Public Signaling in Markets with Information Asymmetry Using a Limited Number of Signals

Xu Zhao[1] , **Ren Liu**[1] , **Weiran Shen**[1,2,3]*
[1]Gaoling School of Artificial Intelligence, Renmin University of China
[2]Beijing Key Laboratory of Research on Large Models and Intelligent Governance
[3]Engineering Research Center of Next-Generation Intelligent Search and Recommendation, MOE
{zhao_xu, liuren201438, shenweiran}@ruc.edu.cn

Abstract

Consider a market with a seller and many buyers. The seller has a kind of item for sale to the buyers. The items have a quality and each buyer has a private type. The quality is only known to the seller, and the buyers only have a prior belief of the quality. A third party (e.g., intermediaries or product reviewers) is able to reveal information about the actual quality by using a so-called signaling scheme. After receiving the information, buyers can update their beliefs accordingly and decide whether to buy the items. We consider the third party's problem of maximizing the purchasing probability by sending signals. However, the optimal signaling scheme has implementation issues, as the number of signals in the optimal scheme is the same as the number of buyer types, which can be exceedingly large or even infinite. We therefore investigate whether a finite and limited set of signals could still approximate the performance of the optimal signaling scheme. Unfortunately, our results show that with a finite number of signals, no signaling scheme can achieve a certain fraction of the performance of the optimal signaling scheme. This limitation persists even with the regularity or the monotone hazard rate assumption. Nevertheless, we identify a mild technical condition under which the third party can approximate the optimal performance within a constant factor by employing only two signals. We also conduct extensive experiments to substantiate our theoretic results. These experiments compare the performance of using a small signal set across different value distributions. Despite the negative results, our experiment results show that using only a small number of signals is able to achieve a fairly reasonable performance in average cases.

1 Introduction

Consider a market where a seller has a kind of items for sale to buyers and sets a fixed price. The items have a quality which is only known to the seller, while the valuation of the item to each buyer depends on the buyer's type which is private information of the buyer. Both the seller and the buyers only hold prior beliefs about the private information of the other party.

In such markets, buyers can only make purchase decisions based on their prior beliefs and the price of the items. The information asymmetry in these markets may lead to adverse selection issues (e.g., the markets for lemons [Akerlof, 1978]), and such inefficiency gives rise to third parties that can bridge the information gap between the two sides. These third parties can be intermediaries or online product reviewers that reveal information about the product quality to the buyers to help the seller better target potential customers. After receiving the information, buyers update their beliefs and decide whether to buy an item accordingly.

This phenomenon is ubiquitous in real-world applications. For instance, in used car markets, a car dealership may act as a third party and provide inspection results (e.g., showing whether each component works) of the cars. The buyers may re-evaluate the cars and make purchase decisions based on the results. Similar roles can also be found in real estate markets, where a real estate agent may provide home inspection reports describing the status of a house to help buyers make better estimation about the house. Websites like Zillow or Realtor even let home buyers search for properties, view property details, and schedule consultations with sellers or their representatives. In social media platforms, people or organizations with expert knowledge in certain fields may post content about certain products or reviews of them to affect the buyer actions of their audience. Such phenomena are sometimes referred to as influencer marketing in the literature.

The various applications mentioned above motivate us to explore how a third party can maximize market trading volume (e.g., the probability of purchases) by revealing information. We employ the Bayesian persuasion model [Kamenica and Gentzkow, 2011] to describe the process of revealing information. In our model, both the quality of the items and the buyer types (their valuations for item) are drawn from publicly known distributions. The third party first commits to a signaling scheme that describes how signals are correlated with the actual item quality. Then the third party assesses the item quality and reveals partial information to the buyers by sending a certain signal. Upon receiving the signal, buyers update their beliefs about the quality of the item and make

*To whom correspondence should be addressed.

purchasing decisions accordingly.

It is worth mentioning that although there are multiple buyers in our setting, which is different from the standard Bayesian persuasion model, we consider the case where the third party uses the same scheme for the buyers and sends the same signal to them. This can be justified by the observation that both intermediaries and product reviewers disclose information publicly, and therefore all buyers receive the same signal. Such signaling schemes are sometimes referred to as public signaling schemes in the literature [Alonso and Câmara, 2016; Dughmi, 2019; Xu, 2020; Zheng and Chen, 2021; Castiglioni *et al.*, 2023].

It is known that in the standard Bayesian persuasion model, an optimal signaling scheme uses the same number of signals as receiver actions [Kamenica and Gentzkow, 2011]. In this paper, we regard all buyers as a meta-buyer since they always receive the same signal. It turns out that the number of signals used by an optimal scheme equals the number of different buyer types, which can be exceedingly large in a big market and can cause implementation issues. Therefore, in this paper, we focus our attention on whether it is possible to achieve a performance comparable to the optimal one by using only a limited number of signals. Our contributions can be summarized as follows:

- We show that in the general distribution case, no signaling scheme that uses finitely many signals can achieve a constant approximation of optimal performance. The result holds even when the value distribution satisfies the regularity condition or the monotone hazard rate condition, which are standard and widely used assumptions in the literature.
- We identify a technical condition under which employing only two signals can already approximate the optimal performance well. This result implies that using a small number of signals can achieve a certain fraction of the optimal performance.
- We also conduct experiments to substantiate our theoretical findings. In the experiments, we compare the performance of using a limited number of signals with the optimal performance. Despite the negative theoretical results, our experiments show that a small number of signals can achieve good performance on average.

