Revenue Maximization Mechanisms for an Uninformed Mediator with Communication Abilities

Zhikang Fan and **Weiran Shen**[*]

Gaoling School of Artificial Intelligence, Renmin University of China

{fanzhikang, shenweiran}@ruc.edu.cn

Abstract

Consider a market where a seller owns an item for sale and a buyer wants to purchase it. Each player has private information, known as their type. It can be costly and difficult for the players to reach an agreement through direct communication. However, with a mediator as a trusted third party, both players can communicate privately with the mediator without worrying about leaking too much or too little information. The mediator can design and commit to a multi-round communication protocol for both players, in which they update their beliefs about the other player's type. The mediator cannot force the players to trade but can influence their behaviors by sending messages to them.

We study the problem of designing revenue-maximizing mechanisms for the mediator. We show that the mediator can, without loss of generality, focus on a set of direct and incentive-compatible mechanisms. We then formulate this problem as a mathematical program and provide an optimal solution in closed form under a regularity condition. Our mechanism is simple and has a threshold structure. We also discuss some interesting properties of the optimal mechanism, such as situations where the mediator may lose money.

1 Introduction

Consider a market where a seller wishes to sell an item to a buyer. The item's quality is only known to the seller, while the buyer's valuation of the item depends on both the item's quality and the buyer's type. Both parties are interested in knowing the other player's private information but may be unwilling to disclose too much about their own. On one hand, revealing too much information gives the other player an informational advantage. If the seller knows the buyer's type, they can set a price accordingly to extract the full surplus as their revenue. Similarly, if the buyer knows the quality of the item, they can certainly utilize the information to make a better purchase decision. On the other hand, revealing no information may hinder the occurrence of a trade. For example, a buyer may deem the price unworthy if they have information about the item's quality. Therefore, this lack of information makes it difficult for both players to come to a trade agreement on their own. Even after the purchase decision is made, the players also need to deal with extensive paperwork, which imposes high costs on both of them.

This problem is ubiquitous in real-world applications and has led to the emergence of mediators between the two sides. For instance, in real estate markets, a broker or a realtor often acts as a mediator between a seller and a buyer. In fact, according to the National Association of Realtors [2022], in 2021, 87% of buyers purchased their homes, and 90% of sellers sold their homes through an agent or a broker in the US. Such a mediator typically benefits both sides and facilitates trades that may not happen on their own.

Apart from real estate markets, mediators also play a crucial role in numerous other applications: online ad exchanges match publishers with advertisers based on the features of both sides; ride-sharing platforms connect passengers with drivers based on their geographical information; lodging services match tourists with rooms based on their needs.

The aforementioned observations prompt us to investigate how a mediator can maximize their revenue by charging fees for providing such a service. We consider a model where the seller and the buyer can interact with a mediator via communication. Both players have private types that are drawn from publicly known distributions. The mediator has no access to the players' private types, but can privately communicate with them based on a specific communication protocol that may involve multiple rounds. Players update their beliefs about the other player's type after each round. In the end, the mediator decides whether to recommend a trade or not, and both players must decide whether to follow the recommendation. We assume there is no outside option for both players since there may be a significantly high cost involved (from negotiation, paperwork, etc.) if they choose to trade on their own.

However, unlike most standard mechanism design problems, the mediator cannot force the players to follow the trade recommendation. As a result, the mediator needs to carefully design the communication protocol and the pricing strategy to incentivize both players to follow their recommendations while maximizing revenue.

In this work, we take the perspective of the mediator and explore the design of a communication protocol that enables

---

[*]To whom correspondence should be addressed.

the mediator to elicit and reveal information. We also aim to determine how the mediator can price their service to maximize the expected revenue.

1.1 Our Contributions

We formally define the mechanism space for this problem. Any mechanism in our mechanism space induces an extensive-form game with incomplete information once the mediator commits to a specific protocol, where each player has a prior belief about the other player's private information. Thus, we use the Perfect Bayesian equilibrium as our solution concept and prove a variant of the celebrated revelation principle, tailored for our setting. We show that the mediator can, without loss of generality, focus on the set of direct and incentive-compatible mechanisms.

Based on the above results, we formulate this problem as a mathematical program and provide an optimal solution in closed form, subject to a regularity condition. Our optimal mechanism falls into the category of *threshold mechanisms*. Lastly, we discuss some interesting properties of the optimal mechanism. For example, a somewhat counterintuitive finding is that sellers with higher-quality items are less likely to make a sale. Additionally, in some cases, the mediator may incur a loss.

1.2 Additional Related Works

Information selling. Closest to our paper in the literature of information selling is the work by Liu *et al.* [2021]. They consider sequential interactions between an information seller and an information buyer and study the optimal pricing problem. However, our paper differs from theirs in that we consider a setting where a mediator, who has no information advantage, interacts with two privately informed players. Babaioff *et al.* [2012] consider a similar setting and provide conditions under which the information seller can achieve optimal revenue by using a one-round mechanism. Chen *et al.* [2020] consider the same setting as [Babaioff *et al.*, 2012], but with additional budget constraints.

