigma \times \Delta\Theta \to \Delta A,$$

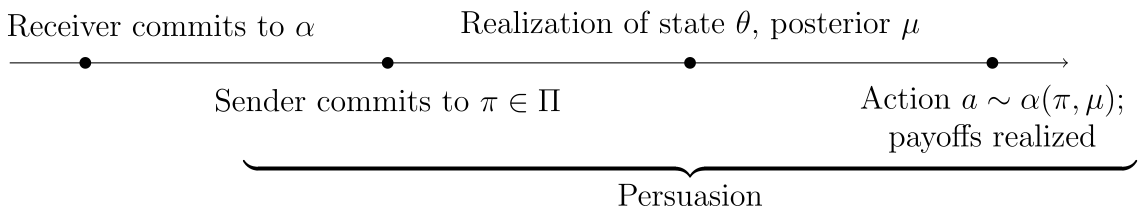

Figure 1: The timeline of the model.

which maps the signal structure that the sender chooses and the receiver's posterior belief after the signal realization to a distribution over actions. Let $\alpha(a|\pi, \mu)$ be the probability of action a being taken according to distribution $\alpha(\pi, \mu)$. To simplify the exposition, here we do not allow the receiver to elicit information about the utility function and the set of available signal structures from the sender. We show that this is in fact without loss of generality in Section 6.

The timing of the game is illustrated in Figure 1 and is formally described as follows:

1. The receiver publicly commits to a decision rule α.

2. The sender publicly commits to a signal structure $\pi \in \Pi$ after observing decision rule α.

3. State θ is realized according to prior ρ, and a signal that leads to posterior belief μ is sent to the receivers according to signal structure π.

4. Action a is chosen according to distribution $\alpha(\pi, \mu)$. Sender receives payoff $v(a)$ and receiver receives payoff $u(\theta, a)$.

In our model, we assume that the sender is a standard Bayesian decision-maker. That is, given decision rule α, she chooses

$$\pi^*(\alpha, v, \Pi) \in \arg\max_{\pi \in \Pi} \int_{\Delta\Theta} \sum_a v(a)\alpha(a|\pi, \mu)\, \mathrm{d}\pi(\mu).$$

to maximize her expected payoff when her utility function is v.[2] As customary in the literature, we assume a tie-breaking rule in favor of the receiver. Thus, letting $u(\mu, \alpha(\pi, \mu)) \triangleq \sum_a \sum_\theta u(\theta, a)\mu(\theta)\alpha(a|\pi, \mu)$, the receiver's expected payoff by committing to decision rule

[2]The maximization problem might not be well-defined (only) when Π is infinite. Our results hold under any behavioral assumption of the sender's choice of signal structures when the max does not exist.

α is

$$U(\alpha, v, \Pi) = \int_{\Delta\Theta} u(\mu, \alpha(\pi, \mu)) \, \mathrm{d}\pi(\mu), \quad \text{where} \quad \pi = \pi^*(\alpha, v, \Pi).$$

Therefore, the optimal payoff of the receiver in this sender-receiver game given utility v and set Π is

$$U^*(v, \Pi) = \sup_{\alpha} U(\alpha, v, \Pi).$$

2.2 Robustness and Regret

Instead of assuming that the receiver is perfectly informed about the utility function v and the set Π of available signal structures of the sender, we consider a robust setting where both and v and Π are not known to the receiver. Therefore, the receiver cannot design decision rules contingent on the true utility v or set Π. For any decision rule α, the regret of the receiver due to his ignorance of utility v and set Π is

$$R(\alpha, v, \Pi) = U^*(v, \Pi) - U(\alpha, v, \Pi). \tag{1}$$

In this robust setting, the goal of the receiver is to design a decision rule that minimizes the worst-case regret

$$R \triangleq \inf_{\alpha} R(\alpha) \triangleq \inf_{\alpha} \sup_{v, \Pi \subseteq \Sigma} R(\alpha, v, \Pi).$$

Generalized Regret For the purpose of delivering intuition, we emphasize regret minimization in this paper. However, we establish our results for the γ-generalized regret. Formally, for any $\gamma \in [0, 1)$, the γ-generalized regret of the receiver is

$$R_{\gamma}(\alpha, v, \Pi) = \gamma \cdot U^*(v, \Pi) - (1 - \gamma) \cdot U(\alpha, v, \Pi). \tag{2}$$

The objective of regret minimization is a special case of γ-generalized regret when $\gamma = \frac{1}{2}$. By varying the parameter γ, we can recover other robust frameworks commonly studied in the literature. For example, when $\gamma = 0$, the objective corresponds to the max-min framework. Moreover, there exists $\hat{\gamma} \in (0, 1/2)$, such that $\hat{\gamma}$-generalized regret coincides with maximizing the worst-case approximation ratio, as discussed in Anunrojwong et al. (2023).

General Decision Rules To simplify the exposition, we assume that the receiver cannot elicit information about (v, Π) from the sender and can only commit to decision rules as a function of the chosen signal structure and realized posterior belief. In Section 6, we show that this is without loss of generality. Even when the decision rule could also depend on the report of the sender about her primitives (v, Π), it is without loss to focus on decision rules that do not depend on it.

3 Managing Bayesian Persuasion

In this section, we consider the Bayesian setting where the utility function and the set of available signal structures of the sender are known to the receiver. We show that in this benchmark, the receiver can implement the first best decision rule via a quota rule.

For any utility function v and set Π of available signal structures of the sender, a naive choice of the receiver is to commit to the decision rule that that maximizes the receiver's expected payoff for any realized posterior regardless of the sender's choice of the experiment (as it is payoff irrelevant):

$$a^*(\mu) \in \arg\max_a u(\mu, a).$$

Ideally, if the sender is not strategic and chooses a signal structure π_F that maximizes the receiver's payoff, the receiver would receive his first-best payoff as

$$U_{\mathrm{F}}(\Pi) = \sup_{\pi\in\Pi}\ u^*(\pi) \triangleq \sup_{\pi\in\Pi} \int_{\Delta\Theta} u(\mu, a^*(\mu))\,\mathrm{d}\pi(\mu).$$

We show that the receiver can carefully design decision rules that attain the first best as the equilibrium payoff when he knows both utility v and the set Π of available signal structures. Specifically, we consider a special class of decision rules called the quota rules.

Definition 1 (Quota Rules)**.**
A decision rule α is a quota rule *with quota $q \in \Delta A$ if for any signal structure π,*[3]

$$\alpha(\pi, \cdot) \in \arg\max_{\hat\alpha\in\Psi(q)} \int_{\Delta\Theta} u(\mu, \hat\alpha(\mu))\,\mathrm{d}\pi(\mu),$$
$$\text{where}\quad \Psi(q) \triangleq \left\{\hat\alpha : \Delta\Theta \to \Delta(A) \,\middle|\, \int_{\Delta\Theta} \hat\alpha(a|\mu)\,\mathrm{d}\pi(\mu) = q(a),\ \forall a\right\}.$$

We also use q to represent the quota rule with quota q without ambiguity. Essentially, the quota rule specifies a quota q as a constraint on the marginal distribution over actions

[3]The existence of the optimal decision rule for this optimization problem is implied by Lemma 1 in later sections.

for the decision rule given any chosen signal structure π. For any π, the quota rule is the optimal decision rule that maximizes the receiver's expected payoff subject to the quota constraints.

Proposition 1 (First Best Implementation)**.**
For any utility v and any set Π of available signal structures, there exists a quota rule q_Π that guarantees the first-best payoff for the receiver:

$$U^*(v,\Pi) = U(q_\Pi, v, \Pi) = U_{\mathrm{F}}(\Pi).$$

Note that the incentives of the sender with state-independent utilities are completely eliminated under quota rules since the distribution over actions for any quota rule is invariant of the chosen signal structure. Therefore, it is incentive-compatible to let the sender choose the signal structure that maximizes the expected payoff of the receiver based on the quota rule. To achieve the first-best payoff, the receiver only needs to commit to a quota rule that is consistent with the marginal distribution over first-best actions of the first-best signal structure in Π. That is, for any action a,

$$q^*_\Pi(a) = \int_{\Delta\Theta} \mathbf{1}_{a=a^*(\mu)}\,\mathrm{d}\hat{\pi}(\mu), \quad \text{where } \hat{\pi} \in \arg\max_{\pi\in\Pi} u^*(\pi).$$

4 Regret Minimizing Decision Rule

In this section, we prove the optimality of quota rules for $\gamma-$generalized regret given any $\gamma \in [0,1)$. In Section 4.1, we show that the loss given any decision rule can be decoupled as the efficiency loss and the agency loss, and the quota rule is the decision rule that completely eliminates agency loss with the sacrifice of the efficiency loss. To establish its optimality, we proceed in two steps. For any $\gamma \in [0,1)$, Section 4.2 characterizes the optimal quota q^*_γ among all quota rules, and Section 4.3 shows that quota rule q^*_γ is optimal among all decision rules for $\gamma-$generalized regret.

