Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Fair Price Discrimination*

Siddhartha Banerjee† Kamesh Munagala‡ Yiheng Shen† Kangning Wang§

Abstract

A seller is pricing identical copies of a good to a stream of unit-demand buyers. Each buyer has a value on the good as his private information. The seller only knows the empirical value distribution of the buyer population and chooses the revenue-optimal price. We consider a widely studied third-degree price discrimination model where an information intermediary with perfect knowledge of the arriving buyer's value sends a signal to the seller, hence changing the seller's posterior and inducing the seller to set a personalized posted price. Prior work of Bergemann, Brooks, and Morris (American Economic Review, 2015) has shown the existence of a signaling scheme that preserves seller revenue, while always selling the item, hence maximizing consumer surplus. In a departure from prior work, we ask whether the consumer surplus generated is fairly distributed among buyers with different values. To this end, we aim to maximize functions of buyers' welfare that reward more balanced surplus allocations.

Our main result is the surprising existence of a novel signaling scheme that *simultaneously* 8-approximates *all* welfare functions that are non-negative, monotonically increasing, symmetric, and concave, compared with any other signaling scheme. Classical examples of such welfare functions include the utilitarian social welfare, the Nash welfare, and the max-min welfare. Such a guarantee cannot be given by consumer-surplus-maximizing schemes – which are the ones typically studied in the literature. In addition, our scheme is socially efficient, and has the fairness property that buyers with higher values enjoy higher expected surplus, which is not always the case for existing schemes.

1 Introduction

Imagine a seller with infinite supply of a good. The seller wishes to sell the good to a population of unit-demand buyers with standard quasi-linear utilities. The seller knows the empirical distribution $\mathcal{D}$ of the buyer valuations and chooses a revenue-maximizing price to sell the good. In this paper, we consider this simple setting, but with a twist: there is an additional information intermediary who can segment the market and help the seller price-discriminate. Information intermediaries for price discrimination were first considered by [4] and our work proposes and studies new desiderata for them.

Such intermediaries are becoming ubiquitous in modern two-sided e-commerce platforms. Consider for example ad exchanges [15, 49, 40, 42], where the platform acts as an intermediary between buyers (in this case advertisers) and sellers (in this case, publishers controlling the ad slot). The intermediary wants the best for both sides; however, as in classical auctions, the seller – not the intermediary – controls the price at which trade happens. Other examples include retail platforms such as Amazon marketplace, who also effectively serve as intermediaries – they merely facilitate the trade, but do not control the prices.

In such settings, the platform can use machine learning and its vast trove of data on buyer behavior to accurately predict the value of buyers. It can then choose to reveal information about the current buyer to the seller in order to influence the trade. This information (or *signal*) leads to the seller updating its prior $\mathcal{D}$ over buyer values to a posterior distribution over values given the signal. The seller now posts the optimal (revenue-maximizing) price for this posterior. Such information revelation is termed signaling or *third-degree*

*The full version of the paper can be accessed at https://arxiv.org/abs/2305.07006.

†School of Operations Research and Information Engineering, Cornell University, Ithaca, NY 14850. Email: sbanerjee@cornell.edu.

‡Department of Computer Science, Duke University, Durham, NC 27708-0129. Emails: kamesh@cs.duke.edu, yiheng.shen@duke.edu. Supported by NSF grant CCF-2113798.

§Department of Computer Science, Stanford University, Stanford, CA 94305. Email: knwang@stanford.edu.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

price discrimination[1] and is a special case of *Bayesian persuasion* [31]. Note that in practice, the seller and the intermediary could be the same entity, such as a retail or ride-share platform that wants to use buyer information to segment the market and perform price discrimination.

To understand this setting better, consider two extremes: At one extreme, the intermediary can choose to reveal no information to the seller, in which case the seller's posterior remains $\mathcal{D}$. Therefore, the seller posts the Myerson price [41] $p^{\mathsf{My}} = \arg\max_p p \cdot \Pr_{v\sim\mathcal{D}}[v \geq p]$ on $\mathcal{D}$ and collects revenue $R^{\mathsf{My}} = p^{\mathsf{My}} \cdot \Pr_{v\sim\mathcal{D}}[v \geq p^{\mathsf{My}}]$. Since trade does not happen if the buyer's private value v is below p^{My}, this scheme is generally *inefficient* – the consumer (buyer) surplus $\mathbb{E}_{v\sim\mathcal{D}}[(v - p^{\mathsf{My}})^+]$ plus seller revenue R^{My} is less than the maximum possible total surplus, $\mathbb{E}_{v\sim\mathcal{D}}[v]$.