Bilateral trade. Another relevant but significantly different direction is bilateral trade. The seminal work [Myerson and Satterthwaite, 1983] characterizes the set of mechanisms for a profit-maximizing monopolistic broker in bilateral trade. However, in their setting, both players cannot opt out of the mechanism after participation, and thus it can be modeled as a single-round mechanism. In contrast, in our setting, players need to decide whether to stay in the mechanism or follow the mediator's recommendations. Therefore, the revelation principle based on Bayesian Nash equilibrium does not apply to our setting. There is also a line of work focusing on improving efficiency through interaction [Alon *et al.*, 2015; Dobzinski *et al.*, 2014]. More recently, Mao *et al.* [2022] consider a setting where a buyer and a seller can establish a communication protocol beforehand, and both of them have commitment power. Following this protocol, the players send verifiable signals to each other in turn. However, in our setting, only the mediator has commitment power, and each player privately communicates with the mediator to update their beliefs.

Information design. Our work is also related to the literature on information design, commonly referred to as "Bayesian persuasion" [Kamenica and Gentzkow, 2011]. In a standard Bayesian persuasion model, an informed sender who has no ability to take actions aims to persuade an uninformed receiver to take actions in their favor by way of signaling. Recently, there have also been studies on its variants, e.g., an informed sender persuades an informed receiver [Kolotilin *et al.*, 2017]; an informed sender persuades a set of uninformed receivers [Castiglioni *et al.*, 2021; Xu, 2020]. In contrast, our problem lies in the area where an information sender with no informational advantage privately sends signals to two informed receivers. There is also a body of work that studies Bayesian persuasion with mediators [Shen *et al.*, 2018; Hennigs, 2019; Mahzoon, 2022; Arieli *et al.*, 2022]. However, the role of the mediators in these settings is significantly different from ours.

Communication equilibrium. Another relevant concept is communication equilibrium [Forges, 1993], where the uninformed mediator can recommend actions to the players based on the information obtained from them. We apply this notion to a real-world application. For a more detailed discussion on the connection between Bayesian persuasion and incomplete information correlated equilibrium, we refer interested readers to [Bergemann and Morris, 2016].

2 Preliminaries

Consider a market with two agents: a *seller s* and a *buyer b*. The seller has an item for sale, while the buyer seeks to purchase it. Let $t \in T$ be the private type of the buyer. Denote by $q \in Q$ the quality of the item, which can be viewed as the private type of the seller. We assume that t and q are random variables independently drawn from publicly known distributions $F(t)$ and $G(q)$, respectively. Additionally, we assume that both $F(t)$ and $G(q)$ are differentiable, with probability density functions $f(t)$ and $g(q)$ that are supported on $[t_1, t_2]$ and $[q_1, q_2]$, respectively.

Let $v(t, q)$ denote the valuation of a buyer of type t obtaining an item of quality q. We assume that the function $v(t, q)$ is monotone non-decreasing in the buyer's type t for any $q \in Q$. In particular, we assume that $v(t, q)$ is linear in t and takes the form $v(t, q) = \alpha_1(q)t + \alpha_2(q)$, with $\alpha_1(q) > 0$ for all $q \in Q$. The seller has a reserve price $r(q)$ for the item, which is assumed to be proportional to the quality of the item, i.e., $r(q) = kq$ for some constant $k \geq 0$. Alternatively, $r(q)$ can be interpreted as the seller's own valuation of the item.

Now suppose that there exists a third party (i.e., the *mediator*) who has a private communication channel through which they can communicate with either the buyer or the seller privately, collecting information from the players or revealing information to them. In the end, the mediator will either recommend "trade" with a suggested price or "not trade" at all. The mediator can also charge them for providing the service.

Note that the revenue received by the mediator comes solely from charging fees to the agents and is not related to the actual trade price. However, we can simplify our analysis by assuming that the buyer pays the trade price to the mediator, who then pays the seller. With this modification, the

utilities of all three parties remain the same, and all monetary transfers occur only between the mediator and the two agents. Throughout the paper, we make the assumption that the seller and the buyer cannot trade with each other without the mediator, as the cost of reaching a trade agreement and handling all the paperwork is formidably high. Therefore, whenever the mediator recommends "not trade", the agents cannot trade on their own, even if the trade would benefit both of them.

3 Mechanism Space

We investigate the problem of designing revenue-maximizing mechanisms for the mediator. In this section, we will describe the set of mechanisms considered in this paper. Therefore, the optimal mechanism is the one that yields the highest expected revenue among the set of mechanisms.