4.1 Effiency Loss and Agency Loss

In Proposition 1 we have shown that, when (v,Π) is known, the receiver can achieve his first-best payoff by committing to quota rule q^*_Π. Therefore, the regret of the receiver defined in (1) given decision rule α simplifies to

$$R(\alpha, v, \Pi) = U_{\mathrm{F}}(\Pi) - U(\alpha, v, \Pi).$$

Moreover, the first best quota q^*_Π varies across different Π, and therefore the receiver can not achieve the first-best payoff in the presence of uncertainty about (v, Π). To have some basic intuition about the trade-offs of the design problem, denote the receiver's actual payoff under decision rule α given signal structure π as

$$U(\alpha, \pi) = \int_{\Delta\Theta} u(\mu, \alpha(\pi, \mu)) \, \mathrm{d}\pi(\mu). \tag{3}$$

Letting $\pi_F \in \arg\sup_{\pi\in\Pi} \; u^*(\pi)$ be a signal structure that maximizes the receiver's first best payoff, we can further decompose the regret of the receiver as

$$R(\alpha, v, \Pi) = \Big[u^*(\pi_F) - U(\alpha, \pi_F)\Big] + \Big[U(\alpha, \pi_F) - U(\alpha, \pi^*(\alpha, v, \Pi))\Big].$$

This equation decomposes the regret function into *efficiency loss* and *agency loss*. In the first difference, the signal structures are the same (π_F) but the decision rules are different. It captures the efficiency loss stemming from sub-optimal actions $\alpha(\pi_F, \cdot)$. In the second difference, the decision rule are the same (α) but the signal structures are different. It captures the agency loss stemming from the sender's strategic choice of a sub-optimal signal structure $\pi^*(\alpha, v, \Pi)$.

Under this view, there are two extremal decision rules. The first one is to completely eliminate the efficiency loss by choosing the myopic best action. To see why this mechanism is sub-optimal, suppose the receiver chooses the myopic optimal decision rule α_0 such that

$$\operatorname{supp}(\alpha_0(\pi, \mu)) \subseteq \arg\max_a u(\mu, a),$$

for any signal structure π and posterior belief μ where $\operatorname{supp}(\cdot)$ is the support of a distribution. Denote $a_0 = \arg\max_a u(\rho, a)$ as the optimal action under the prior, which is generically unique. Then, (one of) the worst-case is (Π_0, v_0) where

$$\Pi_0 = \{\text{Complete Information, Zero Information}\}, \quad v_0(a_0) > v_0(a) \; \forall a \neq a_0.$$

In this case, the sender either chooses to completely reveal the state and let the receiver take the first best action or provides zero information to induce a_0 with probability 1. Because the sender with preference v_0 strictly prefers action a_0 over any other actions, she chooses No Information and the receiver suffers a regret of

$$R(\alpha_0, v_0, \Pi_0) = \int_\Theta \max_a u(\theta, a) \, \mathrm{d}\rho(\theta) - u(\rho, a_0).$$

The receiver suffers from a large regret because he could have enjoyed the first best value if there was no agency problem, but ends up getting no information from the sender.

The second class of extremal mechanisms consists of quota rules (Definition 1), which completely eliminate any agency loss but result in efficiency loss. One might conjecture that the optimal design should maintain a balance between efficiency loss and agency loss. Therefore, it is not immediately clear whether these extremal mechanisms would be optimal.

4.2 Optimal Quota Rules

The choice of the optimal quota $q \in \Delta A$ itself is a non-trivial problem. We denote the $\gamma-$generalized regret of quota rule q given signal structure π as

$$R_\gamma(q,\pi) = \gamma \cdot u^*(\pi) - (1-\gamma) \cdot U(q,\pi),$$

where recall by (3), $U(q,\pi) = \max_{\alpha \in \Psi(q)} \int_{\Delta\Theta} u(\mu,\alpha(\mu))\,\mathrm{d}\pi(\mu)$. Note that the optimization problem of $U(q,\pi)$ can be expressed as an optimal transport problem, where instead of designing decision rule α, the receiver designs the joint distribution F over $\Delta\Theta \times A$ subject to two marginal constraints:

$$\begin{aligned} U(q,\pi) = \max_F &\int_{\Delta\Theta\times A} u(\mu,a)\,\mathrm{d}F(\mu,a), \\ \text{s.t. } & F(\Delta\Theta,a) = q(a), \quad \forall a \in A, \\ & F(N,A) = \pi(N), \quad \forall N \in \mathcal{B}(\Delta\Theta). \end{aligned}$$

We endow the space of $\Delta(\Delta\Theta)$ with the Wasserstein metric d:

$$\begin{aligned} d(\pi_1,\pi_2) = \min_{G\in\Delta(\Delta\Theta\times\Delta\Theta)} &\int_{\Delta\Theta\times\Delta\Theta} |\mu-\nu|_1\,\mathrm{d}G(\mu,\nu), \\ \text{s.t. } & G(N,\Delta\Theta) = \pi_1(N), \quad \forall N \in \mathcal{B}(\Delta\Theta), \\ & G(\Delta\Theta,N) = \pi_2(N), \quad \forall N \in \mathcal{B}(\Delta\Theta), \end{aligned}$$

where $||_1$ refers to the L_1 norm in Euclidean space. Because $\Delta\Theta$ is bounded, the Wasserstein metric induces the vague topology.[4] The vague topology is defined over probability space ΔX. A sequence of measure μ_n converges to μ_0 in the vague topology if and only if for any continuous function $f \in C(X)$, $\int_X f\,\mathrm{d}\mu_n \to \int_X f\,\mathrm{d}\mu$. As a note, the introduction of the Wasserstein metric is purely instrumental. It is used in a couple of following lemmas which help us establish Theorem 1, but our primitives and Theorem 1 themselves do not contain any assumptions on the choice of metric.

[4]See Theorem 6.9 in Villani (2016).

Lemma 1 (Lipschitz Continuity).
The optimal solution for $U(q, \pi)$ exists and hence $R_\gamma(q, \pi)$ is well-defined. Moreover, $R_\gamma(q, \pi)$ is continuous in (q, π), and

$$|R_\gamma(q, \pi') - R_\gamma(q, \pi)| \leq d(\pi, \pi'), \quad \forall q, \pi, \pi'.$$

Now that $R_\gamma(q, \pi)$ is well-defined and continuous, and so the the following concepts are also well-defined since $\Delta\Theta$ is compact under the vague topology:

$$R_\gamma(q) = \max_{\pi \in \Sigma} R_\gamma(q, \pi),$$
$$\Sigma_q = \arg\max_{\pi \in \Sigma} R_\gamma(q, \pi),$$
$$q_\gamma^* = \arg\min_{q \in \Delta A} R_\gamma(q).$$

In the definition of $R_\gamma(q)$, we only consider worst cases $\Pi = \{\pi\}$ that are singletons. This is without loss of generality since under the quota rule the sender always chooses the receiver's favorite signal structure. The set Σ_q, which is closed and hence compact, is the set of worst-case signal structures for quota rule q. We have the following characterization for optimal quota rules.

Lemma 2 (Local Optimality).
For any $\gamma \in [0, 1)$, a quota rule q is the optimal quota rule if and only if there does not exist another quota rule q' such that

$$\max_{\pi \in \Sigma_q} R_\gamma(q', \pi) < R_\gamma(q).$$

Considering the optimal quota rule q_γ^* and its corresponding worst case signal structures $\Sigma_{q_\gamma^*}$, Lemma 2 implies that there does not exist another quota rule q' which can uniformly and strictly reduce the regret on $\Sigma_{q_\gamma^*}$ compared to q_γ^*. Intuitively, if such q' exists, one can slightly adjust the quota q_γ^* towards q'. This modification can lower the regret at all worst cases in $\Sigma_{q_\gamma^*}$ without making significant impact on the regret of other signal structures that are originally slack. As a result, the overall worst-case regret would be lower compared to q_γ^*.

4.3 The Optimality of the Quota Rule

After establishing some basic understanding of the optimal quota within the class of quota rules, we are ready to prove the optimality of quota rules. We start by introducing generalized quota rules.

Definition 2 (Generalized Quota Rules).
The decision rule α is a generalized quota rule *if there exists a function $\bar{q} : \Sigma \to \Delta A$ such that for any signal structure π,*

$$\alpha(\pi, \mu) \in \underset{\alpha' \in \Psi(\bar{q}(\pi))}{\arg\max} \int_{\Delta\Theta} u(\mu, \alpha'(\mu))\, \mathrm{d}\pi(\mu).$$

Intuitively, the generalized quota rule specifies a quota on the marginal distribution over actions separately for each signal structure. Given the signal structure π and the quota $\bar{q}(\pi)$, the decision rule α is the optimal solution that maximizes the receiver's expected payoff subject to the quota constraints.

Because the sender's preference is state-independent, any two decision rules, which lead to the same marginal distribution over actions in any signal structure, are equivalent as far as the sender's incentive is concerned. Thus, the receiver always benefits from choosing the optimal decision rule that respects the quota.

Lemma 3 (Sufficiency of Generalized Quota Rules). *It is without loss of generality to focus on generalized quota rules.*

Now we are ready to prove Theorem 1. We present the full proof in the main text because the argument is considerably simplified by the fact that the mechanism does not elicit information about (Π, v) from the sender, yet the fundamental idea is the same as the proof under more general mechanisms as in Theorem 2 in Section 6, whose proof we relegate to the appendix.