1.1 Related Work

Many papers have explored the impact of third party product reviews and certifications on firms' sales [Eliashberg and Shugan, 1997; Chen and Xie, 2005; Chen *et al.*, 2012; Pei and Mayzlin, 2021; Tian, 2022]. Pei and Mayzlin [2021] examine the optimal level of affiliation from the firm's perspective, since affiliation brings a positive bias to third party's reviews but also decreases the credibility of the reviews. Tian [2022] study how a firm could benefit from collaborating with an honest third party. However, unlike their studies, our paper does not take into account the affiliation between firms and third parties. Since our purpose is to increase the trading volume in a market, we only consider what signaling scheme the third party should adopt to maximize the trading volume.

Our work is also related to the literature on bilateral trade [Blumrosen and Dobzinski, 2014; Colini-Baldeschi *et al.*, 2016; Kang and Vondrák, 2019], which was initially pioneered by Myerson and Satterthwaite [1983]. In contrast, in our paper, the buyers and the seller cannot communicate directly. The buyers can obtain information about the seller's items through signals from the third party, but the seller cannot learn any additional private information about the buyers.

Another relevant topic is the so-called Bayesian persuasion model [Kamenica and Gentzkow, 2011]. The most relevant variant is the public Bayesian persuasion problem [Alonso and Câmara, 2016; Dughmi, 2019; Xu, 2020; Zheng and Chen, 2021; Castiglioni *et al.*, 2023], where a sender sends a public signal to persuade the signal receivers. In the standard Bayesian persuasion model, the sender has the ability to design and commit to any signaling scheme. However, in our work, we consider a limited number of signals and explore what can be achieved by such signaling schemes. Additionally, another line of research studies persuasion with limited communication [Dughmi *et al.*, 2016; Aybas and Turkel, 2019; Gradwohl *et al.*, 2022], where the communication capabilities (e.g., bandwidth) between the sender and the receiver are bounded.

There is also a series of research studying the trade-off between the optimality and complexity of auction mechanisms [Hartline and Roughgarden, 2009; Alaei *et al.*, 2013; Alaei *et al.*, 2019; Cai *et al.*, 2016; Chawla *et al.*, 2014; Huang *et al.*, 2015; Yao, 2014; Shen and Tang, 2017]. Our paper is relevant in the sense that we study what can be achieved by simple signaling schemes, aiming to strike a balance between complexity and performance. For practical purposes, many researchers focus on designing mechanisms that are simple and can guarantee robust performance compared to the optimal one even in the worst case scenarios.

2 Preliminaries

Consider the setting where a seller sells a kind of items to a population of buyers and sets the price to p. We assume that the seller has an unlimited supply, meaning that any buyer who is willing to pay the price p can buy an item. Each item has a quality $q \in Q = [\underline{q}, \overline{q}]$, which is drawn from a distribution $G(q)$. Assume $G(q)$ is differentiable with density function $g(q)$ that has full support on Q. Suppose that the buyers only know the distribution $G(q)$ but have no access to the actual quality q.

Suppose that each buyer has a type $v \in V = [\underline{v}, \overline{v}]$ that is drawn from a publicly known distribution $F(v)$. We also assume $F(v)$ is differentiable with density function $f(v)$ and has full support on V. Following the literature convention [Dughmi *et al.*, 2016], we consider quasi-linear utilities for the buyers defined as:

$$u(v, q) = vq - p. \tag{1}$$

Since the buyers only know the quality distribution $G(q)$, a buyer with type v will buy an item only if the following condition is satisfied:

$$u(v) = \mathbb{E}_{q\sim G}[u(v, q)] = v\,\mathbb{E}_{q\sim G}[q] - p \geq 0. \tag{2}$$

Now, suppose that there is a third party that has access to the quality of the item. After getting the actual quality q, the third party can send messages to the buyers to influence their beliefs about the quality, thereby affecting their purchasing decisions. This third party uses a so-called signaling scheme to inform the buyers. A signaling scheme consists of both a message set and a function that determines how messages are correlated with the quality. In this paper, we use $\sigma = (S, \pi)$ to denote a signaling scheme, where S is the set of possible signals (messages) and $\pi : Q \times S \mapsto \mathbb{R}$ is a signaling function that maps a quality to a distribution of signals. In this paper, we consider both continuous and discrete signal spaces. If the signal space S is discrete, we use $\pi(s|q)$ to denote the probability of sending signal s when the actual quality is q. And if S is continuous, $\pi(s|q)$ becomes a density function of s conditioned on q. For simplicity, we abuse notation and employ $\pi(s, q)$ to denote the joint probability distribution of signal s and quality q. A signaling scheme should naturally satisfy the following feasibility constraint:

$$\int_S \pi(s|q)\,\mathrm{d}s = 1 \quad \text{or} \quad \sum_{s\in S} \pi(s|q) = 1$$

depending on whether the signal space S is discrete or continuous.