At another extreme is full information revelation or first-degree price discrimination, where the intermediary reveals the actual buyer value v to the seller. In this case, the seller's posterior collapses to the deterministic value v. The seller can now post price slightly below v, so that trade always happens. However, this efficiency comes at a cost – the buyer now obtains zero surplus (their value minus price paid), while the seller's revenue becomes equal to the total surplus, $\mathbb{E}_{v\sim\mathcal{D}}[v]$. Note that in the no-signaling case discussed above, the consumer surplus $\mathbb{E}_{v\sim\mathcal{D}}[(v - p^{\mathsf{My}})^+]$ could be positive – thus between these two schemes, from a utilitarian point of view, no-signaling is better for the buyers, while full-revelation is better overall.

1.1 Optimal Signaling and Fairness Signaling clearly helps the seller since they can always obtain at least as much revenue R^{My} as in no-signaling (e.g. by ignoring the signal). What is less clear is whether signaling can improve consumer surplus at all. In a remarkable result, [4] showed the existence of "buyer-optimal" signaling schemes in the following sense: The seller's expected revenue remains the same as in no-signaling (i.e., R^{My}, which is the *minimum* possible under any signaling scheme), while trade is always efficient (i.e., the item always sells), which means that the sum of the consumer surplus and the seller revenue is the maximum possible total surplus, $\mathbb{E}_{v\sim\mathcal{D}}[v]$. Hence, the consumer surplus must be as large as it could possibly be.

This is a beautiful result, but is unsatisfying upon closer inspection. Note that while the proposed scheme maximizes consumer surplus, it is not the unique such scheme [4, 14, 35]. Are all ways of splitting this aggregate surplus among buyers equal, even if this gain in surplus is "subsidized" more by a particular group of buyers? We think not – maximizing the utilitarian total consumer surplus should not be the sole consideration; it is natural to also desire that *price discrimination is fair at the level of individual buyers* – but how should we formalize this?

A first idea is to require some form of **monotonicity** in the surplus split. Let $\mathcal{CS}_v$ be the expected consumer surplus that a signaling scheme provides to a buyer of value v; we could now require that $\mathcal{CS}_v \leq \mathcal{CS}_{v'}$ whenever $v < v'$. This is true in the absence of signaling (as fixed pricing is monotone), and so should perhaps be expected to hold in the presence of signaling. It also captures some sense of envy in price discrimination – a buyer with larger value should not envy the surplus seen by a buyer with smaller value. We show via examples in Section 2.2 that even this very natural constraint rules out some existing buyer-optimal schemes.[2]

An alternative and more wide-reaching fairness requirement is given by the following paradigm:

Universal Welfare Maximization (and Majorization). Consider the surplus vector where its j^{th} dimension is the expected surplus of the j^{th} buyer. Note that the surplus vector is not the vector of *different* surplus values. Instead, each dimension represents the expected surplus of a single buyer. For example, if there is 1 buyer with expected surplus 1 and 9 buyers with expected surplus 2, we are transforming the surplus into a 10-dimension vector rather than a 2-dimension vector. A general welfare function[3] takes the surplus vector as input, and outputs a non-negative real number (higher is better). We restrict to welfare functions that are symmetric, non-decreasing, and concave: Symmetry ensures equal treatment to all buyers; non-decreasing ensures that Pareto improvements are desirable; and concavity is a common fairness consideration favoring balanced allocations. Common examples of such welfare functions include the utilitarian social welfare function, the Nash

[1] It is termed "third-degree" price discrimination because the seller or intermediary divides the market into segments, each with its own price. In contrast, in first-degree price discrimination, the seller has perfect information and charges buyers exactly their value, while in second-degree price discrimination, the seller sells similar yet 'different' goods (differing in quality/quantity) at different prices.

[2] We note that the scheme for continuous priors in [4] is both monotone and buyer-optimal.

[3] Our welfare notion is based only on each buyer's expected surplus, and does not consider the seller's surplus. We can easily extend our welfare notion to also take into account the seller's surplus and obtain similar results, by mixing our schemes with the seller-optimal one.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

welfare function, and the max-min welfare function. A fair signaling scheme could be defined as one which maximizes such a welfare function; however, it is unclear how to unambiguously pick one welfare function among the numerous possibilities.

What would be ideal is if there is a *universal* scheme that is optimal (or at least, approximately optimal) for *all* such welfare functions. This universal maximization of concave functions is closely related to *majorization* [32, 27] and its approximate form: α-majorization (see e.g. [24, 23]).

1.2 Our Results The main question we ask is:

> In third-degree price discrimination, how close can we get to a *universally-fair* signaling scheme, i.e., one which is (monotone and) near-optimal for any welfare function?

At the outset, one might be pessimistic: For resource allocation and stochastic optimization problems [24, 37, 9], typically $\alpha = \omega(1)$, where α is the approximation factor for majorization (and hence for universal welfare maximization). Indeed, as we show in Section 4, any buyer-optimal signaling scheme in the sense of [4] cannot be α-majorized for *any* given constant α. Given this, one may wonder if universal welfare maximization is too strong a condition to expect.