Theorem 1 (Optimality of Quota Rules).
For any $\gamma \in [0, 1)$, the quota rule q^*_γ is optimal for the receiver.

Proof. We first introduce several notations. Let π^N be the information structure that reveals no information. Recall that by definition, $U(q, \pi)$ is monotonously increasing in the Blackwell order of π and

$$R_\gamma(q, \pi) = \gamma \cdot u^*(\pi) - (1 - \gamma) \cdot U(q, \pi).$$

We prove the theorem by contradiction. Suppose there is a decision rule that induces strictly less regret than the quota rule q^*_γ. By Lemma 3, it is without loss of optimality to assume this improvement is attained by a generalized quota rule $\bar{q}$, which further implies that

$$R_\gamma(\bar{q}(\pi), \pi) < R_\gamma(q_\gamma^*), \quad \forall \pi \in \Sigma.$$

If $\Sigma_{q_\gamma^*}$ is a singleton, denote it as $\Sigma_{q_\gamma^*} = \{\bar{\pi}\}$, then by Lemma 2,

$$R_\gamma(q_\gamma^*) \leq \max_{\pi \in \Sigma_{q_\gamma^*}} R_\gamma(q', \pi) = R_\gamma(q', \bar{\pi}), \quad \forall q'.$$

This contradicts with $\bar{q}$ being a strict improvement.

If $\Sigma_{q_\gamma^*}$ is not a singleton, then by Lemma 2, $\bar{q}(\pi^N)$ can not uniformly improve q_γ^*. Thus, there exists an information structure $\pi' \in \Sigma_{q_\gamma^*}$ such that

$$R_\gamma(\bar{q}(\pi^N), \pi') \geq R_\gamma(q_\gamma^*).$$

This implies that $U(\bar{q}(\pi^N), \pi') \leq U(q_\gamma^*, \pi')$ and $\bar{q}(\pi') \neq \bar{q}(\pi^N)$ since $\bar{q}$ is a strict improvement. In this case, when the set of information structures available to the sender is $\{\pi', \pi^N\}$, there exists a utility v of the sender such that the sender prefers $\bar{q}(\pi^N)$ over $\bar{q}(\pi')$ so she chooses π^N. The regret of the receiver is

$$\begin{aligned}
&\gamma \cdot u^*(\pi') - (1-\gamma) \cdot U(\bar{q}(\pi^N), \pi^N) \\
&\geq \gamma \cdot u^*(\pi') - (1-\gamma) \cdot U(\bar{q}(\pi^N), \pi') \\
&\geq \gamma \cdot u^*(\pi') - (1-\gamma) \cdot U(q_\gamma^*, \pi') = R_\gamma(q_\gamma^*),
\end{aligned}$$

a contradiction to the assumption that $\bar{q}$ is a strict improvement. □

5 Binary Actions

In Section 4 we proved the optimality of quota rules but did not determine which quota rule is optimal within this class. To deliver more insight on how model primitives determine the optimal quota in robust environments, we characterize the optimal quota rules in a binary action model.

5.1 Optimal Binary Quota Rules

In a binary-action model, the receiver's preference for each state is summarized by its expected utility difference between action 0 and action 1. Therefore, it is without loss of generality to assume that $\Theta \subseteq [-1,1]$. Moreover,

$$u(\theta, a) = \begin{cases} \theta, & a = 1; \\ 0, & a = 0. \end{cases}$$

Let $\Theta_0 \subseteq \Theta$ be the set of states such that $\theta < 0$ and $\Theta_1 \subseteq \Theta$ be the set of states such that $\theta > 0$. We assume there exists θ, θ' in the support of ρ such that $\theta \in \Theta_0$ and $\theta' \in \Theta_1$ (otherwise the information is without value).

We first characterize the worst case signal structure given a quota rule.

Definition 3 (Binary Monotone Partition)**.**
An signal structure π is a binary monotone partition *if the signal space is binary $\{0,1\}$, and letting $\hat{\Theta}_s$ be the set of states that send signal s with strictly positive probability, either one of the sets $\hat{\Theta}_0, \hat{\Theta}_1$ is empty, or for any $\theta \in \hat{\Theta}_0$ and $\theta' \in \hat{\Theta}_1$, we have $\theta \leq \theta'$.*

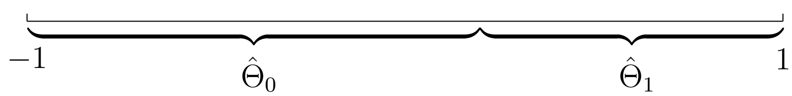

Figure 2: An illustration of binary monotone partition.

Intuitively, a binary monotone partition signal structure sends signal 0 for all low states and signal 1 for all high states. This is illustrated in Figure 2 for state space $\Theta = [-1, 1]$. In the case where the state distribution is discrete, there may exist a middle type such that the signal structure randomizes the signals for this type. In the binary action environment we can identify the quota rule q simply by the probability that action $a = 1$ is chosen:

$$q = Pr(a = 1).$$

Lemma 4 (Worst Case Signal Structure)**.**
For any $\gamma \in [0, 1)$ and any quota q, there exists a binary monotone partition signal structure that maximizes the principal's regret.

In a binary action model, the regret of the receiver in any binary monotone partition signal structure can be classified into two categories, the left-biased error and the right-biased error. The classification depends on whether p_1, the probability of signal $s = 1$, is larger than the quota q. Specifically, for any quota rule $q = Pr(a = 1)$, the regret given a signal structure with signal probabilities p_1 is called a left-biased error if $q \leq p_1$ and a right-biased error if $q \geq p_1$. This is illustrated in Figure 3. In the left-biased error,

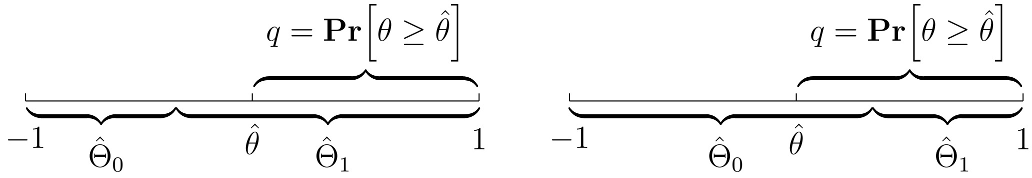

(a) left-biased error

(b) right-biased error

Figure 3: An illustration of two types of errors.

$q \leq p_1$ and hence the receiver has to take action 0 with a higher probability than in the first best solution. Consequently, he takes action 0 for sure when he observes signal 0, while mixing his action between 0 and 1 when he observes signal 1, in order to satisfy the quota restriction. Similarly, in the right-biased error, $q \geq p_1$ and hence the receiver has to take action 1 with a higher probability than in the first best solution. Consequently, he then chooses action 1 for sure when he observes signal 1, while mixing his action between 0 and 1 when he observes signal 0.

For any quota rule q, let $\Sigma_L(q)$ be the set of binary monotone partition signal structures with left-biased errors. The worst case left-biased error is

$$L_\gamma(q) \triangleq \max_{\pi \in \Sigma_L(q)} R_\gamma(q, \pi).$$

Similarly, $\Sigma_R(q)$ is the set of binary monotone partition signal structures with right-biased errors and the worst case right-biased error is

$$R_\gamma(q) \triangleq \max_{\pi \in \Sigma_R(q)} R_\gamma(q, \pi).$$

Let F be the cumulative distribution function of the prior ρ and we define $F^{-1}(p) \triangleq \inf\{\theta \mid F(\theta) \geq p\}$. Note that for any binary monotone partition signal structure π, letting p_0 be the probability of the low signal in π, we have

$$R_\gamma(q, \pi) = \left(\gamma - (1-\gamma) \cdot \frac{1 - p_0 - q}{1 - p_0}\right) \cdot \int_{p_0}^{1} F^{-1}(p)\, \mathrm{d}p,$$

if signal structure π is left-biased for quota q, and

$$R_\gamma(q, \pi) = (2\gamma - 1) \cdot \int_{p_0}^{1} F^{-1}(p)\, \mathrm{d}p - (1-\gamma) \cdot \frac{p_0 - 1 + q}{p_0} \cdot \int_{0}^{p_0} F^{-1}(p)\, \mathrm{d}p,$$

if signal structure π is right-biased for quota q. Moreover, let

$$z_1 = \sup\left\{ z \,\middle|\, \int_0^z F^{-1}(p)\,\mathrm{d}p \le 0 \right\},$$
$$z_0 = \inf\left\{ z \,\middle|\, \int_z^1 F^{-1}(p)\,\mathrm{d}p \ge 0 \right\}.$$

Since signal structure π is left-biased for quota rule q if $p_0 \in [z_0, z_1]$ and $p_0 \le 1-q$, and is right-biased for quota rule q if $p_0 \in [z_0, z_1]$ and $p_0 \ge 1-q$, the hence the right-biased error and the right-biased error are

$$L_\gamma(q) = \max_{p_0 \in [z_0, z_1],\, p_0 \le 1-q} \left(\gamma - (1-\gamma)\cdot \frac{1-p_0-q}{1-p_0}\right)\cdot \int_{p_0}^1 F^{-1}(p)\,\mathrm{d}p,$$
$$R_\gamma(q) = \max_{p_0 \in [z_0, z_1],\, p_0 \ge 1-q} (2\gamma-1)\cdot \int_{p_0}^1 F^{-1}(p)\,\mathrm{d}p - (1-\gamma)\cdot \frac{p_0-1+q}{p_0}\cdot \int_0^{p_0} F^{-1}(p)\,\mathrm{d}p.$$

Let m_ρ be the prior mean and m^* be the posterior mean such that the receiver is indifferent between two actions.