We assume that the third party has commitment power, i.e., the third party samples a signal from S exactly based the distribution $\pi(s|q)$ upon accessing q. The third party first announces their signaling scheme $\pi(s|q)$ to the buyers, and then observes q and samples a signal s according to $\pi(s|q)$ and sends s to all the buyers. Since the third party has commitment power, after receiving the signal from the third party, the buyers update their belief and get a posterior distribution about the quality of the item by applying the Bayes rule:

$$g(q|s) = \frac{\pi(s|q)g(q)}{\int_{q\in Q} \pi(s|q)g(q)\,\mathrm{d}q},$$

where $g(q|s)$ is the posterior belief of the buyers when receiving signal s. Then a buyer with type v makes a purchase decision based on whether the following term is non-negative:

$$\mathbb{E}[u(v, q)|s] = v\,\mathbb{E}[q|s] - p = v \cdot \frac{\int_{q\in Q} q\pi(s|q)g(q)\,\mathrm{d}q}{\int_{q\in Q} \pi(s|q)g(q)\,\mathrm{d}q} - p.$$

In this paper, we consider the case where the third party can only send public signals, i.e., the third party announces the signal publicly and all the buyers receive the same signal. Such a third party can be a product reviewer that reveals additional information to potential customers, or a quality certification organization that issues certificates to products with certain quality levels. Such information is usually publicly available and thus can be seen by all buyers.

Note that sending public signals is different from using the same signaling scheme for all the buyers, where the third party can still send different signals to different buyers, as long as these signals are sampled according the same signal scheme. In our setting, the third party must send exactly the same signal to all buyers. Such a signal can be a random variable. But once realized, all buyers should receive the same signal.

Suppose that the third party extracts a certain fraction of the payment as its revenue. Therefore, the goal of the third party is to attract as many buyers as possible to buy the item as possible. Or equivalently, for a random buyer, the goal of the third party is to maximize the buying probability of the buyer.

Let $\Pr(buy|s)$ denote the probability of a random buyer purchasing an item after receiving signal s. Since only buyers with value $v\,\mathbb{E}[q|s] \geq p$ will purchase an item, we have:

$$\Pr(buy|s) = 1 - F\left(\frac{p}{\mathbb{E}[q|s]}\right).$$

The utility of the third party can be written as:

$$\int_{s\in S} \pi(s) \Pr(buy|s)\,\mathrm{d}s = \int_{s\in S} \pi(s)\left[1 - F\left(\frac{p}{\mathbb{E}[q|s]}\right)\right]\mathrm{d}s,$$

or

$$\sum_{s\in S} \pi(s) \Pr(buy|s) = \sum_{s\in S} \pi(s)\left[1 - F\left(\frac{p}{\mathbb{E}[q|s]}\right)\right],$$

where $\pi(s)$ is the marginal distribution of sending signal s:

$$\pi(s) = \int_{q\in Q} \pi(s|q)g(q)\,\mathrm{d}q. \tag{3}$$

In a standard Bayesian persuasion setting, there is a signal sender and a signal receiver. And there exists an optimal signaling scheme where the number of signals is at most the number of possible actions of the receiver. Each signal leads to the receiver playing a certain action, and thus the signals can be viewed as "action recommendations". However, in our setting, there is a population of receivers and different receivers may take different actions. Nonetheless, we can view the population as a single "meta receiver", since all buyers receive the same signal. The meta receiver's action can be indexed by a buyer type v, which corresponds to a meta action: all buyers with a type at least v decide to buy the items. In this case, a signal can be viewed as an "outcome indicator". Therefore, an optimal signaling scheme may feature a continuous signal space, as the buyer type v is continuous.

Since each signal s leads to an expected quality of the item, which completely determines a buyer's purchase decision. We can then use such an expectation to index the corresponding signal. If two signals lead to the same expected quality, we can simply merge these two signals and obtain exactly the same expected quality. In this case, all three parties' utilities stay the same. Therefore, we require the signaling scheme to satisfy:

$$\mathbb{E}[q|s] = \frac{\int_{q\in Q} q\pi(s|q)g(q)\,\mathrm{d}q}{\int_{q\in Q} \pi(s|q)g(q)\,\mathrm{d}q} = s,$$

or equivalently,

$$\int_{q\in Q} q\pi(s|q)g(q)\,\mathrm{d}q = s \int_{q\in Q} \pi(s|q)g(q)\,\mathrm{d}q.$$

We can then formulate the problem as the following mathematical program:

$$\begin{aligned}
\max_{\pi} \quad & \int_{s\in S} \pi(s)\left[1 - F\left(\frac{p}{\mathbb{E}[q|s]}\right)\right] \mathrm{d}s \\
\text{s.t.} \quad & \int_S \pi(s|q)\,\mathrm{d}s = 1, && \forall q, \\
& \pi(s|q) \geq 0, && \forall s, \forall q, \\
& \int_{q\in Q} q\pi(s|q)g(q)\,\mathrm{d}q = s\int_{q\in Q} \pi(s|q)g(q)\,\mathrm{d}q, && \forall s
\end{aligned} \tag{4}$$

Such a continuous signaling scheme is clearly impossible to implement in reality. It is infeasible to even represent such a signaling scheme, let alone compute the optimal one. Therefore, in this paper, we aim to study what performance can be achieved by only using a limited number of signals. A signaling scheme $\sigma = (S, \pi)$ is called an n-signal scheme, if the signal set S contains only n signals.