Definition 1 (General communication protocol). *A general communication protocol is a mechanism that induces a finite extensive-form game between the agents. The game can be described by a game tree, where each non-leaf node h is one of the following types:*

- *buyer node, where the buyer chooses an action;*

- *seller node, where the seller chooses an action;*

- *signaling-buyer node, where the mediator privately sends a message to the buyer. Any such node h has a prescription of the mediator's behavior, which is a probability distribution ψ_h over the set of its children $C(h)$;*

- *signaling-seller node, which is analogous to the signaling-buyer node but has prescribed behavior ψ_h;*

- *charging-buyer node, which has only one child node and is associated with a monetary transfer (possibly negative) t_h from the buyer to the mediator;*

- *paying-seller node, which is similar to the above node except that the transfer is from the mediator to the seller;*

- *recommendation node, where the mediator decides whether to recommend the agents to trade. We allow for randomized recommendations. Therefore, each recommendation node has two children (one for "trade" and the other for "not trade"), and is associated with the mediator's behavior, which is a probability distribution over the two children;*

- *buyer decision node, where the buyer decides whether to trade when receiving "trade" recommendation. Once the buyer refuses to trade, the game ends and goes to a leaf (terminal) node;*

- *seller decision node, which is analogous to the above node except that the decision is made by the seller.*

We require that the path from the root to any terminal node include one and only one recommendation node. Also, we require at least one buyer decision node and one seller decision node after the child node "trade" of any recommendation node to guarantee the players' right to reject. Except at the recommendation node, any interaction between the buyer and the mediator cannot be observed by the seller, and vice versa.

Our mechanism space generalizes a similar definition called the "generic interactive protocol" proposed by Babaioff *et al.* [2012]. Similar to other mechanism design problems, our mechanism proceeds as follows:

1. The mediator announces the mechanism to the agents;

2. The agents play the game described by the game tree.

In our definition, each node involving the mediator is associated with a specific behavior. Therefore, the induced game tree is only well-defined once both the seller and the buyer are aware of the mediator's intended behaviors. We make the standard assumption that the mediator has commitment power and will behave exactly as committed.

Note that the buyer node and the seller node can be regarded as nodes where the two agents send messages to the mediator since their actions rely on their private information. In fact, in a direct mechanism (which will be defined later), the action sets for the agents are precisely their type sets.

3.1 Game Re-Formulation and Solution Concept

Since we assume that the mediator has commitment power, the induced game is actually played between the seller and the buyer. Thus, given a mechanism, we can re-formulate the game and view the mediator's actions as chance moves.

Formally, given a mechanism M, a corresponding two-player extensive-form game can be defined as a tuple $\Gamma = \langle N, T, H, Z, \mathcal{H}, A, \mathcal{I}, u \rangle$, where

- $N = \{b, s\}$ is the set of players;

- T is a game tree, obtained by adding a *nature* node to the tree described by M, where nature chooses the types of players at the beginning of the game;

- H is a set of non-terminal nodes;

- Z is a set of terminal nodes;

- $\mathcal{H} = \{H_0, H_s, H_b\}$ is a partition of H, where H_0 is the set of chance move nodes, and H_s and H_b are the sets of nodes where the seller and the buyer moves, respectively;

- A is a function that maps a non-terminal node $h \in H$ to a set of available actions $A(h)$;

- $\mathcal{I} = \{\mathcal{I}_i\}_{i \in N}$ is a collection of information partitions, where each $\mathcal{I}_i$ is a partition of H_i. An element $I \in \mathcal{I}_i$ is called an *information set* of player i. Every information set I satisfies $A(h) = A(h'), \forall h, h' \in I$. Thus for any information set $I \in \mathcal{I}_i$, we denote by $A(I)$ the set $A(h)$ for all $h \in I$;

- $u = (u_i)_{i \in N}$ is a collection of payoff functions, where each $u_i : Z \mapsto \mathbb{R}$ is the payoff function of player i.

In the game tree T, any node can be uniquely determined by the path from the root to the node. This path is also called the history of the node. So with a slight abuse of notation, we use h or z to denote both a node and its history.

We assume quasi-linear utility for both the buyer and the seller. Thus, we have

$$u_b(z) = \begin{cases} v(t, q) - \tau_b(z) & \text{if } z \text{ leads to a trade} \\ -\tau_b(z) & \text{otherwise} \end{cases},$$

where $\tau_b(z)$ is the total payments made by the buyer along the path from the root to node z, and t and q can be inferred from z, as z is a history including chance moves. Note that trade occurs if and only if the mediator recommends "trade" and both agents choose to follow the recommendation. Similarly, for the seller, we have:

$$u_s(z) = \begin{cases} \tau_s(z) & \text{if } z \text{ leads to a trade} \\ \tau_s(z) + r(q) & \text{otherwise} \end{cases},$$

where $\tau_s(z)$ is the total payments from the mediator to the seller.

A *strategy* $\beta_i(I)$ of player i is a function that assigns a probability distribution over $A(I)$ to each information set $I \in \mathcal{I}_i$. Let B_i denote the set of all possible strategies of player i. Given a strategy profile $\beta = (\beta_i)_{i \in N}$, one can easily calculate the probability $\rho(z|\beta)$ of reaching a terminal node z using the structure of the game tree.

A *belief* is a function $\mu : H \mapsto [0,1]$ that assigns to each node $h \in H$ a probability such that the probabilities of the node in any information set sum up to 1, i.e., $\sum_{h \in I} \mu(h) = 1, \forall I \in \mathcal{I}_i$, for all players $i \in N$. The belief $\mu(h)$ can be interpreted as the probability that the player believes they are at node h when I is reached.