Proposition 2 (Optimal Quota Rules)**.**
For any $\gamma \in [0,1)$ and any prior ρ, the left-biased error is weakly decreasing in quota q and the right-biased error is weakly increasing in quota q, and inequalities are strict if $m_\rho \ne m^*$. The optimal quota rule q satisfies $L_\gamma(q) = R_\gamma(q)$.

Corollary 1 (Interior Optimal Quota Rule)**.**
The optimal quota is in the interior for $\gamma \in (0,1)$. The optimal quota is unique if $m_\rho \ne m^*$.

The optimal quota is in the interior except when $\gamma = 0$. In this case, our model is equivalent to the max-min framework where the receiver evaluates the decision rule according to his equilibrium payoff in the worst case. Thus, the optimal quota is simply choosing the optimal action according to his prior ρ. When $m_\rho \ne m^*$, any action is optimal according to the prior.

5.2 Comparative Statics

In general it is hard to have a closed-form solution for the optimal quota rule even with the nice characterization of the worst-case signal structure. To have more insights on what determines the optimal quota rule, we perform two comparative analysis with respect to the prior ρ, which is only relevant model primitive. Note that ρ is a joint representation of the receiver's utility function and the common prior about the state, so one can have different interpretations.

Our first result shows that the optimal quota on one action is increasing if the random reward of that action increases in the first-order stochastic order. This is intuitive since the receiver will always become more favorable to the action with higher rewards even with uncertainty over the primitives.

Proposition 3 (First-Order Stochastic Dominance)**.**
When $\gamma = 1/2$, the optimal quota under ρ is weakly lower than the optimal quota under $\hat{\rho}$ if $\hat{\rho}$ first-order stochastically dominates ρ.

We believe the quantifier $\gamma = 1/2$ is not necessary but the formal proof is currently missing. In addition, we show that the optimal quota also exhibits responsiveness to changes in second-order stochastic dominance. To present the result, let us first introduce the following notion.

Definition 4 (Mean-Preserving Spread)**.**
A distribution ρ is a mean-preserving spread of ρ' in Θ_i for $i \in \{0, 1\}$ if $\rho(z) = \rho'(z)$ for any $z \subseteq \Theta_{1-i}$ and ρ is a mean-preserving spread of ρ'.

Proposition 4 (Second-Order Stochastic Dominance)**.**
The optimal quota under ρ is weakly lower (higher) than the optimal quota under ρ' if ρ is a mean-preserving spread of ρ' in Θ_1 (Θ_0).

Intuitively, Proposition 4 implies that the receiver should try to avoid an action when its relative advantage over other actions becomes more obscure. Note that this does not mean the receiver should avoid an action if it is riskier. To see this, note that θ is the payoff difference of the two actions. Now, suppose we fix the common prior and fix the payoff of action 1 while change the payoff of action 0 to make it riskier in Θ_1 (in states where action 1 is better). This would lead to a mean-preserving spread of the induced ρ in Θ_1. Thus, Proposition 4 predicts the optimal quota for action 0 increases, although the reward of action 0 is riskier in the normal sense.

6 Extensions

In this section, we extend the optimality of quota rules to a more general environment. The extension includes two parts: (i) relaxing the set of possible signal structures, and (2) allowing general mechanisms for eliciting the sender's private information.

Possible Signal Structures In the main specification, we assume the set of all possible signal structures is

$$\Sigma = \left\{ \pi \in \Delta(\Delta\Theta) \,\middle|\, \int_{\Delta\Theta} \mu \,\mathrm{d}\pi(\mu) = \rho \right\},$$

and the set of signal structure Π can be any subset of Σ. Namely, the receiver thinks any signal structure is possible. In this section, we assume Π can be any closed subset of $\Sigma_0 \neq \Sigma$, where Σ_0 has the following property.

Definition 5 (Minimal Element)**.**
The set Σ_0 has a minimal element π_0, if $\pi_0 \in \Sigma_0$ and any $\pi \in \Sigma_0$ Blackwell dominates π_0.

In the main specification, where $\Sigma_0 = \Sigma$, the minimal element is the zero information π^N. We interpret the existence of the minimal element π_0 as there being common knowledge about a minimal source of information that the sender could get. What remains uncertain to the receiver is how much additional information the sender could access to.

General Mechanisms In the main specification, we consider decision rules

$$\alpha(\pi, \mu) : \Sigma \times \Delta\Theta \to \Delta A.$$

That is, the action taken by the receiver only depends on the signal structure chosen by the sender and the signal realization. In this subsection, we consider the general mechanism (π, α) with a closed subset $\Sigma_0 \subseteq \Sigma$ that has a minimal element π_0. In the general mechanism, the sender first reports the set of available signal structures $\Pi \in 2^{\Sigma_0}$ and her utility $v \in V$, and based on the report, the mechanism chooses one signal structure $\pi \in \Pi$ and the corresponding decision rule α. Specifically,

$$\pi(\Pi, v) : 2^{\Sigma_0} \times V \to \Sigma_0,$$
$$\alpha(\Pi, v, \mu) : 2^{\Sigma_0} \times V \times \Delta\Theta \to \Delta A.$$

We assume the sender's report on the set of available signal structures $\hat{\Pi}$ must be authentic in the sense that $\hat{\Pi} \subseteq \Pi$. Invoking the revelation principle, the design problem of the receiver is to minimize regret subjective to incentive constraints.

$$\inf_{\pi(),\alpha()} \sup_{v,\Pi\subseteq\Sigma_0} R_\gamma(\pi, \alpha, v, \Pi) = \gamma \cdot U^*(v, \Pi) - (1-\gamma) \cdot U(\alpha(\Pi, v), \pi(\Pi, v)),$$
$$\text{s.t.} \quad (\Pi, v) \in \arg\max_{\hat{v},\hat{\Pi}\subseteq\Pi} \int_{\Delta\Theta} \sum_a v(a)\alpha(a|\hat{\Pi}, \hat{v}, \mu) \,\mathrm{d}\pi(\hat{\Pi}, \hat{v})(\mu).$$

To introduce a natural definition of the quota rule in the general environment, recall in the main specification $\alpha^q(\pi, \mu)$ is called a quota rule with quota q if

$$\alpha^q(\pi, \cdot) = \arg\max_{\alpha' \in \Psi(q)} \int_{\Delta\Theta} u(\mu, \alpha'(\mu)) \, \mathrm{d}\pi(\mu), \quad \forall \pi.$$

Definition 6 (Quota Rule for General Mechanisms).
The general mechanism (π, α) *is a* quota rule *with* $q \in \Delta A$*, if for any* (Π, v),

$$\begin{aligned} \pi(\Pi, v, \mu) &\in \arg\max_{\pi \in \Pi} \int_{\Delta\Theta} u(\mu, \alpha^q(\pi, \mu)) \, \mathrm{d}\pi(\mu), \\ \alpha(\Pi, v, \mu) &= \alpha^q(\pi(\Pi, v, \mu), \mu). \end{aligned}$$

Namely, for any signal structure π that is eventually chosen, the receiver commits to the same quota rule $\alpha^q(\pi, \mu)$ as described in the main specification. For any reported Π, the mechanism chooses a signal structure π that maximizes the receiver's payoff, where π is evaluated by its performance in the quota rule $\alpha^q(\pi, \mu)$. This general mechanism can be implemented by the quota rule in the main specification because the sender can directly choose $\pi(\Pi, v)$ on behalf of the receiver without reporting (Π, v).

Theorem 2 (General Optimality of Quota Rules).
The quota rule q^*_γ *is the optimal general mechanism, where*

$$q^*_\gamma = \arg\min_{q \in \Delta A} \max_{\pi \in \Sigma_0} R_\gamma(q, \pi).$$

The proof of Theorem 2 is more involved than that of Theorem 1 but the basic intuition is the same, so we relegate it to the appendix. Note that if $\Sigma_0 = \Sigma$, then the optimal mechanism in both theorems are equivalent. The fact that the optimal quota rule q^*_γ varies as the possible set of signal structures Σ_0 changes shall not be surprising. For example, if $\Sigma_0 = \{\pi_0\}$, i.e., the receiver is sure that there is only one possible signal structure, then clearly the optimal quota rule is the first-bests quota $q^*_{\{\pi_0\}}$ as in Proposition 1.

7 Conclusion

We considered strategic communication between a sender and a receiver when both sides had the ability to commit to a policy, an information rule and a decision rule, respectively. We were interested in finding the optimal decision rule for the sender when he faced uncertainty about the objective and the instruments of the sender. In particular, the receiver was uncertain about the state independent preferences and the set of signal structures that are feasible for the sender. The literature on strategic communication has typically considered forms of communication where at least one of the players had no ability to commit, as illustrated in Table 1. In the current paper, we offer both sender and receiver some level of commitment power.