Our main result is a surprising new signaling scheme that shows the following theorem:

THEOREM 1.1. (FORMALLY STATED AS THEOREM 3.1) *For any prior $\mathcal{D}$, there is a signaling scheme that is* 8*-majorized, and hence it simultaneously* 8*-approximates all non-negative, increasing, symmetric and concave welfare functions, compared with any other signaling scheme. Further, this scheme is monotone and efficient, and can be computed in time polynomial in the size of the support of $\mathcal{D}$.*

Our main theorem therefore shows that we can be (near)-universally-fair (i.e., near-optimal for any welfare function). This signaling scheme sacrifices some consumer surplus to achieve this guarantee; however, as mentioned above, this sacrifice is necessary – as we show in Section 4, any exactly buyer-optimal scheme is not α-majorized for any constant α. Further note that by definition, our scheme is also 8-approximately buyer-optimal. We complement our 8-approximation with a lower bound of 1.5 in Section 4: There is no signaling scheme that is α-majorized by every other signaling scheme for $\alpha < 1.5$.

At a technical level, the proof of Theorem 1.1 constructs a very different signaling scheme from prior work on price discrimination. The scheme is composed of signals such that each of them induces a posterior as a distribution over at most two values. We first decompose the prior $\mathcal{D}$ into a collection of such signals, and show that this collection 4-approximates prefix sums of consumer surplus when sorted on buyer value. We then apply a novel ironing procedure to modify the signals so that the resulting scheme is approximately majorized, while losing an additional factor of 2. Both steps are non-trivial, and together yields an 8-majorized scheme that is also monotone (and socially efficient).

1.3 Related Work Our model of third-degree price discrimination is a special case of *information design* (see [7]) where an information mediator provides information to impact the behavior of agents. This has also been termed *signaling* or *persuasion* in literature. (See [16].) In *Bayesian persuasion* first proposed by [31], there is one agent called the *receiver* who receives additional information from a better-informed *sender*. Given the signal, the receiver computes their posterior over the state of nature and chooses an action to maximize their own utility. The sender can design the signals so that the receiver, acting in her own interest, maximizes some utility function the sender cares about. This problem has been widely studied in various contexts [17, 18, 2, 4, 10, 47, 26].

In the setting we consider, the sender is the intermediary, while the receiver is the seller that maximizes their revenue given the signal. This was first considered by [4], who showed buyer-optimal signaling schemes that preserve seller revenue while transferring the rest of the surplus to the buyers. Subsequently, it was shown by [14] that the set of all buyer-optimal signaling schemes can be specified by a linear program. Several works [45, 8, 39, 6, 1, 35] consider various extensions to the basic single seller/buyer setting, and show exact/approximate buyer-optimality under various assumptions.

The concept of majorized vectors has existed for a long time [32, 27], and is equivalent to solutions that simultaneously maximize symmetric concave functions of the coordinates. In the context of resource allocation and routing problems, an approximate version of this concept was defined by [24], and subsequently shown by [23] to be equivalent to solutions that simultaneously approximately maximize every symmetric concave function of

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

the coordinates; see also [37, 9]. It was shown by [24] that the best approximation factor is the solution of a linear program. However, the approximation factor is problem-dependent and typically logarithmic in the number of coordinates for general routing problems. The surprising aspect of our paper is that this factor is only a constant for the price discrimination problem, and is achieved by a very simple signaling scheme. This is similar in spirit to recent results in metric distortion of voting rules, where the Copeland rule is 5-majorized [22], and a novel rule of [21] is 3-majorized, by any other rule.

Apart from approximate majorization, researchers have proposed other fairness notions in advertising auctions, such as enforcing similar position allocations for similar bids [11, 12, 20], and similar assignment of ads [30] or payoffs [46] for similar users.

[48] consider fairness in price discrimination by imposing a bound on the ratio or difference in prices that the seller is allowed to charge to different buyers (akin to monotonicity). They assume a perfectly informed seller (first-degree price discrimimation) and derive a unique optimal pricing strategy as well as characterize the tradeoffs for different buyer value distributions. Our work in contrast focuses on the more involved objective of majorization, and furthermore, we do not assume the seller is perfectly informed (third-degree price discrimination).

Our work connects to the larger body of work on fairness in machine learning, where again, optimality in the sense of overall risk minimization (ERM) can lead to systematic unfairness [33, 34, 29, 19, 36]. Much of this work focuses on the tradeoffs between efficiency and fairness. As machine learning systems become more pervasive, it becomes important to consider not just their direct impact, but also their impact to downstream applications when they are embedded in a larger system. In our case, the larger system is a marketplace platform that uses machine learning to predict buyer values and help sellers price-discriminate. Our results show that naïvely maximizing surplus can be unfair, while different mechanisms can achieve good tradeoffs between efficiency and fairness.