An *assessment* is a combination of a strategy profile and a belief (β, μ). Given an information set I of player i, let Z_I be the set of all terminal nodes that can be reached from some nodes in I. Given an assessment (β, μ) and information set $I \in \mathcal{I}_i$, the conditional probability of reaching a terminal node $z \in Z_I$ is denoted as $\rho(z|\beta, \mu, I)$. Then, the conditional expected payoff of player i for (β, μ) at information set I is defined as $U_{i,I}(\beta|\mu) = \sum_{z \in Z_I} \rho(z|\beta, \mu, I)u_i(z)$.

Now we are ready to define the solution concept used in this paper.

Definition 2 (Perfect Bayesian Equilibrium, PBE). *An assessment (β^*, μ^*) is a perfect Bayesian equilibrium of Γ if the following conditions hold for each player $i \in N$.*

- Sequential rationality. *For every information set $I \in \mathcal{I}_i$ and every $\beta_i \in B_i$,*

$$U_{i,I}(\beta_i^*, \beta_{-i}^*|\mu^*) \geq U_{i,I}(\beta_i, \beta_{-i}^*|\mu^*);$$

- Consistency. *For every information set $I \in \mathcal{I}_i$ and every $h \in I$,*

$$\mu^*(h) = \frac{\rho(h|\beta^*)}{\sum_{h' \in I} \rho(h'|\beta^*)},$$

where $\rho(h'|\beta^*)$ is the probability of reaching node $h' \in I$ given β^*.

Remark. *Note that the above belief update equation is only well-defined if $\sum_{h' \in I} \rho(h'|\beta^*) > 0$. Otherwise, $\mu^*(h)$ can be arbitrary, as the probability of reaching the information set is 0.*

Simply put, the players are sequentially rational if they always maximize their utilities based on their beliefs, and consistency requires that the players update their beliefs using the Bayes rule.

3.2 Direct Mechanisms and the Revelation Principle

In this section, we recast the celebrated revelation principle in our setting, showing that the mediator can narrow their focus to the set of direct and incentive-compatible mechanisms without loss of generality.

Before defining direct mechanisms, we first introduce signaling schemes, which formalize how the mediator reveals information. Given a signal set Σ, a signaling scheme $\pi : T \times Q \mapsto \Delta(\Sigma)$ is a mapping from the players' type profile to a distribution over the signal set Σ. This way of revealing information is also called "Bayesian persuasion" or an "experiment" in some research [Kamenica and Gentzkow, 2011].

Upon receiving signal σ, the buyer with type t updates their belief about the seller's type q using the standard Bayes rule:

$$\begin{aligned} g(q|\sigma, t) &= \frac{\pi(\sigma|t, q)g(q)f(t)}{\int_{q' \in Q} \pi(\sigma|t, q')g(q')f(t)\,\mathrm{d}q'} \\ &= \frac{\pi(\sigma|t, q)g(q)}{\int_{q' \in Q} \pi(\sigma|t, q')g(q')\,\mathrm{d}q'}. \end{aligned}$$

Similarly, the seller can also obtain a posterior belief:

$$f(t|\sigma, q) = \frac{\pi(\sigma|t, q)f(t)}{\int_{t' \in T} \pi(\sigma|t', q)f(t')\,\mathrm{d}t'}.$$

In this paper, we use signal schemes to describe how the mediator sends recommendations. The signal set Σ has only two signals[1], i.e., $\Sigma = \{0, 1\}$, where 0 corresponds to "not trade" while 1 corresponds to "trade". This is because we assume that the two agents cannot trade without the mediator. Thus the mediator should clearly specify whether or not they would recommend them to trade.

We define the payment functions for the two players as follows: $P_b : T \mapsto \mathbb{R}$ is the payment made by the buyer of type t to the mediator, and $P_s : Q \mapsto \mathbb{R}$ is the payment from the mediator to the seller with quality q.

Now we are ready to define direct mechanisms.

Definition 3 (Direct Mechanism). *A direct mechanism is described by a tuple (π, P_b, P_s), and proceeds as follows:*

1. *The mediator announces π, P_b and P_s;*

2. *The buyer and seller privately report (possibly untruthfully) their types t and q to the mediator;*

3. *The mediator decides whether to recommend the agents to trade according to $\pi(t, q)$;*

4. *Upon receiving the signal, the players update their belief, and decide whether to follow the recommendation;*

5. *If the trade occurs, the mediator charges the buyer $P_b(t)$ and pays the seller $P_s(q)$.*

---

[1]One might think that the reason for this is that we only need to consider the so-called "responsive experiments" as shown in [Bergemann *et al.*, 2018]. However, this is not true, since in our paper, the receivers' available actions depend on the actual signal sent by the sender. Thus their result does not apply here.

Remark. *In the above definition, although the mediator already knows both t and q when calculating the payments, the payment functions $P_b(t)$ and $P_s(q)$ still only depend on their respective types t and q. This is because the dependency on both t and q would reveal additional information.*

Definition 4 (Incentive Compatibility)**.** *A direct mechanism is incentive compatible if, for both players, reporting their true types and following the mediator's recommendation form a PBE of the game induced by the mechanism.*

One can easily verify that a direct mechanism can also be represented as a general communication protocol. We now show that the mediator can, without loss of generality, focus solely on direct and incentive compatible mechanisms.