Sender \ Receiver	No Commitment	Commitment
No Commitment	Cheap Talk (Crawford and Sobel, 1982)	Constrained Delegation (Holmstrom, 1980)
Commitment	Bayesian Persuasion (Kamenica and Gentzkow, 2011)	this paper

Table 1: Four Forms of Strategic Communication

Our analysis proposes a robust decision rule to deal with the uncertainty about the sender's state-independent preference and set of feasible signal structures. There are a number of possible variations for future work. First, what if the sender's preference is known so the only uncertainty is about the set of possible signal structures? We conjecture that in this case the optimal decision rule might reward the sender if the proposed signal structure is more precise. Second, what if the sender's preferences are state-dependent. One may then ask whether a suitable generalization of the quota rule exists. In the absence of uncertainty about the feasible signal structure, the sender could then commit to a decision rule that would penalize the sender if she were to choose any signal structure other than the receiver's preferred signal structure. This indeed suggests that some level of uncertainty is necessary to weaken an overly strong implication of commitment. A possible generalization of the quota rule might offer transfer rate across pairs of action state vectors that would guarantee a certain utility level for the receiver.

References

Allouah, A. and Besbes, O. (2020). Prior-independent optimal auctions. *Management Science*, 66(10):4417–4432.

Anunrojwong, J., Balseiro, S. R., and Besbes, O. (2023). Robust auction design with support information. *arXiv preprint arXiv:2305.09065*.

Babichenko, Y., Talgam-Cohen, I., Xu, H., and Zabarnyi, K. (2022). Regret-minimizing bayesian persuasion. *Games and Economic Behavior*, 136:226–248.

Bergemann, D. and Schlag, K. (2011). Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543.

Beviá, C. and Corchón, L. (2022). Contests with dominant strategies. *Economic Theory*, 74(4):1–19.

Brooks, B. and Du, S. (2021). Optimal auction design with common values: An informationally robust approach. *Econometrica*, 89(3):1313–1360.

Carroll, G. (2017). Robustness and separation in multidimensional screening. *Econometrica*, 85(2):453–488.

Carroll, G. (2019). Robust incentives for information acquisition. *Journal of Economic Theory*, 181:382–420.

Chakraborty, A. and Harbaugh, R. (2007). Comparative cheap talk. *Journal of Economic Theory*, 132(1):70–94.

Chen, B. (2012). All-or-nothing payments. *Journal of Mathematical Economics*, 48(3):133–142.

Chen, Y. and Yu, F.-Y. (2021). Optimal scoring rule design. *arXiv preprint arXiv:2107.07420*.

Clark, A. and Reggiani, G. (2021). Contracts for acquiring information. *arXiv preprint arXiv:2103.03911*.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

de Clippel, G. and Zhang, X. (2022). Non-bayesian persuasion. *Journal of Political Economy*, 130(10):2594–2642.

Deb, R., Pai, M. M., and Said, M. (2018). Evaluating strategic forecasters. *American Economic Review*, 108(10):3057–3103.

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

Ghossoub, M. and Saunders, D. (2021). On the continuity of the feasible set mapping in optimal transport. *Economic Theory Bulletin*, 9(1):113–117.

Guo, Y. and Shmaya, E. (2019). Robust monopoly regulation. *arXiv preprint arXiv:1910.04260*.

Guo, Y. and Shmaya, E. (2022). Regret-minimizing project choice. *accepted at Econometrica*.

Hartline, J., Johnsen, A., and Li, Y. (2023). Scale-robust auctions. *working paper*.

Hartline, J. D. and Lucier, B. (2015). Non-optimal mechanism design. *American Economic Review*, 105(10):3102–3124.

Holmstrom, B. (1980). On The Theory of Delegation. Discussion Papers 438, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

Jackson, M. O. and Sonnenschein, H. F. (2007). Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Laux, C. (2001). Limited-liability and incentive contracting with multiple projects. *RAND Journal of Economics*, pages 514–526.

Li, Y., Hartline, J. D., Shan, L., and Wu, Y. (2022). Optimization of scoring rules. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 988–989.

Lin, X. and Liu, C. (2022). Credible Persuasion. Papers 2205.03495, arXiv.org.

Slivkins, A. et al. (2019). Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286.

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

Villani, C. (2003). *Topics in Optimal Transportation.* Graduate studies in mathematics. American Mathematical Society.

Villani, C. (2016). *Optimal Transport: Old and New.* Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg.

Zermeno, L. (2011). A principal-expert model and the value of menus. *unpublished paper, Massachusetts Institute of Technology.*

Zermeno, L. (2012). The role of authority in a general principal-expert model. *unpublished paper, Massachusetts Institute of Technology.*

Zhao, R. R. (2008). All-or-nothing monitoring. *American Economic Review*, 98(4):1619–1628.

A Missing Proofs

Proof of Lemma 1. All proof except for the Lipschitz continuity can be found in the literature. The classic reference for the existence is Villani (2003) (pp.32) or Villani (2016) (Theorem 4.11). Villani's proof might be a bit hand-waving for some readers, so we add one explanatory note here. One thing that is not very transparent in the reference is why when a sequence of feasible solutions $F_n \to F$ in the vague topology implies F is also feasible (namely has the marginal distribution (π_1, π_2)). This is because 1) vague convergence $F_n \to F$ implies the marginal distribution also converges in the vague topology; 2) the marginal distribution of F_n is just (π_1, π_2) for any n, so the sequence of marginal distribution must converge to (π_1, π_2) in value topology; 3) if two measures over a compact space coincide on any bounded continuous function, then the two measures coincide.

The proof of $R_\gamma(q, \pi)$ being continuous in (q, π), which is essentially a careful application of Berge's Maximum Theorem, can be found in Ghossoub and Saunders (2021).

Finally, we prove that $R_\gamma(q, \pi)$ is Lipschitz-1 continuous. We first prove that

$$
\begin{aligned}
U(q, \pi) = \max_F \quad & \int_{\Delta\Theta \times A} u(\mu, a) \, \mathrm{d}F(\mu, a), \\
\text{s.t.} \quad & F(\Delta\Theta, a) = q(a), \quad \forall a \in A, \\
& F(N, A) = \pi(N), \quad \forall N \in \mathcal{B}(\Delta\Theta),
\end{aligned}
$$

is Lipschitz-1 continuous. Recall we denote the Wasserstein metric d as:

$$
\begin{aligned}
d(\pi_1, \pi_2) = \min_{G \in \Delta(\Delta\Theta \times \Delta\Theta)} & \int_{\Delta\Theta \times \Delta\Theta} |\mu - \nu|_1 \, \mathrm{d}G(\mu, \nu), \\
\text{s.t.} \quad & G(N, \Delta\Theta) = \pi_1(N), \quad \forall N \in \mathcal{B}(\Delta\Theta), \\
& G(\Delta\Theta, N) = \pi_2(N), \quad \forall N \in \mathcal{B}(\Delta\Theta),
\end{aligned}
$$

For any $\pi \neq \pi'$, denote G as the optimal solution of $d(\pi, \pi')$ (the existence of G follows from the existence result). G, as a joint distribution on $\Delta\Theta \times \Delta\Theta$, induces the conditional probability

$$
\begin{aligned}
& P^{\pi}_{\pi'}(N|\mu) : \mathcal{B}(\Delta\Theta) \times \Delta\Theta \to [0, 1], \\
\text{where} \quad & \int_{\Delta\Theta} P^{\pi}_{\pi'}(N|\mu) \, \mathrm{d}\pi' = \pi(N), \quad \forall N \in \mathcal{B}(\Delta\Theta), \\
& P^{\pi}_{\pi'}(\cdot|\mu) \text{ is a probability measure on } \Delta\Theta, \quad \forall \mu.
\end{aligned}
$$

Now denote F^*_π, and $F^*_{\pi'}$ as the optimal solution of $R_\gamma(q, \pi)$ and $R_\gamma(q, \pi')$ respectively.