2 Preliminaries

2.1 Basic Setting

Seller and Buyers. A monopolistic seller of a good has infinite supply, and wants to price them so as to maximize her revenue. There are a finite number of buyers in the market. Each buyer is interested in buying at most one copy of the good, and has a value for the good given by a positive real number. A buyer chooses to buy if and only if the price is at most his value. We henceforth focus on discrete empirical distributions over buyer valuations; in particular, we consider distributions with support size n over values $v_1 < v_2 < \cdots < v_n$. (We write $v_0 := 0$ to simplify notation.) For any distribution $\mathcal{P}$, we use $f_{\mathcal{P}}(v)$ to denote the probability mass function: $f_{\mathcal{P}}(v) := \Pr_{v' \sim \mathcal{P}}[v' = v]$, and define the cumulative distribution function (CDF) $F_{\mathcal{P}}(v) := \Pr_{v' \sim \mathcal{P}}[v' \leq v]$ and complementary CDF $G_{\mathcal{P}}(v) := \Pr_{v' \sim \mathcal{P}}[v' \geq v]$.

Let $\mathcal{D}$ denote the empirical distribution of buyer valuations. The seller knows the distribution $\mathcal{D}$, but not the actual value of each buyer. Consequently, without additional information, the seller chooses a common price p^{My} (sometimes called the Myerson price [41]) for all buyers such that $p = p^{\mathsf{My}}$ maximizes $p \cdot G_{\mathcal{D}}(p)$.

Price Discrimination via an Information Intermediary. The main idea in the work of [4] is that in this setting, one can model the effects of *price discrimination* by considering an exogenous intermediary who provides some additional signal to the seller about each buyer, enabling the seller to modify the price offered to that buyer. We now formalize this as a game among the intermediary, the seller and the buyers.

We assume the information intermediary knows $\mathcal{D}$ as well as the exact value of each buyer. Independently, for each buyer, the intermediary sends a *signal* about the buyer's value to the seller via some chosen *signaling scheme*: a (potentially randomized) mapping from a value in $\{v_1, v_2, \ldots, v_n\}$ to some set of signals $[Q]$. Crucially, the intermediary commits to a scheme upfront, and the scheme is known to the seller.

From the perspective of the seller, since all agents are a priori indistinguishable, the effect of receiving a signal is to update the seller's belief over the buyer's value from $\mathcal{D}$ to some new distribution $\mathcal{S}_q$ ($q \in [Q]$) over possible values. Consequently, with a slight abuse of notation, instead of defining a signaling scheme in terms of the mapping from value to signal, we directly define it in terms of the resulting posterior distributions $\mathcal{S}_q$ corresponding to each q, as well as the resulting distribution over these signals. Formally:

DEFINITION 2.1. (**Signal; Signaling Scheme**) *A* signal $\mathcal{S}$ *updates the seller's belief over a buyer's value from* $\mathcal{D}$ *to some new distribution* $\mathcal{S}$. *A* signaling scheme $\mathcal{Z} = \{(\mathcal{S}_q, \gamma_q)\}_{q \in [Q]}$ *is a collection of* Q *signals* $\{\mathcal{S}_q\}_{q \in [Q]}$ *with weights* $\{\gamma_q\}_{q \in [Q]}$ *that satisfy: (1)* $\gamma_q \geq 0$ *and* $\sum_{q \in [Q]} \gamma_q = 1$*; and moreover (2)* $\sum_{q \in [Q]} \gamma_q \mathcal{S}_q = \mathcal{D}$.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

We note again that $\mathcal{S}_q$ denotes both the q^{th} signal in the signaling scheme, and the posterior of the seller after receiving the q^{th} signal. The constraints in Definition 2.1 ensure that the signaling scheme $\mathcal{Z}$ is *Bayes plausible* [31], i.e., that the expected posterior is equal to the prior. Given a signaling scheme as defined above, it is easy to construct the random mapping from values to signals: each v is mapped to $\mathcal{S}_q$ with probability $\gamma_q f_{\mathcal{S}_q}(v)/f_{\mathcal{D}}(v)$.

Outcomes under Signaling. After receiving signal $\mathcal{S}_q$ from the intermediary, the seller offers the buyer a new price $p^*_{\mathcal{S}_q}$ based on the new posterior $\mathcal{S}_q$ satisfying

$$p^*_{\mathcal{S}_q} = \arg\max_v v \cdot G_{\mathcal{S}_q}(v).$$

The resulting expected gains from trade are split between the buyer and the seller as:

- *Producer surplus* (or revenue) of the seller: $\mathcal{R}(\mathcal{S}_q) = \max_v v \cdot G_{\mathcal{S}_q}(v) = p^*_{\mathcal{S}_q} \cdot G_{\mathcal{S}_q}(p^*_{\mathcal{S}_q})$.
- *Consumer surplus* of the buyer with value v: $\mathcal{CS}_v(\mathcal{S}_q) = \mathbb{1}[v \geq p^*_{\mathcal{S}_q}] \cdot (v - p^*_{\mathcal{S}_q})$.