Theorem 1. *For any general mechanism M, there exists a direct, incentive compatible mechanism that achieves the same expected revenue as in any PBE of the game induced by M.*

4 Problem Analysis

In this section, we analyze the problem and formulate it as a mathematical program. For simplicity, we use $\pi(t, q)$ to denote the probability of sending signal 1, when the reported type profile is (t, q).

The mediator's goal is to design $\pi(t, q)$, $P_b(t)$ and $P_s(q)$ to maximize their expected revenue:

$$\int_{q \in Q} \int_{t \in T} \pi(t, q)[P_b(t) - P_s(q)]f(t)g(q)\,\mathrm{d}t\mathrm{d}q.$$

We also need to pose some constraints to ensure that the agents are willing to participate and report truthfully.

Individual rationality (IR). The utility of a buyer who chooses not to participate is 0. Therefore, the expected utility of a buyer with type t who reports truthfully and follows the recommendation is:

$$U_b(t) = \operatorname*{\mathbf{E}}_{q \sim G}[\pi(t, q)(v(t, q) - P_b(t))]$$
$$= \int_{q \in Q} \pi(t, q)[v(t, q) - P_b(t)]g(q)\,\mathrm{d}q. \qquad (1)$$

Note that the seller values the item $r(q)$ even if the trade does not occur. Thus the expected utility of a seller with type q reporting truthfully and following the recommendation is:

$$U_s(q) = \operatorname*{\mathbf{E}}_{t \sim F}[\pi(t, q)P_s(q)] + \operatorname*{\mathbf{E}}_{t \sim F}[(1 - \pi(t, q))r(q)]$$
$$= \int_{t \in T} \pi(t, q)[P_s(q) - r(q)]f(t)\,\mathrm{d}t + r(q). \qquad (2)$$

Therefore, to ensure the buyer's participation, we require:

$$\int_{q \in Q} \pi(t, q)[v(t, q) - P_b(t)]g(q)\,\mathrm{d}q \geq 0. \qquad (3)$$

Similarly, for the seller, we have:

$$\int_{t \in T} \pi(t, q)[P_s(q) - r(q)]f(t)\,\mathrm{d}t + r(q) \geq r(q). \qquad (4)$$

For convenience, we define the notion of "seller surplus", which expresses the seller's additional utility gain from participating in the mechanism, as follows:

$$SU(q) = \int_{t \in T} \pi(t, q)[P_s(q) - r(q)]f(t)\,\mathrm{d}t. \qquad (5)$$

Therefore, the IR constraint for the seller is equivalent to non-negative seller surplus.

Incentive compatibility (IC). To satisfy the IC constraint, we need two steps. The first step is to ensure that following the mediator's recommendation is the best option for both agents, assuming that they truthfully reported their types in previous steps. We refer to this property as the *obedience* property. The second step is to ensure that both agents obtain a higher utility from truthfully reporting their types than misreporting any other types, even if they choose not to follow the recommendation.

The first step requires that for each agent, choosing to trade maximizes the utilities when signal 1 (the "trade" signal) is received. For the buyer, we have:

$$\int_{q \in Q} g(q|1, t)v(t, q)\,\mathrm{d}q - P_b(t) \geq 0.$$

With some simple algebraic manipulations, we obtain:

$$\int_{q \in Q} \pi(t, q)[v(t, q) - P_b(t)]g(q)\,\mathrm{d}q \geq 0. \qquad (6)$$

Similarly, for the seller, we have:

$$\int_{t \in T} \pi(t, q)[P_s(q) - r(q)]f(t)\,\mathrm{d}t \geq 0. \qquad (7)$$

Note that when the mediator sends signal 0 (the "not trade" signal), we do not need to guarantee obedience, as we assume that the agents cannot trade without the mediator.

It is interesting to observe that the above obedience constraints turn out to be exactly the same as the IR constraints. Thus we can safely ignore the obedience constraints.

For the second step, we first consider the constraints for the buyer. We still only need to consider the case when signal 1 is received, as the agents cannot trade without the mediator. Upon receiving signal 1, the expected utility of trade for a buyer with type t misreporting t' is given by:

$$U_b(t'; t) = \int_{q \in Q} \pi(t', q)[v(t, q) - P_b(t')]g(q)\,\mathrm{d}q. \qquad (8)$$

Note that for a buyer of type t, $\pi(t', q)$ may not be obedient anymore, i.e., after misreporting their type, the buyer may choose not to follow the trade recommendation if doing so leads to a negative utility. So the expected utility for the buyer is $\max\{U_b(t'; t), 0\}$. Thus the incentive compatibility constraints for the buyer become the following:

$$U_b(t) \geq \max\{U_b(t'; t), 0\}. \qquad (9)$$

Interestingly, $U_b(t) \geq 0$ is already implied by the IR constraint (Constraint (3)), so the only useful constraint is

$$U_b(t) \geq U_b(t'; t). \qquad (10)$$

For the seller, upon receiving signal 1, the expected utility of trade for a seller with type q misreporting q' is

$$U_s(q';q) = \int_{t \in T} \pi(t,q')[P_s(q') - r(q)]f(t)\,\mathrm{d}t + r(q), \tag{11}$$

where $\pi(t,q')$ may also not be obedient for the seller. With arguments similar to that for the buyer, we obtain that the IC constraint for the seller is the following:

$$U_s(q) \geq U_s(q';q). \tag{12}$$

The seller's IC constraint can also be equivalent to misreporting does not result in a higher seller surplus, that is:

$$SU(q) \geq SU(q';q), \tag{13}$$

where $SU(q';q) = \int_{t \in T} \pi(t,q')[P_s(q') - r(q)]f(t)\,\mathrm{d}t$.