Because $F^*_{\pi'}$ is a joint probability measure on $\Delta\Theta \times A$ with marginal distribution π', we can combine $P^{\pi}_{\pi'}(\cdot|\mu)$ and $F^*_{\pi'}$ to induce a probability measure on $\Delta\Theta \times \Delta\Theta \times A$, and denote its marginal distribution on $\Delta\Theta \times A$ as F_π, whose marginal distribution on $\Delta\Theta$ is π. By construction

$$F_\pi(N, a') = \int_{\Delta\Theta \times A} P^{\pi}_{\pi'}(N|\mu) 1_{a=a'} \,\mathrm{d}F^*_{\pi'}(\mu, a), \quad \forall a' \in A, N \in \mathcal{B}(\Delta\Theta).$$

Thus,

$$\begin{aligned}
&\left| \int_{\Delta\Theta\times A} u(\mu, a) \,\mathrm{d}F^*_{\pi'}(\mu, a) - \int_{\Delta\Theta\times A} u(\mu, a) \,\mathrm{d}F_\pi(\mu, a) \right| \\
&= \left| \int_{\Delta\Theta\times A} \left[u(\mu, a) - \int_{\Delta\Theta} u(\mu', a) \,\mathrm{d}P^{\pi}_{\pi'} \left(\mu'|\mu\right) \right] \mathrm{d}F^*_{\pi'}(\mu, a) \right| \\
&\leq \int_{\Delta\Theta\times A} \int_{\Delta\Theta} \left| u(\mu, a) - u(\mu', a) \right| \,\mathrm{d}P^{\pi}_{\pi'} \left(\mu'|\mu\right) \,\mathrm{d}F^*_{\pi'}(\mu, a) \\
&\leq \int_{\Delta\Theta\times A} \int_{\Delta\Theta} |\mu - \mu'|_1 \,\mathrm{d}P^{\pi}_{\pi'} \left(\mu'|\mu\right) \,\mathrm{d}F^*_{\pi'}(\mu, a), \\
&= \int_{\Delta\Theta} \int_{\Delta\Theta} |\mu - \mu'|_1 \,\mathrm{d}P^{\pi}_{\pi'} \left(\mu'|\mu\right) \,\mathrm{d}\pi'(\mu) = d(\pi, \pi').
\end{aligned}$$

The last inequality comes from the boundedness of u: $u(\theta, a) \in [0, 1]$. Because F_π has marginal distribution π on $\Delta\Theta$, it is a feasible solution for the optimization problem of $u(q, \pi)$. This means

$$U(q, \pi) \geq U(q, \pi') - d(\pi, \pi').$$

We can prove the other direction using a symmetric argument and in conclusion

$$|U(q, \pi') - U(q, \pi)| \leq d(\pi, \pi').$$

Using a similar argument one can show the first best payoff

$$u^*(\pi) = \int_{\Delta\Theta} \max_a u(\mu, a) \,\mathrm{d}\pi(\mu) = \max_F \begin{Bmatrix} \int_{\Delta\Theta\times A} u(\mu, a) \,\mathrm{d}F(\mu, a), \\ \text{s.t.} \quad F(N, A) = \pi(N), \quad \forall N \in \mathcal{B}(\Delta\Theta), \end{Bmatrix}$$

is also Lipschitz-1 continuous in π, and so is $R_\gamma(q, \pi)$. □

Proof of Lemma 2. The only if direction is trivial. We now prove the if direction.

Suppose by contradiction there exists a quota rule q' with strictly lower worst-case regret on set Σ_q. Let

$$\delta \triangleq R_\gamma(q) - \max_{\pi \in \Sigma_q} R_\gamma(q', \pi) > 0.$$

For any information structure $\pi \in \Sigma$ and any $\varepsilon > 0$, let $B_{\pi,\varepsilon}$ be an open ball around information structure π with radius ε. Because $R_\gamma(q, \pi)$ is Lipschitz-1 continuous according to Lemma 1, take $\varepsilon = \delta/2$. We know for any $\pi \in \Sigma_q$ and any $\pi' \in B_{\pi,\varepsilon}$,

$$R_\gamma(q', \pi') \leq R_\gamma(q', \pi) + \frac{\delta}{2}.$$

Let $B_\varepsilon \triangleq \cup_{\pi \in \Sigma_q} B_{\pi,\varepsilon}$ and let

$$\hat{\delta} \triangleq R_\gamma(q) - \max_{\pi \in \Sigma \setminus B_\varepsilon} R_\gamma(q, \pi).$$

Since B_ε is an open set, $\Sigma \setminus B_\varepsilon$ is closed. A closed subset of a compact set is compact so supreme is attained with a maximizer $\pi^* \in \Sigma \setminus B_\varepsilon$. Therefore,

$$\hat{\delta} = R_\gamma(q) - R_\gamma(q, \pi^*) > 0.$$

Consider the quota rule

$$q'' = \frac{\hat{\delta}}{2(1 - R_\gamma(q) + \hat{\delta})} \cdot q' + \left(1 - \frac{\hat{\delta}}{2(1 - R_\gamma(q) + \hat{\delta})}\right) \cdot q.$$

For any $\pi \in B_\varepsilon$, we have

$$\begin{aligned} R_\gamma(q'', \pi) &\leq \frac{\hat{\delta}}{2(1 - R_\gamma(q) + \hat{\delta})} \cdot \left(R_\gamma(q) - \frac{\delta}{2}\right) + \left(1 - \frac{\hat{\delta}}{2(1 - R_\gamma(q) + \hat{\delta})}\right) \cdot R_\gamma(q) \\ &= R_\gamma(q) - \frac{\hat{\delta} \cdot \delta}{4(1 - R_\gamma(q) + \hat{\delta})}. \end{aligned}$$

The first inequality holds since the principal can optimize the payoff given the quota rule constraints for q'', which is weakly better than optimizing them separately given constraints for q and q'. Similarly, for any $\pi \in \Sigma \setminus B_\varepsilon$, we have

$$\begin{aligned} R_\gamma(q'', \pi) &\leq \frac{\hat{\delta}}{2(1 - R_\gamma(q) + \hat{\delta})} \cdot 1 + \left(1 - \frac{\hat{\delta}}{2(1 - R_\gamma(q) + \hat{\delta})}\right) \cdot \left(R_\gamma(q) - \hat{\delta}\right) \\ &= R_\gamma(q) - \frac{\hat{\delta}}{2}. \end{aligned}$$

Therefore, the worst-case regret for q'' is strictly lower than q for all information structures, a contradiction. $\square$

Proof of Lemma 4. We first show that there exists a signal structure with binary signals that maximizes the principal's regret.

For any signal structure π, let S_0 be the set of posterior beliefs such that action 0 is the weakly optimal action, and S_1 be the the set of posterior beliefs such that action 1 is the uniquely optimal action. Consider another signal structure π' with binary signal $\{0, 1\}$ by pooling all signals in S_0 into 0 and all signals in S_1 into 1. Note that $u^*(\pi) = u^*(\pi')$ since pooling posteriors with the same optimal action does not change the optimal expected payoff of the receiver. Moreover, pooling signals weakly decreases the expected payoff of any quota rule. Thus, the expected regret weakly increases.

Next we show that there exists a binary partition signal structure that maximizes the principal's regret. Let $\chi(\theta, s)$ be the joint distribution over states and signals for any $\theta \in \Theta$ and $s \in \{0, 1\}$. Let $\hat{\Theta}_s = \{\Theta : \chi(\theta, s) > 0\}$ for any $s \in \{0, 1\}$. Denote the posterior mean given signal 0 and 1 by m_0 and m_1 respectively. Note that it is without loss of generality to assume that $m_0 \le m_1$. Let m^* be the posterior mean such that the receiver is indifferent between action 0 and 1. If $m_0 \ge m^*$ or $m_1 \le m^*$, by pooling two signals into one signal, i.e, by considering another signal structure $\hat{\chi}(\theta, 0) = 0$ and $\hat{\chi}(\theta, 1) = 1$ for all $\theta \in \Theta$, the optimal payoff remains unchanged while the expected payoff given the quota rule weakly decreases. Thus the regret weakly increases. Note that signal structure $\hat{\chi}$ corresponds to the degenerate case of binary partition signal structure where one of the signals occurs with probability 0.

Now we focus on the case where $m_0 < m^* < m_1$. Let $p_s = \sum_{\theta \in \Theta} \chi(\theta, s)$ be the probability signal s is sent. We first consider the case where the quota $q \ge p_1$. The case where $q \le p_1$ can be proved analogously.

Let $\underline{\theta}_s$ be the smallest state in $\hat{\Theta}_s$ and $\bar{\theta}_s$ be the largest state in $\hat{\Theta}_s$. Suppose by contradiction that $\bar{\theta}_0 > \underline{\theta}_1$. If $\underline{\theta}_1 > m_0$, we have $\underline{\theta}_0 \le m_0 < \underline{\theta}_1$ and there exists sufficiently small constant $\varepsilon > 0$ such that:

(i) $\varepsilon \le \chi(\underline{\theta}_1, 1)$;

(ii) $\varepsilon \cdot \frac{\underline{\theta}_1 - \underline{\theta}_0}{\bar{\theta}_0 - \underline{\theta}_0} \le \chi(\bar{\theta}_0, 0)$; and

(iii) $\varepsilon \cdot \frac{\bar{\theta}_0 - \underline{\theta}_1}{\bar{\theta}_0 - \underline{\theta}_0} \le \chi(\underline{\theta}_0, 0)$.

Let $\hat{f}(\bar{\theta}_0) = \varepsilon \cdot \frac{\underline{\theta}_1 - \underline{\theta}_0}{\bar{\theta}_0 - \underline{\theta}_0}$, $\hat{f}(\underline{\theta}_0) = \varepsilon \cdot \frac{\bar{\theta}_0 - \underline{\theta}_1}{\bar{\theta}_0 - \underline{\theta}_0}$, $\hat{f}(\underline{\theta}_1) = -\varepsilon$, and $\hat{f}(\theta) = 0$ for any $\theta \notin \{\underline{\theta}_0, \bar{\theta}_0, \underline{\theta}_1\}$.