Besides, we assume that $\overline{v} \leq \frac{p}{\underline{q}}$ and $\underline{v} \geq \frac{p}{\overline{q}}$, as buyers with a valuation $v > \frac{p}{\underline{q}}$ (or $v < \frac{p}{\overline{q}}$) will always choose to buy (or not buy) an item, making it unnecessary to send them a signal. Additionally, we assume $\underline{q} \geq \frac{p}{\overline{v}}$, since, a priori, an item with an expected quality $q < \frac{p}{\overline{v}}$ would not be purchased by any buyer and would thus be excluded from the market.

3 Theoretic Analysis

In this section, we analyze the problem theoretically. We first show that, without loss of generality, we can consider only a special type of schemes called n-segmentations (Definition 2). Then, we compare the performance (the probability of purchases) of using finitely many signals to that of the optimal scheme under different conditions.

We use T^* to denote the overall buying probability of the optimal signaling scheme when using a continuous signaling space, i.e., T^* is the optimal objective value of Program (4). Let T_n^* be the overall buying probability of the optimal signaling scheme when using only n signals (i.e., $|S| = n$). Clearly, we have $T_n^* \leq T^*$.

A useful signaling scheme that will be considered in our paper is the *full information revelation scheme*, which always directly reveals the actual quality q to the buyer population.

Definition 1 (Full Information Revelation)**.** *A signaling scheme is called full information revelation, if the third party always directly discloses the quality q of the item, i.e., a full information revelation scheme is $\sigma_{full} = (S_{full}, \pi_{full})$ with*

$$S_{full} = Q,$$
$$\pi_{full}(q'|q) = \delta(q' - q),$$

where $\delta(\cdot)$ is the Dirac delta function.

When the third party uses the full information revelation scheme, the buyers will know the exact quality q after receiving the signal. There is no need to calculate the posterior belief, or equivalently, the posterior belief is a single point mass at q. Similarly, we use T_{full} to denote the overall buying probability of the population under the full information revelation scheme. Clearly, we also have $T_{full} \leq T^*$.

We now define a special type of signaling scheme that will be useful for later arguments.

Definition 2 (Segmentation)**.** *A signaling scheme $\sigma = (S, \pi)$ is called a segmentation (or an n-segmentation), if there exist $\underline{q} = q_1 < q_2 < ... < q_n < q_{n+1} = \overline{q}$ with finite n, such that*

$$S = \{s_i\}_{i=1}^n,$$
$$\pi(s_i|q) = \begin{cases} 1 & q \in (q_i, q_{i+1}] \\ 0 & \textit{otherwise} \end{cases}, \forall i.$$

The points $q_1, q_2, \ldots, q_{n+1}$ are called separating points for the segmentation.

In the above definition, we ignore the point $\underline{q}$ as the probability is 0. We now provide a characterization for an n-signal scheme to be a segmentation.

Lemma 1. *A feasible n-signal scheme $\sigma = (S, \pi)$ is a segmentation if and only if, for any two signals $s_i, s_j \in S$, there exists $q' \in (\underline{q}, \overline{q})$ that perfectly separates the two signals, i.e., one of the following conditions must hold:*

- *$\pi(s_i|q) = 0$ when $q \leq q'$ and $\pi(s_j|q) = 0$ when $q > q'$;*
- *$\pi(s_i|q) = 0$ when $q > q'$ and $\pi(s_j|q) = 0$ when $q \leq q'$.*

It should be straightforward to see the equivalence between Definition 2 and the stated conditions. We now show that, for any finite integer n, it is without loss of generality to only consider n-segmentations.

Lemma 2. *If function $H(q) = 1 - F(\frac{p}{q})$ is convex, then for any finite number n, there exists an optimal n-signal scheme that is a segmentation.*

With the above results, we now present the theoretical analysis for the general distributions.

3.1 Performance Analysis for General Distributions

In this subsection, we show that, unfortunately, the performance of using any finitely number of signals can be arbitrarily worse compared to that of the optimal solution. We now present a special problem instance in Example 1. The instance is determined by a parameter k, and we prove the result based on this example by pushing the parameter k towards 0.

Example 1. *Consider an example with $\overline{q} = \frac{p}{\underline{v}}$ and $\underline{q} = \frac{p}{\overline{v}}$. Let $k \in (0, 1)$ be a parameter. the quality distribution $G(q)$ and the buyer type distribution are:*

$$G(q) = \frac{k^q - k^{\underline{q}}}{k^{\overline{q}} - k^{\underline{q}}} \quad \textit{and} \quad F(v) = 1 - Ak^{-\frac{p}{v}} - B,$$

where:

$$A = \frac{k^{\frac{p}{\overline{v}}} k^{\frac{p}{\underline{v}}}}{k^{\frac{p}{\overline{v}}} - k^{\frac{p}{\underline{v}}}} \quad \textit{and} \quad B = -\frac{k^{\frac{p}{\underline{v}}}}{k^{\frac{p}{\overline{v}}} - k^{\frac{p}{\underline{v}}}}.$$

The corresponding density functions are:

$$g(q) = \frac{\ln k}{k^{\overline{q}} - k^{\underline{q}}} k^q \quad \textit{and} \quad f(v) = -pA(\ln k)\frac{k^{-\frac{p}{v}}}{v^2}.$$

Lemma 3. *In Example 1, any n-signal segmentation $\sigma = (S, \pi)$ satisfies:*

$$\lim_{k \to 0_+} \frac{H(\mathbb{E}[q|s_i])\pi(s_i)}{\mathbb{E}[H(q)\pi(s_i|q)]} = 0, \forall s_i \in S,$$

where all expectations are taken over q and $H(q) = 1 - F\left(\frac{p}{q}\right)$.