In the event that $v \cdot G_{\mathcal{S}_q}(v)$ are maximized at multiple points, we assume the seller breaks ties by choosing the lowest tied price.[4] Furthermore, we can now define the *expected* outcomes under a given signaling scheme: The expected consumer surplus of a buyer with value v is the expectation of that buyer's surplus on all signals that the seller might receive from the intermediary. Similarly we can define the expected seller revenue.

Definition 2.2. (Expected Outcomes under Signaling) *Given a signaling scheme* $\mathcal{Z} = \{(\mathcal{S}_q, \gamma_q)\}_{q\in[Q]}$, *the* expected consumer surplus *of a buyer with value* v *under* $\mathcal{Z}$ *is:*

$$\mathcal{CS}_v(\mathcal{Z}) = \sum_{q\in[Q]} \mathcal{CS}_v(\mathcal{S}_q) \cdot \frac{\gamma_q \cdot f_{\mathcal{S}_q}(v)}{f_{\mathcal{D}}(v)}. \tag{2.1}$$

Moreover, the overall expected consumer surplus under $\mathcal{Z}$ *is* $\mathcal{CS}(\mathcal{Z}) = \sum_v f_{\mathcal{D}}(v) \cdot \mathcal{CS}_v(\mathcal{Z})$.

Similarly, the seller's expected revenue *is given by:*

$$\mathcal{R}(\mathcal{Z}) = \sum_q \mathcal{R}(\mathcal{S}_q) \cdot \gamma_q. \tag{2.2}$$

We illustrate our setting, signaling schemes, and the above metrics with the following running example; Figs. 1 and 2 show different signaling schemes for this setting.

Example. [Running example] The buyer values are given by $\langle 1, 2, 5, 6\rangle$ with distribution on this support being $\mathcal{D} = \langle 0.25, 0.25, 0.25, 0.25\rangle$. The revenue under each of the posted prices is $\langle 1, 1.5, 2.5, 1.5\rangle$, and thus 5 is the Myerson price under $\mathcal{D}$, resulting in revenue $\mathcal{R}^{\mathsf{My}} = 2.5$.

Fig. 1 illustrates one particular signaling scheme $\mathcal{Z}_1^{\mathsf{Ex}}$ for this setting (based on the construction of [4]). Here, it is easy to check that 5 is an optimal price in all the signals, resulting in (seller-optimal) revenue of $\mathcal{R}(\mathcal{Z}_1^{\mathsf{Ex}}) = 5 \cdot (f_{\mathcal{D}}(5) + f_{\mathcal{D}}(6)) = 2.5$. To compute the consumer surplus, take $v_3 = 5$ as an example: the expected consumer surplus of a buyer with value 5 is

$$\mathcal{CS}_5(\mathcal{Z}_1^{\mathsf{Ex}}) = \left(\frac{1}{60} \cdot (5-1) + \frac{1}{90} \cdot (5-2)\right) \Big/ \left(\frac{1}{4}\right) = 0.4.$$

□

2.2 Global and Per-Agent Performance Metrics of Signaling Schemes

[4] Note that we can avoid ties by slightly perturbing each signal, without changing the message of our results.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

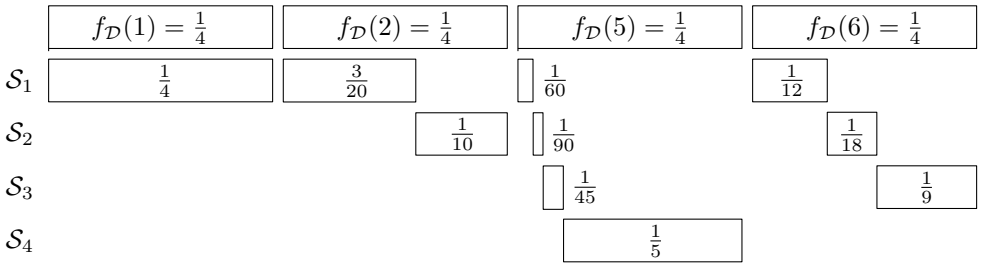

Figure 1: *Signaling scheme in Section 2.1: The distribution $\mathcal{Z}_1^{\mathsf{Ex}}$ is drawn as rectangles in the first row. Each subsequent row corresponds to a signal under $\mathcal{Z}$. The rectangles under each v_i's column indicate the mass $\gamma_q f_{\mathcal{S}_q}(v_i)$ placed on v_i in each signal. We can recover each signal and its weight by normalizing; for example, signal $\mathcal{S}_1$ satisfies $f_{\mathcal{S}_1}(1) = \frac{1}{2}$, $f_{\mathcal{S}_1}(2) = \frac{3}{10}$, $f_{\mathcal{S}_1}(5) = \frac{1}{30}$, $f_{\mathcal{S}_1}(6) = \frac{1}{6}$ and has weight $\gamma_1 = \frac{1}{2}$.*