Combining all the constraints together, we can formulate the mechanism design problem as the following mathematical program:

$$\begin{aligned}
\text{maximize} \quad & \int_{q \in Q} \int_{t \in T} \pi(t,q)[P_b(t) - P_s(q)]f(t)g(q)\,\mathrm{d}t\mathrm{d}q \\
\text{subject to} \quad & U_b(t) \geq 0, && \forall t \in T \\
& SU(q) \geq 0, && \forall q \in Q \\
& U_b(t) \geq U_b(t';t), && \forall t, t' \in T \\
& SU(q) \geq SU(q';q), && \forall q, q' \in Q
\end{aligned} \tag{14}$$

5 The Optimal Mechanism

In this section, we apply a variant of the Myersonian approach and give an optimal solution in closed form to the optimization problem (14).

Before we present our main result, we need the following definition.

Definition 5 (Virtual Value Function and Virtual Cost Function). *For any random variable w with PDF $x(w)$ and CDF $X(w)$, the virtual value function $\phi^-(w)$ and the virtual cost function $\phi^+(w)$ are defined as:*

$$\phi^-(w) = w - \frac{1 - X(w)}{x(w)}, \quad \phi^+(w) = w + \frac{X(w)}{x(w)}.$$

The definition of the virtual value function is the same as in [Myerson, 1981], and the virtual cost function is also standard in the literature (see, e.g., [Manelli and Vincent, 1995; Myerson and Satterthwaite, 1983]).

Let $\phi_b^-(t)$ be the buyer's virtual value function and $\phi_s^+(q)$ the seller's virtual cost function. Our main result is built upon the following regularity condition:

Definition 6 (Regularity). *A problem instance is regular if both functions $\phi_b^-(t)$ and $\frac{k\phi_s^+(q) - \alpha_2(q)}{\alpha_1(q)}$ are monotone non-decreasing.*

This regularity condition is also standard in the literature [Myerson, 1981; Cai and Daskalakis, 2011; Krishna and Maenner, 2001] and is satisfied for a wide range of distributions. Our solution belongs to the following category of *threshold mechanisms* when the problem instance is *regular*.

Definition 7. *A mechanism (π, P_b, P_s) is called a threshold mechanism if there exist monotone functions $\lambda(t)$ and $\eta(q)$, such that the mediator recommends "trade" as long as $\lambda(t) \geq \eta(q)$, i.e.,*

$$\pi(t,q) = \begin{cases} 1 & \text{if } \lambda(t) \geq \eta(q) \\ 0 & \text{otherwise} \end{cases}.$$

In this case, $\pi(t,q)$ is fully described by function $\lambda(t)$ and $\eta(q)$.

Now we are ready to present our optimal mechanism.

Theorem 2. *Suppose that a problem instance satisfies the regularity condition. Then the threshold mechanism with the following threshold functions and payment functions is an optimal mechanism:*

$$\lambda(t) = \phi_b^-(t) \quad and \quad \eta(q) = \frac{k\phi_s^+(q) - \alpha_2(q)}{\alpha_1(q)},$$

$$\begin{aligned}
P_b^*(t) = &\frac{1}{\int_{q \in Q} \pi^*(t,q)g(q)\,\mathrm{d}q} \left[\int_{q \in Q} \pi^*(t,q)v(t,q)g(q)\,\mathrm{d}q \right. \\
& \left. - \int_{t_1}^{t} \int_{q \in Q} \alpha_1(q)\pi^*(x,q)g(q)\,\mathrm{d}q\mathrm{d}x \right],
\end{aligned} \tag{15}$$

$$\begin{aligned}
P_s^*(q) = &\frac{1}{\int_{t \in T} \pi^*(t,q)f(t)\,\mathrm{d}t} \left[\int_{t \in T} \pi^*(t,q)r(q)f(t)\,\mathrm{d}t \right. \\
& \left. + k \int_{q}^{q_2} \int_{t \in T} \pi^*(t,x)f(t)\,\mathrm{d}t\mathrm{d}x \right].
\end{aligned} \tag{16}$$

Example 1. *Consider an example where both the buyer's type and the seller's quality are uniformly distributed in the interval $[1,2]$, i.e., $T = Q = [1,2]$, and $f(t) = g(q) = 1, \forall t, q$. We set $\alpha_1(q) = q$ and $\alpha_2(q) = 0$. Let $r(q) = 1.5q$, i.e., $k = 1.5$.*

In this example, we have $\phi_b^-(t) = t - \frac{1-F(t)}{f(t)} = 2t - 2, \phi_s^+(q) = q + \frac{G(q)}{g(q)} = 2q - 1$. Therefore $\lambda(t) = \phi_b^-(t) = 2t - 2$ is non-decreasing in t and $\eta(q) = k\frac{\phi_s^+(q)}{q} = 3 - \frac{1.5}{q}$ is also non-decreasing in q, satisfying the regularity condition. In the optimal mechanism, the mediator will recommend them to trade if $\lambda(t) \geq \eta(q)$, or equivalently, when $t \geq 2.5 - \frac{1.5}{2q}$. For the buyer, there are two cases.