Lemma 3 will be used in the following result. We are now ready to present the first negative result. The following lemma shows that we cannot guarantee to obtain a certain fraction of the overall buying probability T_{full} of the full information revelation scheme with any signaling scheme that uses finitely many signals.

Lemma 4. *For any n and any $\epsilon > 0$, there exists a problem instance such that:*

$$\frac{T_n^*}{T_{full}} < \epsilon.$$

We provide a sketch of the proof of Lemma 4 here. It is clear that T_n^* is the sum of the performance from n signals under the optimal signaling scheme and each signal is distributed over a certain interval of q according to Lemma 2. Thus, we represent T_{full} as the sum of the corresponding performance from the n intervals. According to Lemma 3, as $k \to 0$, the performance of each signal in T_n^* becomes infinitesimally small compared to the corresponding interval's performance in T_{full}. This completes the proof of Lemma 4. Since lemma 4 has been established, it is obvious that we can obtain the following negative result.

Theorem 1. *For any n and any $\epsilon > 0$, there exists a problem instance such that:*

$$\frac{T_n^*}{T^*} < \epsilon.$$

Proof. The proof is straightforward by considering Example 1. Since $T_{full} \leq T^*$, we clearly have:

$$\lim_{k \to 0_+} \frac{T_n^*}{T^*} \leq \lim_{k \to 0_+} \frac{T_n^*}{T_{full}} = 0. \tag{5}$$

□

3.2 Performance Analysis under Regularity and MHR Conditions

The regularity condition and the MHR condition are standard assumptions in the literature and most commonly used distributions satisfy these conditions. However, even with the regularity condition or the *monotone hazard rate* (MHR) condition, any signaling scheme with finitely many signals still cannot achieve performance comparable to the optimal one.

Definition 3 (Regularity). *A distribution $F(v)$ with density $f(v)$ is said to satisfy the regularity condition, if $v - \frac{1-F(v)}{f(v)}$ is monotonically increasing in v.*

Definition 4 (Monotone Hazard Rate (MHR)). *A distribution $F(v)$ with density $f(v)$ is said to satisfy the monotone hazard rate condition if $\frac{1-F(v)}{f(v)}$ is monotonically decreasing in v.*

Note that if a distribution satisfies the MHR condition, it is also regular. We only show that a distribution satisfying MHR cannot achieve a certain fraction of the optimal performance, and the result for the regularity condition directly follows.

Lemma 5. *Suppose the buyer type's distribution $F(v)$ satisfies the MHR condition. For any n and any $\epsilon > 0$, there exists a problem instance such that:*

$$\frac{T_n^*}{T^*} < \epsilon.$$

Since Lemma 5 is established and MHR condition suffices for regularity, it directly implies the following corollary.

Corollary 1. *Suppose the buyer type's distribution $F(v)$ satisfies the regularity condition. For any n and any $\epsilon > 0$, there exists a problem instance such that:*

$$\frac{T_n^*}{T^*} < \epsilon.$$

3.3 Performance Analysis for (L, c)-Flat Distributions

Despite the negative results presented in Section 3.1 and 3.2, in this subsection, we introduce a mild technical condition called (L, c)-flatness (Definition 5) and show that if the value distribution $f(v)$ is (L, c)-flat, using only two signals already achieves a certain fraction of the optimal performance. It follows that any optimal n-signal scheme where $n \geq 2$ can also achieve performance comparable to the optimal scheme. We also identify a condition under which we can achieve the optimal performance using just a single signal.

Definition 5 ((L, c)-Flatness). *Let L and c be positive constants. A distribution F with density function f is (L, c)-flat, if $f(v)$ is L-Lipschitz and lower bounded by c, i.e., $|f(v_1) - f(v_2)| \leq L|v_1 - v_2|, \forall v_1, v_2 \in V$, and $f(v) \geq c, \forall v \in V$.*

We remark that (L, c)-flatness is a mild condition, as buyers' values are bounded, and many widely used distributions satisfy this condition when truncated to a bounded interval. Examples include truncated Gaussian, lognormal, exponential, and Poisson distributions.

Let T_2^* denote the overall buying probability under the optimal signaling scheme when only two signals are used (i.e., $|S| = 2$). Now, we are ready to present the main result in this section.

Theorem 2. *If the buyer type's density function $f(v)$ is (L, c)-flat, then using only two signals achieves an α fraction of the optimal performance T^*, i.e.,*

$$T_2^* \geq \alpha T^*,$$

where $\alpha = \frac{(M-\underline{q})^2 cp}{(\overline{q}-\underline{q})M^2}$ and $M = \frac{p}{\overline{v}^2(L(\overline{v}-\underline{v})+c)} + \underline{q}$.

Proof. Since it is difficult to directly compare T_2^* and T^*, we derive an upper bound for T^* and a lower bound for T_2^*, and compare these bounds instead.