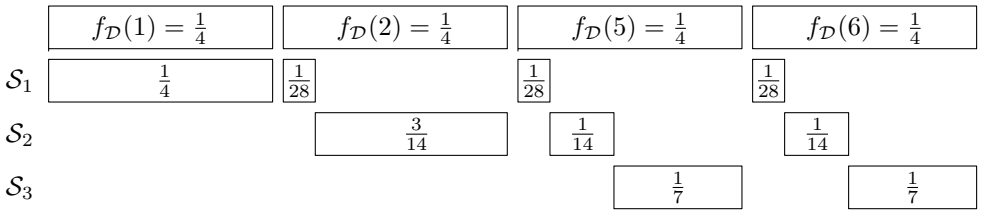

Figure 2: *A completely different signaling scheme $\mathcal{Z}_2^{\mathsf{Ex}}$ for the instance given in Section 2.1: The notations are the same as those in Fig. 1. Note that while the two schemes are very different, they both are efficient (i.e., the item is always sold), and have the same revenue and same overall consumer surplus (although the distribution of the overall surplus between different values is different; see Section 2.2).*

Utilitarian Metrics (Efficiency and Consumer Surplus). For a signaling scheme $\mathcal{Z}$ to be efficient (i.e., to maximize the gains from trade), it needs to ensure the item is always sold. This corresponds to requiring that for each signal $q \in [Q]$, the optimal price posted by the seller under $\mathcal{S}_q$ is the smallest value with non-zero probability in the support of $\mathcal{S}_q$. If this holds, then any buyer will always accept the posted price and the item is always sold.

Given any discrete distribution $\mathcal{P}$, the *lowest value in the support* of $\mathcal{P}$ is defined as $\underline{v}_{\mathcal{P}} := \min\{v \mid f_{\mathcal{P}}(v) > 0\}$. Note that since we focus on distributions over a finite support, the minimum exists. Now we can formally define the condition for a signaling scheme to be efficient:

Definition 2.3. (Efficient Signaling Scheme) *A signaling scheme $\mathcal{Z} = \{(\mathcal{S}_q, \gamma_q)\}_{q\in[Q]}$ is* efficient *if*

$$\underline{v}_{\mathcal{S}_q} = \arg\max_v v \cdot G_{\mathcal{S}_q}(v) \qquad \forall\, q \in [Q].$$

While efficiency ensures that a signaling scheme maximizes the overall gains-from-trade, it does not specify how the surplus is divided. In particular, revealing the buyer's true value to the seller is an efficient scheme, but results in the seller getting the full surplus. An alternative is to maximize the overall consumer surplus $\mathcal{CS}(\mathcal{Z})$. From the above definitions, it is easy to see that given $\mathcal{D}$ one can write the problem of constructing a signaling scheme $\mathcal{Z}$ that maximizes $\mathcal{CS}(\mathcal{Z})$ via a linear program; it is not clear however what guarantees such a scheme has, or even, if it is efficient. The surprising result of [4] is that this is indeed the case:

Proposition 2.1. (From [4]) *For any given $\mathcal{D}$, let $\mathcal{R}^{\mathsf{My}} = \max_v v \cdot G_{\mathcal{D}}(v)$ denote the optimal revenue without signaling (i.e., the "Myerson revenue"). Then there exist efficient signaling schemes $\mathcal{Z}$ under which $\mathcal{R}(\mathcal{Z}) = \mathcal{R}^{\mathsf{My}}$.*

Note that since the seller can always get $\mathcal{R}^{\mathsf{My}}$ under any signaling scheme (by ignoring the signal and posting p^{My}), and since the signaling scheme is efficient, it must have maximized $\mathcal{CS}(\mathcal{Z})$. We also note that [4] in fact construct an explicit signaling scheme that achieves this result; since then, alternative constructions have been found [14, 35] which also maximize $\mathcal{CS}(\mathcal{Z})$, with other additional desirable properties; moreover, any convex mixture of such schemes leads to new signaling schemes which all maximize $\mathcal{CS}(\mathcal{Z})$.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Example. Continuing our running example from Section 2.1, note that under the signaling scheme $\mathcal{Z}_1^{\mathsf{Ex}}$, the item is always sold, meaning that $\mathcal{Z}_1^{\mathsf{Ex}}$ is efficient. Moreover, since $\mathcal{R}(\mathcal{Z}_1^{\mathsf{Ex}}) = \mathcal{R}^{\mathsf{My}}$, we have that $\mathcal{Z}_1^{\mathsf{Ex}}$ maximizes $\mathcal{CS}(\mathcal{Z})$, and the maximum surplus consumer is 1.