- *When $t < 1.75$, we have $t < 2.5 - \frac{1.5}{2q}, \forall q$. This means that the mechanism will never recommend trade, or equivalently, $\pi^*(t,q) = 0, \forall q$. The payment is $P_b^*(t) = 0$. For these buyers, there is no trade the mediator charges nothing.*

- *When $t \geq 1.75$, the mechanism will recommend trade when $q \leq \frac{1.5}{5-2t}$. In this case, the buyer's payment function becomes:*

$$P_b^*(t) = 0.5 + \frac{0.9375}{2.5 - t}.$$

For the seller, there are also two cases.

- *When $q > 1.5$, then $2.5 - \frac{1.5}{2q} > 2$. This means the mediator will never recommend trade since t is at most 2. So we have $\pi^*(t, q) = 0$ and $P_s^*(q) = 0$ for all t in this case.*

- *When $q \leq 1.5$, we have that $\pi^*(t, q) = 1$ when $t \geq 2.5 - \frac{1.5}{2q}$. If $q < 1.5$, the payment function for the seller $P_s^*(q)$ becomes:*

$$P_s^*(q) = \frac{9q}{6 - 4q} \ln\left(\frac{1.5}{q}\right).$$

And if $q = 1.5$, we can take the limit and get $P_s^*(q) = 2.25$.

The rest of the section is devoted to the proof of Theorem 2. We first introduce the following quantities:

$$R_b^\pi(t) = \int_{q \in Q} \alpha_1(q)\pi(t, q)g(q) \, \mathrm{d}q, \tag{17}$$

$$R_s^\pi(q) = \int_{t \in T} \pi(t, q)f(t) \, \mathrm{d}t. \tag{18}$$

Note that $R_b^\pi(t)$ and $R_s^\pi(q)$ can be viewed as the expected probability of being recommended to trade, except that $R_b^\pi(t)$ is weighted by $\alpha_1(q)$.

We call a mechanism (π, P_b, P_s) *feasible* if it satisfies the constraints in Program (14). The following lemma gives a characterization of feasible mechanisms.

Lemma 1. *A mechanism (π, P_b, P_s) is feasible if and only if it satisfies the following constraints:*

$$R_b^\pi(t) \text{ is monotone non-decreasing in } t. \tag{19}$$
$$R_s^\pi(q) \text{ is monotone non-increasing in } q. \tag{20}$$

$$U_b(t) = U_b(t_1) + \int_{t_1}^{t} R_b^\pi(x) \, \mathrm{d}x \tag{21}$$

$$SU(q) = SU(q_1) - k \int_{q_1}^{q} R_s^\pi(x) \, \mathrm{d}x \tag{22}$$

$$U_b(t_1) \geq 0 \tag{23}$$
$$SU(q_2) \geq 0 \tag{24}$$

Before we derive the optimal mechanism, we first need to rewrite the revenue of the mediator.

Lemma 2. *The mediator's expected revenue of any feasible mechanism (π, P_b, P_s) can be written as:*

$$Rev(\pi, P_b, P_s)$$
$$= \int_{q_1}^{q_2} \int_{t_1}^{t_2} \pi(t, q)[\alpha_1(q)\phi_b^-(t) + \alpha_2(q) - k\phi_s^+(q)]$$
$$\cdot f(t)g(q) \, \mathrm{d}t\mathrm{d}q - U_b(t_1) - SU(q_2). \tag{25}$$

The following result shows that the mechanism described in Theorem 2 is feasible.

Lemma 3. *For a regular problem instance, the mechanism (π^*, P_b^*, P_s^*) defined in Theorem 2 is feasible.*

We only provide a proof sketch of Theorem 2 due to space limit.

Proof sketch of Theorem 2. According to Lemma 3, our mechanism is feasible. The definition of π^* implies that our mechanism maximizes the first term in Equation (25). And using the definition of $P_b^*(t)$ and $P_s^*(q)$, we can show that $U_b(t_1) = 0$ and $SU(q_2) = 0$. This means that our mechanism optimizes the 3 terms in Equation (25) simultaneously, and hence optimal. $\square$

Our main result reveals some interesting facts about the optimal mechanism.

Fact 1. *For a regular problem instance, the higher the seller's type is, the less likely their item will be sold.*

For a regular instance, the probability of $\lambda(t) \geq \eta(q)$ becomes smaller as $\eta(q)$ is monotone non-decreasing in q. Sellers with high-quality items also have higher reserve prices, which means that fewer buyers can afford to buy the item.