We first show that under the L-Lipschitz continuity condition, the function $f(v)$ is bounded. Note that:

$$|f(v_1) - f(v_2)| \le L|v_1 - v_2|, \forall v_1, v_2 \in V.$$

And $f(v) \ge c, \forall v \in [\underline{v}, \overline{v}]$, we have:

$$c \le f(v) \le L(\overline{v} - \underline{v}) + c. \tag{6}$$

Since we have assumed that $\underline{q} \ge \frac{p}{\overline{v}}$ before, we have $\frac{p}{q} \le \overline{v}, \forall q \in [\underline{q}, \overline{q}]$. It follows that:

$$H'(q) = \frac{p}{q^2} f\left(\frac{p}{q}\right) \le \frac{\overline{v}^2 \left(L(\overline{v} - \underline{v}) + c\right)}{p}.$$

For convenience, we let $H'_{max} = \frac{\overline{v}^2(L(\overline{v}-\underline{v})+c)}{p}$ and construct a function $\overline{H}(q)$ as follows:

$$\overline{H}(q) = \begin{cases} (q - \underline{q})H'_{max} & q \in [\underline{q}, M] \\ 1 & q \in (M, \overline{q}] \end{cases},$$

where $M = \frac{p}{\overline{v}^2(L(\overline{v}-\underline{v})+c)} + \underline{q} \le \overline{q}$. Note that $\overline{H}(\underline{q}) = H(\underline{q}) = 0$ and $\overline{H}(M) = 1$. It is easy to check that $\overline{H}(q)$ is concave and satisfies $\overline{H}(q) \ge H(q)$.

Given the optimal signaling scheme $\sigma = (S, \pi)$, the probability of a random buyer purchasing an item is:

$$\begin{aligned} T^* &= \int_{s \in S} \pi(s) H(\mathbb{E}[q|s])\, \mathrm{d}s \le \int_{s \in S} \pi(s) \overline{H}(\mathbb{E}[q|s])\, \mathrm{d}s \\ &\le \overline{H}\left(\int_{s \in S} \pi(s)\, \mathbb{E}[q|s]\, \mathrm{d}s\right) = \overline{H}(\mathbb{E}[q]), \end{aligned}$$

where the first inequality is due to $H(q) \le \overline{H}(q)$, and the second inequality follows from the concavity of $\overline{H}(q)$. The above inequality gives an upper bound for T^*.

We next derive the lower bound for T_2^*. When $q \le M$, according to Equation 6, we can obtain the lower bound of $H'(q)$ as follows:

$$H'(q) = \frac{p}{q^2} f\left(\frac{p}{q}\right) \ge \frac{pc}{M^2}.$$

Let $H'_{min} = \frac{pc}{M^2}$. Then, we construct the the following function $\underline{H}(q)$:

$$\underline{H}(q) = \begin{cases} (q - \underline{q})H'_{min} & q \in [\underline{q}, M] \\ (M - \underline{q})H'_{min} & q \in (M, \overline{q}] \end{cases}.$$

One can easily verify that $\underline{H}(q) \le H(q)$.

Now, we construct two signals s_1 and s_2, where:

$$\pi(s_1|q) = \begin{cases} 1 & q \in [\underline{q}, M] \\ 0 & q \in [M, \overline{q}] \end{cases} \text{ and } \pi(s_2|q) = \begin{cases} 0 & q \in [\underline{q}, M] \\ 1 & q \in [M, \overline{q}] \end{cases}.$$

Assume that the overall buying probability under signals s_1 and s_2 is T_2, we have $T_2 \le T_2^*$ and:

$$T_2 = \Pr(s_1) H\left(\mathbb{E}\left[q|s_1\right]\right) + \Pr(s_2) H\left(\mathbb{E}\left[q|s_2\right]\right).$$

We construct the following linear function $\hat{H}(q)$:

$$\hat{H}(q) = (q - \underline{q})\hat{H}',$$

where $\hat{H}' = \frac{(M-\underline{q})cp}{(\overline{q}-\underline{q})M^2}$. It is easy to check that $\hat{H}(q) \le \underline{H}(q) \le H(q)$. Then we have:

$$\begin{aligned} T_2^* \ge T_2 &= \Pr(s_1) H\left(\mathbb{E}\left[q|s_1\right]\right) + \Pr(s_2) H\left(\mathbb{E}\left[q|s_2\right]\right) \\ &\ge \Pr(s_1) \hat{H}\left(\mathbb{E}\left[q|s_1\right]\right) + \Pr(s_2) \hat{H}\left(\mathbb{E}\left[q|s_2\right]\right) \\ &\ge \hat{H}\left(\mathbb{E}[q]\right). \end{aligned}$$

For convenience, we let $\alpha = \frac{(M-\underline{q})^2 cp}{(\overline{q}-\underline{q})M^2}$, which is a positive constant. Then we can get that:

$$\frac{T_2^*}{T^*} \ge \frac{\hat{H}\left(\mathbb{E}[q]\right)}{\overline{H}\left(\mathbb{E}[q]\right)} \ge \alpha.$$

□

Now we identify a special case where revealing no information is optimal.