Fig. 2 illustrates an alternative efficient signaling scheme $\mathcal{Z}_2^{\mathsf{Ex}}$ (based on the construction of [35]), which we call "remove from bottom". Again, we can check that $\mathcal{CS}(\mathcal{Z}_2^{\mathsf{Ex}}) = 1$ and the item always sells, so that this scheme is buyer-optimal. ∎

Fairness Metrics (Monotonicity and Equitable Welfare Functions). The main problem with focusing on utilitarian metrics alone is that they do not give good guarantees for each individual agent's surplus. To understand how *fair* a given signaling scheme is, we need to consider additional performance metrics. The simplest of these is monotonicity: we say a signaling scheme $\mathcal{Z}$ is *monotone* if buyers with larger values gain larger expected surplus from $\mathcal{Z}$:

Definition 2.4. (Monotonicity) *A signaling scheme $\mathcal{Z}$ is* monotone *if for any ordered pair of values $v_i < v_j$, we have $\mathcal{CS}_{v_i}(\mathcal{Z}) \le \mathcal{CS}_{v_j}(\mathcal{Z})$.*

Our running example shows that not all schemes satisfying Proposition 2.1 (i.e., efficient and consumer surplus maximizing) are monotone.

Example. In Section 2.2 under $\mathcal{Z}_2^{\mathsf{Ex}}$, the expected consumer surplus of buyers with values $\langle 1, 2, 5, 6\rangle$ are respectively $\langle 0, \frac{1}{7}, \frac{10}{7}, \frac{17}{7}\rangle$. This vector is monotone; however, $\mathcal{Z}_1^{\mathsf{Ex}}$ in Section 2.1 has surplus vector $\langle 0, 0.6, 0.4, 3\rangle$, which is not monotone since $\mathcal{CS}_2(\mathcal{Z}_1^{\mathsf{Ex}}) = 0.6 > 0.4 = \mathcal{CS}_5(\mathcal{Z}_1^{\mathsf{Ex}})$. ∎

What can we say about what a *fair* signaling scheme is, beyond the above metrics (efficiency, maximizing consumer surplus, monotonicity)? One option that is often used is to maximize an alternative *equitable* welfare function – one which promotes a more balanced solution. Such a welfare function $\mathcal{W}$ takes as input the surplus vector $\vec{u}$ containing the expected surplus under each value, and outputs a real number; moreover, $\mathcal{W}$ satisfies the following natural properties:

- (Symmetry) For any $\vec{u}$ and any permutation σ, $\mathcal{W}(\sigma(\vec{u})) = \mathcal{W}(\vec{u})$. In other words, it treats the buyers equally.
- (Non-decreasing) For any $\vec{u_1} \le \vec{u_2}$, $\mathcal{W}(\vec{u_1}) \le \mathcal{W}(\vec{u_2})$. In other words, it weakly prefers Pareto improvements.
- (Concavity) $\mathcal{W}$ is concave. In other words, it weakly prefers a balanced allocation to a convex combination of extremes with the same expected allocation.
- (Normalization) $\mathcal{W}(\vec{0}) = 0$. (It also suffices to alternatively require non-negativity: $\mathcal{W}(\vec{u}) \ge 0$ for any $\vec{u} \ge 0$.)

This definition captures many common welfare functions, such as the utilitarian social welfare function that outputs the sum, the Nash welfare function that outputs the geometric mean, and the max-min (a.k.a. egalitarian) welfare function that outputs the minimum. It will be clear that we cannot hope for similar results if we drop any of these four conditions.

We will show the surprising existence of a universal scheme – we do not need to know $\mathcal{W}$ in order to approximately optimize it. Our technical tool to deal with the unknown $\mathcal{W}$ is *majorization.*[5] Below we define it with the related notions which we will need later.

Majorization. Given a signaling scheme $\mathcal{Z}$, we define its *surplus-mass function* to be a step function over $(0, 1]$ taking value $\mathcal{CS}_{v_i}(\mathcal{Z})$ on the interval $(F_{\mathcal{D}}(v_{i-1}), F_{\mathcal{D}}(v_i)]$. Formally, we have:

Definition 2.5. (Surplus-Mass Function) *Given a signaling scheme $\mathcal{Z}$, the* surplus-mass function *induced from $\mathcal{Z}$ is a step function $s_{\mathcal{Z}} : (0, 1] \to \mathbb{R}_{\ge 0}$ that satisfies $\forall\, x \in (0, 1]$ and $i \in [n]$:*

$$s_{\mathcal{Z}}(x) = \mathcal{CS}_{v_i}(\mathcal{Z}),\ \forall x \in (F_{\mathcal{D}}(v_{i-1}), F_{\mathcal{D}}(v_i)]\,.$$

[5] More accurately, we use the notion of *majorization from above*, a.k.a. *supermajorization*. For simplicity, we use the term *majorization* throughout this paper.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

That is, the surplus-mass function maps a quantile in the value distribution to the expected surplus of the buyer with that value.