Fact 2. *In the optimal mechanism, for each buyer of type t, there is a certain threshold q', such that the buyer can trade with the seller only if the seller's type q satisfies $q \leq q'$.*

Fact 2 shows that even a buyer with a high type may still buy a low-quality item, but it is impossible for a low-typed buyer to buy a high-quality item.

Fact 3. *In the optimal mechanism, the mediator may lose money for some type profiles (t, q).*

This happens whenever the mediator's payment to the seller is higher than the payment collected from the buyer, i.e., $P_s^*(q) > P_b^*(t)$. Let us re-consider Example 1. As shown in Figure 1, the mediator will recommend trade in the area above the blue curve. And the equation $P_s^*(q) = P_b^*(t)$ is the orange curve. Therefore, in the area between the blue curve and the orange curve, the mediator recommends trade but loses money.
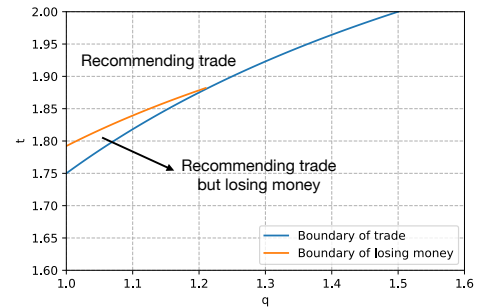

Figure 1: The threshold mechanism in Example 1.

6 Conclusion

We studied the problem of designing revenue-maximizing mechanisms for the mediator. We formally defined the set of mechanisms that can be used by the mediator, and proved that the mediator can, without loss of generality, consider the set of direct and incentive compatible mechanisms. We showed that under the regularity condition, the optimal mechanism is a threshold mechanism. Our closed-form solution also reveals some interesting properties of the optimal mechanism.

Acknowledgements

References

[Alon *et al.*, 2015] Noga Alon, Noam Nisan, Ran Raz, and Omri Weinstein. Welfare maximization with limited interaction. In *Proceedings of the 2015 IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1499–1512, 2015.

[Arieli *et al.*, 2022] Itai Arieli, Yakov Babichenko, and Fedor Sandomirskiy. Bayesian persuasion with mediators. *ArXiv*, abs/2203.04285, 2022.

[Babaioff *et al.*, 2012] Moshe Babaioff, Robert Kleinberg, and Renato Paes Leme. Optimal mechanisms for selling information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 92–109, 2012.

[Bergemann and Morris, 2016] Dirk Bergemann and Stephen Morris. Information design, bayesian persuasion and bayes correlated equilibrium. *Econometrics: Econometric & Statistical Methods - General eJournal*, 2016.

[Bergemann *et al.*, 2018] Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American economic review*, 108(1):1–48, 2018.

[Cai and Daskalakis, 2011] Yang Cai and Constantinos Daskalakis. Extreme-value theorems for optimal multidimensional pricing. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 522–531. IEEE, 2011.

[Castiglioni *et al.*, 2021] Matteo Castiglioni, A. Celli, Alberto Marchesi, and Nicola Gatti. Signaling in bayesian network congestion games: the subtle power of symmetry. In *AAAI*, 2021.

[Chen *et al.*, 2020] Yiling Chen, Haifeng Xu, and Shuran Zheng. Selling information through consulting. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2412–2431. SIAM, 2020.

[Dobzinski *et al.*, 2014] Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 233–242, 2014.

[Forges, 1993] Françoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and decision*, 35(3):277–310, 1993.

[Hennigs, 2019] Raphaela Hennigs. Conflict prevention by bayesian persuasion. *Conflict Studies: Prevention*, 2019.

[Kamenica and Gentzkow, 2011] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[Kolotilin *et al.*, 2017] Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.

[Krishna and Maenner, 2001] Vijay Krishna and Eliot Maenner. Convex potentials with an application to mechanism design. *Econometrica*, 69(4):1113–1119, 2001.

[Liu *et al.*, 2021] Shuze Liu, Weiran Shen, and Haifeng Xu. Optimal pricing of information. *Proceedings of the 22nd ACM Conference on Economics and Computation*, 2021.

[Mahzoon, 2022] Majid Mahzoon. Hierarchical bayesian persuasion: Importance of vice presidents. *ArXiv*, abs/2204.05304, 2022.

[Manelli and Vincent, 1995] Alejandro M Manelli and Daniel R Vincent. Optimal procurement mechanisms. *Econometrica: Journal of the Econometric Society*, pages 591–620, 1995.

[Mao *et al.*, 2022] Jieming Mao, Renato Paes Leme, and Kangning Wang. Interactive communication in bilateral trade. In *ITCS*, 2022.

[Myerson and Satterthwaite, 1983] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.

[Myerson, 1981] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

[of Realtors, 2022] National Association of Realtors. National association of realtors statistics. https://www.nar.realtor, 2022. Accessed: 2022-8-5.

[Shen *et al.*, 2018] Weiran Shen, Pingzhong Tang, and Yulong Zeng. A closed-form characterization of buyer signaling schemes in monopoly pricing. In *AAAS*, 2018.

[Xu, 2020] Haifeng Xu. On the tractability of public persuasion with no externalities. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2708–2727, 2020.