Lemma 6. *If $v^2 f(v)$ is monotonically increasing in v, then using just a single signal (which is equivalent to revealing no information) achieves the optimal performance.*

4 Experiments

We conduct experiments and report the results in this section. All experiments are run on a server equipped with a 13th Gen Intel(R) Core(TM) i9-13900K CPU and 128 GB of RAM. Since the computer cannot handle continuous distribution, we discretize the item quality space Q into m_q different qualities and the buyers' value space V into m_v different values. The discrete problem can be formulated as a linear program.

We also give an approximation result for the discrete case:

Theorem 3. *Let m_v be the number of different values. Then for any discrete problem instance, we have*

$$\frac{T_2^*}{T^*} \ge \frac{1}{m_v},$$

where T_2^* is the performance of the optimal 2-signal scheme and T^* is the optimal performance. And the bound is tight.

Since the bound in Theorem 3 is tight, there exist problem instances with $\frac{T_2^*}{T^*} = \frac{1}{m_v}$. As m_v approaches infinity, the performance ratio goes to 0 for such instances, which is negative. Therefore, Theorem 3 can be viewed as the discrete version of Theorem 1.

We next present our experiment results. All LPs are solved with Gurobi 11.0 [Gurobi Optimization, LLC, 2023]. We conduct experiments on (i) instances satisfying (L, c)-flatness, and (ii) instances that are generated randomly without any constraints.

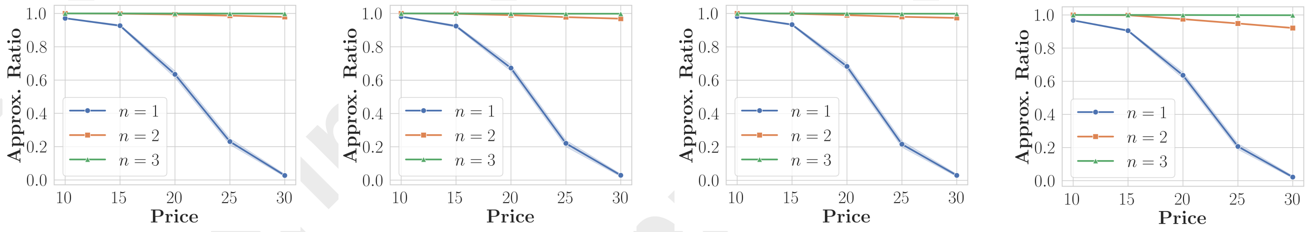

(a) $L = 1$ and $c = 10^{-3}$ (b) $L = 10$ and $c = 10^{-4}$ (c) $L = 100$ and $c = 10^{-5}$ (d) General distributions

Figure 1: The approximation ratio of a small number of signals under (L, c)-flat distributions and general distributions.

4.1 Experiment Results for (L, c)-Flat Distributions

In this experiment, we analyze the approximation ratio of a limited number of signals (ranging from 1 to 3) compared to the optimal performance when the buyers' value distribution f is (L, c)-flat. The experiments are conducted with different prices $p \in \{10, 15, 20, 25, 30\}$ and set $m_v = m_q = 20$. We randomly sample $\underline{v} \in [1, 3], \overline{v} \in [6, 8], \underline{q} \in [1, 3]$, $\overline{q} \in [6, 8]$ and equidistantly choose value and quality from their corresponding space.

We enumerate L over the set $\{1, 10, 100\}$ and c over the set $\{10^{-3}, 10^{-4}, 10^{-5}\}$, conducting experiments for each combination. We randomly sample 1000 instances for each combination of p, L and c to obtain the average results with statistical confidence. For each instance, we calculate the approximation ratio as the ratio between the optimal value with a limited number of signals and the optimal value with sufficient signals. The optimal value for different numbers of signals is computed by solving a series of linear programs.

The results for combinations $(L = 1, c = 10^{-3})$, $(L = 10, c = 10^{-4})$ and $(L = 100, c = 10^{-5})$ are presented in Figures 1a, 1b, and 1c, while results for other combinations exhibit similar outcomes. It is evident that, as the price increases, the approximation ratio for a single signal decreases and eventually approaches zero. In contrast, the approximation ratio of two or three signals remains consistently high, always exceeding 0.95. This indicates that a third party (e.g., a used car platform) can achieve near-optimal performance using only two signals, which, in practical scenarios, correspond to recommendations for buyers to either purchase or not purchase. Note that a single signal implies that the third party reveals no information, which means that revealing information to buyers significantly increases the probability of buyers making a purchase under (L, c)-flat distributions.

4.2 Experiment Results for General Distributions

In the previous section, we examined the approximation ratio of a limited number of signals under (L, c)-flat distributions. However, we are also interested in the average approximation ratio under general distributions without the (L, c)-flatness constraint. Although the theoretical results in Theorem 1 indicate that the approximation ratio can be arbitrarily worse in certain problem instances, it's important to note that these instances represent extreme cases. To evaluate the average approximation ratio under general distributions, we randomly generate distributions for each instance and conduct experiments across a large number of instances.

In this experiment, for each instance, we also randomly sample $\underline{v} \in [1, 3], \overline{v} \in [6, 8], \underline{q} \in [1, 3]$, $\overline{q} \in [6, 8]$, and set $m_v = m_q = 20$. Then, we equidistantly select value and quality from their corresponding space. To ensure the randomness of the generated distributions, we sample f and g from $[0, 1]$ and subsequently normalize them. Additionally, we also experiment with different prices $p \in \{10, 15, 20, 25, 30\}$. For each price p, we randomly sample 1000 instances and conduct experiments to calculate the average results with statistical confidence.