Consider another signal structure $\hat{\chi}$ such that

$$\hat{\chi}(\theta, s) = \begin{cases} \chi(\theta, s) - \hat{f}(\theta) & s = 0 \\ \chi(\theta, s) + \hat{f}(\theta) & s = 1. \end{cases}$$

Given signal structure $\hat{\chi}$, the probability of each signal and their posteriors remain unchanged, and hence the regret remains the same. Moreover, in signal structure $\hat{\chi}$, the lowest state that sends signal 1 is below m_0. Therefore, we can without loss focus on the case where $\underline{\theta}_1 \leq m_0$ given signal structure χ.

Since $\bar{\theta}_1 \geq m_1 > m_0$, there exists sufficient small constant $\bar{\varepsilon} > 0$ such that

(i) $\bar{\varepsilon} \cdot \frac{m_0 - \underline{\theta}_1}{\bar{\theta}_1 - \underline{\theta}_1} \leq \chi(\bar{\theta}_1, 1)$; and

(ii) $\bar{\varepsilon} \cdot \frac{\bar{\theta}_1 - m_0}{\bar{\theta}_1 - \underline{\theta}_1} \leq \chi(\underline{\theta}_1, 1)$.

Let $\bar{f}(\bar{\theta}_1) = \bar{\varepsilon} \cdot \frac{m_0 - \underline{\theta}_1}{\bar{\theta}_1 - \underline{\theta}_1}$, and $\bar{f}(\underline{\theta}_1) = \bar{\varepsilon} \cdot \frac{\bar{\theta}_1 - m_0}{\bar{\theta}_1 - \underline{\theta}_1}$. Consider another signal structure $\bar{\chi}$ such that

$$\bar{\chi}(\theta, s) = \begin{cases} \chi(\theta, s) + \bar{f}(\theta) & s = 0 \\ \chi(\theta, s) - \bar{f}(\theta) & s = 1. \end{cases}$$

That is, signal structure $\bar{\chi}$ shifts probability mass $\bar{f}$ from signal 1 to signal 0. We have the following two observations:

1. The optimal payoff strictly increases given $\bar{\chi}$. This is because the unique optimal action for probability mass $\bar{f}$ is 0 since its conditional expectation is $m_0 < m^*$. However, the action chosen for this probability mass is 1 given signal structure χ leading to a strict payoff loss;

2. The expected payoff given the quota rule remains unchanged. This is because when $q \geq p_1$, action 1 is chosen when the signal is 1 in both χ and $\bar{\chi}$. By moving probability mass $\bar{f}$ from signal 1 to signal 0, since the posterior mean given $\bar{f}$ coincides with m_0, it is without loss to assign action 1 for probability mass $\bar{f}$ given signal structure $\bar{\chi}$ and quota q, leading to the same distribution over outcomes, and hence the same expected payoff.

Therefore, when $\gamma \in (0, 1)$, signal structure $\bar{\chi}$ leads to strictly higher regret, a contradiction. If $\gamma = 0$, then the objective coincides with max-min, and it is easy to verify that no information maximizes regret, a special case of binary partition signal structures. $\square$

Proof of Proposition 2. Let $\hat{\pi} \in \Sigma_L(q)$ be the regret maximizing left-biased signal structure for quota rule q. For any $q' < q$, we have $\Sigma_L(q) \subseteq \Sigma_L(q')$ since any signal structure that is left-biased for q is also left-biased for q'. Therefore $\hat{\pi} \in \Sigma_L(q')$.

Let $m_0 \leq m^* \leq m_1$ be the posterior means of signals 0 and 1 in signal structure $\hat{\pi}$. First note that it cannot be the case that $m_0 < m^* = m_1$. This is because by pooling two signals, the optimal payoff remains unchanged and the expected payoff from the quota rule strictly decreases, contradicting the assumption that $\hat{\pi}$ maximizes the regret.

In the case when $m_0 = m^* = m_1$, since the signal structure is left-biased, we must have $m_\rho \geq m^*$. By decreasing the quota for action 1 from q to q', the expected payoff of the receiver strictly decreases if $m_\rho > m^*$ and remains the same if $m_\rho = m^*$. Therefore,

$$R_\gamma(q, \hat{\pi}) \leq R_\gamma(q', \hat{\pi}),$$

and the inequality is strict if $m_\rho > m^*$. If $m^* < m_1$, by decreasing the quota for action 1 from q to q', the expected payoff of the receiver strictly decreases. Again we have $R_\gamma(q, \hat{\pi}) < R_\gamma(q', \hat{\pi})$. Therefore, the left-biased error is weakly decreasing in q and strictly decreasing if $m_\rho \neq m^*$. Similarly, the right-biased error is weakly increasing in q and strictly increasing if $m_\rho \neq m^*$.

To minimize the maximum regret, the quota rule needs to balance the left-biased error and right-biased error, and the optimal quota is obtained when two errors are equal. □

Proof of Corollary 1. We first show that the optimal quota is in the interior. Note that, with the help of Proposition 2, it is sufficient to show that $L_\gamma(0) > R_\gamma(0)$ and $L_\gamma(1) < R_\gamma(1)$. Indeed, when quota $q = 0$, the only feasible signal structure for right-biased error is no information, and the worst case regret in this case is

$$R_\gamma(0) = \gamma \cdot u^*(\pi^N) - (1-\gamma) \cdot u(0, \pi^N).$$

However, consider the signal structure $\hat{\pi}$ that reveals whether $\theta > 0$. This signal structure is left-biased, and hence the left-biased error is

$$L_\gamma(0) \geq \gamma \cdot u^*(\hat{\pi}) - (1-\gamma) \cdot u(0, \hat{\pi}) > \gamma \cdot u^*(\pi^N) - (1-\gamma) \cdot u(0, \pi^N) = R_\gamma(0),$$

where the last inequality holds since $u^*(\hat{\pi}) > u^*(\pi^N)$ under the assumption that there exist state θ in the support of the prior such that $\theta \in \Theta_1$, and $u(0, \hat{\pi}) = u(0, \pi^N)$ by Bayesian consistency. Similarly, we can show that $L_\gamma(1) < R_\gamma(1)$ and Corollary 1 holds.

Finally, the uniqueness comes from the fact that both the left-biased error and the right-biased error are strictly monotone in the quota when $m_\rho \neq m^*$.

$\square$

Proof of Proposition 4. The proof of Proposition 4 relies on a refined characterization of worst-case binary partition signal structure for any given quota rule.

Lemma 5. *For any $\gamma \in [0, 1)$ and any quota q, there exists a binary partition signal structure with $\min \hat{\Theta}_1 > 0$ that maximizes the right-biased regret, and there exists a binary partition signal structure with $\max \hat{\Theta}_0 < 0$ that maximizes the left-biased regret.*

Proof of Lemma 5. We prove the characterization for right-biased regret, and the other case holds analogously. To maximize the right-biased error, the lowest type in $\hat{\Theta}_1$ is positive. For any prior ρ, let $\hat{\theta}_\rho$ be the minimum state that is strictly positive in the support of ρ. For any quota rule q and any binary partition signal structure π with cutoff type $\hat{\theta} \leq 0$ (lowest type in $\hat{\Theta}_1$) such that π is right-biased for quota rule q, let π' be the binary partition signal structure with cutoff type $\hat{\theta}_\rho$, that is, the sender sends signal 1 if and only if the state is strictly positive. Note that π' is also right-biased since the probability of choosing action 1 in the optimal strategy is smaller. Moreover,

$$R_\gamma(q, \pi) = \gamma \cdot u^*(\pi) - (1 - \gamma) \cdot u(q, \pi) \leq \gamma \cdot u^*(\pi') - (1 - \gamma) \cdot u(q, \pi') = R_\gamma(q, \pi'),$$

where the inequality holds because $u^*(\pi) \leq u^*(\pi')$, since $u^*(\pi')$ achieves the first best payoff for the receiver, and $u(q, \pi) \geq u(q, \pi')$, since by increasing the cutoff the allocation is less assortative for non-positive types, leading to lower expected payoff. $\square$

Now we prove Proposition 4. We will show that the optimal quota is weakly increasing in ρ in second-order stochastic dominance in Θ_0. The other direction holds analogously.

For any type $\hat{\theta} > 0$, let π be the binary partition signal structure with cutoff type $\hat{\theta}$ given prior ρ and let π' be the binary partition signal structure with cutoff type $\hat{\theta}$ given prior ρ', where ρ second-order stochastic dominates ρ' in Θ_0. Note that signal structure π is right-biased for q if and only if π' is right-biased for q since both ρ and ρ' coincide for types in Θ_1. Therefore, the regret of the agent is

$$R_\gamma(q, \pi) = \gamma \cdot u^*(\pi) - (1 - \gamma) \cdot u(q, \pi) = \gamma \cdot u^*(\pi') - (1 - \gamma) \cdot u(q, \pi') = R_\gamma(q, \pi'),$$

where the second equality holds since having a mean preserving spreading for types in Θ_0 does not affect either the optimal payoff or the expected payoff under a fixed quota. By Lemma 5, it is sufficient to consider signal structures with strictly positive cutoffs to maximize the right-biased regret. Thus, for any quota q, the right-biased error $R_\gamma(q)$ remains unchanged in ρ in second-order stochastic dominance in Θ_0.

Finally, having a mean preserving spread in Θ_0 weakly enriches the set of possible signal structures that is left-biased for any quota q. Therefore, for any quota q, the left-biased error $L_\gamma(q)$ is weakly increasing in ρ in second-order stochastic dominance in Θ_0. By Proposition 2, the optimal quota equalizes two errors, and hence the optimal quota is weakly increasing in second-order stochastic dominance in Θ_0. □

Proof of Proposition 3. We omit the subscript of γ in notations as $\gamma = \frac{1}{2}$. In this case, the left-biased error and the right-biased error are simplified to

$$L(q) = \max_{p_0 \in [z_0, z_1],\, p_0 \le 1-q} \frac{q}{2(1-p_0)} \cdot \int_{p_0}^{1} F^{-1}(p)\,\mathrm{d}p,$$
$$R(q) = \max_{p_0 \in [z_0, z_1],\, p_0 \ge 1-q} -\frac{p_0 - 1 + q}{2p_0} \cdot \int_{0}^{p_0} F^{-1}(p)\,\mathrm{d}p.$$

Note that prior $\hat{\rho}$ first-order stochastically dominates ρ if and only if $\hat{F}^{-1}(p) \ge \hat{F}^{-1}(p)$ for any $p \in [0, 1]$. Therefore, the thresholds satisfy $\hat{z}_0 \le z_0$ and $\hat{z}_1 \le z_1$. Let q be the optimal quota rule for prior ρ and $\hat{q}$ be the optimal quota rule for prior $\hat{\rho}$. If $\hat{z}_1 < 1 - q$, in order to equalize the left-biased error and right-biased error given prior $\hat{\rho}$, we have $1-\hat{q} \le \hat{z}_1 < 1-q$, and hence $\hat{q} \ge q$.

Thus it is sufficient to focus on the case when $\hat{z}_1 \ge 1-q$. Let p_0^L be the probability that maximizes the left-biased error given prior ρ. Since $p_0^L \le 1 - q \le \hat{z}_1$, p_0^L is also a feasible choice for left-biased error given prior $\hat{\rho}$. Since $\hat{F}^{-1}(p) \ge \hat{F}^{-1}(p)$ for any $p \in [0, 1]$, the left-biased error given the choice of p_0^L is larger given prior $\hat{\rho}$ compared to given prior ρ. Therefore, the left-biased error is larger in $\hat{\rho}$.

Let $y = \sup\left\{p \,\middle|\, F^{-1}(p) \le 0\right\}$. It is easy to verify that $z_0 \le y \le z_1$. For any $p_0 \in [\hat{z}_0, \hat{z}_1]$, since $p_0^L \le 1 - q \le \hat{z}_1$, the right-biased error given the choice of p_0 is smaller when given prior $\hat{\rho}$ than when given prior ρ. Moreover, for any $p_0 \in [\hat{z}_0, y]$, the right-biased error of p_0 is smaller than the right-biased error of y given prior ρ. Therefore, the right-biased error is larger in ρ. By Proposition 2, the optimal quota rule must equalize the left-biased error and right-biased error, and hence $\hat{q} \ge q$. □

Proof of Theorem 2 The proof is a modified version of Theorem 1. We first extend the notion of generalized quota rules as in Lemma 3.

Lemma 6. *It is without loss of generality to focus on generalized quota rules (π, α) where there exists $q(\Pi, v)$ such that*

$$\alpha(\Pi, v, \cdot) = \alpha^{q(\Pi,v)}(\pi(\Pi, v), \cdot), \quad \forall(\Pi, v).$$

The proof of Lemma 6 is obvious. We next introduce the following lemma:

Lemma 7. *$U(q,\pi)$ is concave in q, so $R_\gamma(q,\pi)$ is convex in q.*

Proof of Lemma 7. Recall that

$$\begin{aligned} U(q,\pi) = \max_F \int_{\Delta\Theta\times A} & u(\mu,a)\,\mathrm{d}F(\mu,a), \\ \text{s.t. } & F(\Delta\Theta,a) = q(a), \quad \forall a\in A, \\ & F(N,A) = \pi(N), \quad \forall N\in\mathcal{B}(\Delta\Theta). \end{aligned}$$

Fixing π, denote F_q^* as the optimal solution of $U(q,\pi)$. For any q_1, q_2, $\lambda\in(0,1)$ and $q_\lambda = \lambda q_1 + (1-\lambda)q_2$, we know

$$F_\lambda = \lambda F_{q_1}^* + (1-\lambda)F_{q_2}^*$$

as a joint distribution over $\Delta\Theta\times A$ is a feasible solution of $U(q_\lambda,\pi)$. Thus we know

$$U(q_\lambda,\pi) \geq \lambda U(q_1,\pi) + (1-\lambda)U(q_2,\pi). \qquad \square$$

Proof of Theorem 2. Recall π_0 is the minimum element in Σ_0. Recall that by definition, $U(q,\pi)$ is monotonously increasing in the Blackwell order of π and

$$R_\gamma(q,\pi) = \gamma\cdot u^*(\pi) - (1-\gamma)\cdot U(q,\pi).$$

We prove the theorem by contradiction. Suppose there is a general mechanism that induces strictly less regret than the quota rule q_γ^*. By Lemma 6, it is without loss of optimality to assume this improvement is attained by a generalized quota rule $(\bar{q}(),\bar{\pi}())$. Denote

$$Q(\Pi) = \{q \mid \exists v \text{ s.t. } q = \bar{q}(\Pi,v)\}.$$

Because the general quota rule is a strict improvement, we know

$$R_\gamma(q,\pi) < R_\gamma(q_\gamma^*), \quad \forall\pi\in\Sigma, \forall q\in Q(\pi).$$

Now if $\Sigma_{q_\gamma^*}$ is a singleton, denote it as $\Sigma_{q_\gamma^*} = \{\pi^*\}$, then by Lemma 2,

$$R_\gamma(q_\gamma^*) \leq \max_{\pi\in\Sigma_{q_\gamma^*}} R_\gamma(q',\pi) = R_\gamma(q',\pi^*), \quad \forall q'.$$

This contradicts with $(\bar{q}, \bar{\pi})$ being a strict improvement.

If $\Sigma_{q^*_\gamma}$ is not a singleton, denote $q_0 \in Q(\{\pi_0\})$. According to Lemma 2, there exists $\pi' \in \Sigma_{q^*_\gamma}$ such that

$$R_\gamma(q^*_\gamma) \leq R_\gamma(q_0, \pi').$$

Since $R_\gamma(q, \pi)$ is convex in q, the lower-contour set is convex. There exists a separating hyperplane characterized by v' and c such that $q_0 \cdot v' = c$ and

$$R_\gamma(q, \pi') \geq R_\gamma(q_0, \pi'), \quad \forall q \ s.t. \ q \cdot v' \geq c. \tag{4}$$

Next, we claim that $q_0 \in \overline{\text{Conv}}(Q(\{\pi', \pi_0\}))$, i.e., the closure of the convex hull of $Q(\{\pi', \pi_0\})$. Suppose not, according to Separating Hyperplane Theorem, there exists a utility v'' such that

$$v'' \cdot q_0 > v'' \cdot q, \quad \forall q \in Q(\{\pi', \pi_0\}).$$

This violates the incentive compatibility of the general mechanism, as the sender with utility v'' will report $\hat{\Pi} = \{\pi_0\}$ (with an appropriate report on v) when $\Pi = \{\pi', \pi_0\}$.

Now that $q_0 \in \overline{\text{Conv}}(Q(\{\pi', \pi_0\}))$, we know $\bar{q}(\{\pi', \pi_0\}, v')$ must satisfies:

$$\bar{q}(\{\pi', \pi_0\}, v') \cdot v' = \max_{q \in Q(\{\pi', \pi_0\})} q \cdot v' = \max_{q \in \overline{\text{Conv}}(Q(\{\pi', \pi_0\}))} q \cdot v' \geq q_0 \cdot v' = c.$$

According to the separating hyperplane characterized in Eq. (4),

$$R_\gamma(\bar{q}(\{\pi', \pi_0\}, v'), \pi') \geq R_\gamma(q_0, \pi').$$

We already know $U(q, \pi)$ is increasing in the Blackwell order, so

$$U(\bar{q}(\{\pi', \pi_0\}, v'), \pi_0) \leq U(\bar{q}(\{\pi', \pi_0\}, v'), \pi').$$

Thus, regardless of the choice of $\bar{\pi}(\{\pi', \pi_0\})$, we must have

$$\begin{aligned}
&R_\gamma(\bar{q}(\{\pi', \pi_0\}, v'), \bar{\pi}(\{\pi', \pi_0\}, v')) \\
=&\gamma \cdot u^*(\pi') - (1-\gamma) \cdot U(\bar{q}(\{\pi', \pi_0\}, v'), \bar{\pi}(\{\pi', \pi_0\}, v')) \\
\geq&\gamma \cdot u^*(\pi') - (1-\gamma) \cdot U(\bar{q}(\{\pi', \pi_0\}, v'), \pi') \\
=&R_\gamma(\bar{q}(\{\pi', \pi_0\}, v'), \pi') \\
\geq&R_\gamma(q_0, \pi') \geq R_\gamma(q^*_\gamma).
\end{aligned}$$

This contradicts to the hypothesis that the general mechanism $(\bar{q}(), \bar{\pi}())$ is a strict improvement. $\square$