As shown in Figure 1d, under general distributions, the approximation ratio for two signals, although not as strong as the results under the (L, c)-flat distributions, consistently exceeds 0.8. This finding suggests that, even under general distributions, a third party can still use just two or three signals to nearly approximate the optimal performance, despite the theoretical possibility of extremely poor approximation ratios in worst-case scenarios. Similarly, as the price increases, the approximation ratio of a single signal (equivalent to revealing no information) decreases and eventually approaches zero. This highlights the importance of a third party revealing information to buyers, as it significantly increases the probability of purchases even under general distributions.

5 Conclusion

We explored the extent to which a third party can achieve a constant approximation of the optimal performance through a signaling scheme with a limited number of signals. We first showed that no signaling scheme can achieve a certain fraction of the optimal performance, even if the buyers' value distribution satisfies the commonly used regularity or monotone hazard rate conditions. However, we demonstrated that when the value distribution is (L, c)-flat, the third party can achieve a certain fraction of the optimal performance by using only two signals. Additionally, we conducted extensive experiments to verify our theoretical results. Our experiments indicated that even under general distributions, a small number of signals (two or three signals) can nearly approximate the optimal performance. This finding suggests that the existence of a third party that reveals information to buyers plays a crucial role in increasing the probability of purchases.

Acknowledgments

We thank all the anonymous reviewers for their helpful comments. This work is supported in part by the National Natural Science Foundation of China (No. 72192805), Public Computing Cloud, Renmin University of China, the fund for building world-class universities (disciplines) of Renmin University of China.

References

[Akerlof, 1978] George A Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in economics*, pages 235–251. Elsevier, 1978.

[Alaei *et al.*, 2013] Saeed Alaei, Hu Fu, Nima Haghpanah, and Jason Hartline. The simple economics of approximately optimal auctions. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 628–637. IEEE, 2013.

[Alaei *et al.*, 2019] Saeed Alaei, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*, 118:494–510, 2019.

[Alonso and Câmara, 2016] Ricardo Alonso and Odilon Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.

[Aybas and Turkel, 2019] Yunus C Aybas and Eray Turkel. Persuasion with coarse communication. *arXiv preprint arXiv:1910.13547*, 2019.

[Blumrosen and Dobzinski, 2014] Liad Blumrosen and Shahar Dobzinski. Reallocation mechanisms. *arXiv preprint arXiv:1404.6786*, 2014.

[Cai *et al.*, 2016] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 926–939, 2016.

[Castiglioni *et al.*, 2023] Matteo Castiglioni, Andrea Celli, and Nicola Gatti. Public bayesian persuasion: being almost optimal and almost persuasive. *Algorithmica*, 85(9):2885–2921, 2023.

[Chawla *et al.*, 2014] Shuchi Chawla, Hu Fu, and Anna Karlin. Approximate revenue maximization in interdependent value settings. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 277–294, 2014.

[Chen and Xie, 2005] Yubo Chen and Jinhong Xie. Third-party product review and firm marketing strategy. *Marketing science*, 24(2):218–240, 2005.

[Chen *et al.*, 2012] Yubo Chen, Yong Liu, and Jurui Zhang. When do third-party product reviews affect firm value and what can firms do? the case of media critics and professional movie reviews. *Journal of Marketing*, 76(2):116–134, 2012.

[Colini-Baldeschi *et al.*, 2016] Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1424–1443. SIAM, 2016.

[Dughmi *et al.*, 2016] Shaddin Dughmi, David Kempe, and Ruixin Qiang. Persuasion with limited communication. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 663–680, 2016.

[Dughmi, 2019] Shaddin Dughmi. On the hardness of designing public signals. *Games and Economic Behavior*, 118:609–625, 2019.

[Eliashberg and Shugan, 1997] Jehoshua Eliashberg and Steven M Shugan. Film critics: Influencers or predictors? *Journal of marketing*, 61(2):68–78, 1997.

[Gradwohl *et al.*, 2022] Ronen Gradwohl, Niklas Hahn, Martin Hoefer, and Rann Smorodinsky. Algorithms for persuasion with limited communication. *Mathematics of Operations Research*, 47(3):2520–2545, 2022.

[Gurobi Optimization, LLC, 2023] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.

[Hartline and Roughgarden, 2009] Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009.

[Huang *et al.*, 2015] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 45–60, 2015.

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

[Kang and Vondrák, 2019] Zi Yang Kang and Jan Vondrák. Fixed-price approximations to optimal efficiency in bilateral trade. *Available at SSRN 3460336*, 2019.

[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.

[Pei and Mayzlin, 2021] Amy Pei and Dina Mayzlin. Influencing the influencers. *Available at SSRN 3376904*, 2021.

[Shen and Tang, 2017] Weiran Shen and Pingzhong Tang. Practical versus optimal mechanisms. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 78–86, 2017.

[Tian, 2022] Zijun Tian. How should firms cooperate with honest influencers? *Available at SSRN 4173460*, 2022.

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

[Yao, 2014] Andrew Chi-Chih Yao. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 92–109. SIAM, 2014.

[Zheng and Chen, 2021] Shuran Zheng and Yiling Chen. Optimal advertising for information products. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 888–906, 2021.