Definition 2.6. (Integration Prefix Sum) *Given a function $f : (0,1] \to \mathbb{R}_{\geq 0}$ and $m \in (0,1]$. The* m-integration prefix sum *of f is*

$$\textsc{IPrefix}(f, m) = \int_0^m f(x)\,\mathrm{d}x.$$

Next, we define the *sorted m-prefix sum* of any step function f as the area under the curve over the leftmost m-length interval of the "sorted function" obtained by sorting the segments of f.

Definition 2.7. (Sorted Prefix Sum) *Given a step function $f : (0,1] \to \mathbb{R}_{\geq 0}$ with finite steps and a real number $m \in (0,1]$, define a new* sorted *function $f_{\text{sorted}}(x)$ by rearranging the segments in f in the ascending order of $f(x)$ (while keeping the domain $(0,1]$ unchanged). The* sorted m-prefix sum *of f is*

$$\textsc{SPrefix}(f, m) = \int_0^m f_{\text{sorted}}(x)\,\mathrm{d}x.$$

In other words, the sorted m-prefix sum outputs the minimum possible (over S) integral of $f(x)$ on $x \in S$, where $S \subseteq [0,1]$ is a finite union of disjoint intervals with total length of m.

We now define the majorization relation between two signaling schemes as follows:

Definition 2.8. (Majorization Relation) *A signaling scheme $\mathcal{Z}_1$ is* majorized *by another signaling scheme $\mathcal{Z}_2$ if*

$$\forall\, m \in (0,1],\ \textsc{SPrefix}(s_{\mathcal{Z}_1}, m) \geq \textsc{SPrefix}(s_{\mathcal{Z}_2}, m),$$

where $s_{\mathcal{Z}_1}, s_{\mathcal{Z}_2}$ are the surplus-mass functions induced under schemes $\mathcal{Z}_1$ and $\mathcal{Z}_2$ respectively. A signaling scheme is said to be majorized *if it is majorized by every other signaling scheme.*

Example. In our running example, the expected consumer surplus under $\mathcal{Z}_1^{\mathsf{Ex}}$ (resp. $\mathcal{Z}_2^{\mathsf{Ex}}$) is $\langle 0, 0.6, 0.4, 3\rangle$ (resp. $\langle 0, \frac{1}{7}, \frac{10}{7}, \frac{17}{7}\rangle$). Each of these surplus values occupies mass of $1/4$. Thus,

$$\textsc{SPrefix}\Big(\mathcal{Z}_2^{\mathsf{Ex}}, \frac{1}{2}\Big) = \frac{\sum_{v\in\{1,2\}} \mathcal{CS}_v(\mathcal{Z}_2^{\mathsf{Ex}})}{4} = \frac{1}{28} < \frac{1}{10} = \frac{\sum_{v\in\{1,3\}} \mathcal{CS}_v(\mathcal{Z}_1^{\mathsf{Ex}})}{4} = \textsc{SPrefix}\big(\mathcal{Z}_1^{\mathsf{Ex}}, \frac{1}{2}\big);$$
$$\textsc{SPrefix}\Big(\mathcal{Z}_1^{\mathsf{Ex}}, \frac{3}{4}\Big) = \frac{\sum_{v\in\{1,2,3\}} \mathcal{CS}_v(\mathcal{Z}_1^{\mathsf{Ex}})}{4} = \frac{1}{4} < \frac{11}{28} = \frac{\sum_{v\in\{1,2,3\}} \mathcal{CS}_v(\mathcal{Z}_2^{\mathsf{Ex}})}{4} = \textsc{SPrefix}\Big(\mathcal{Z}_2^{\mathsf{Ex}}, \frac{3}{4}\Big).$$

Thus $\mathcal{Z}_2^{\mathsf{Ex}}$ is not majorized by $\mathcal{Z}_1^{\mathsf{Ex}}$ and $\mathcal{Z}_1^{\mathsf{Ex}}$ is not majorized by $\mathcal{Z}_2^{\mathsf{Ex}}$, and hence neither $\mathcal{Z}_1^{\mathsf{Ex}}$ nor $\mathcal{Z}_2^{\mathsf{Ex}}$ can be majorized by every other signaling scheme. ∎

Indeed, in Section 4, we show that (exact) majorization is unattainable – there are instances where no signaling scheme is majorized by every other signaling scheme. Given this, we define the following approximation version of majorization.

Definition 2.9. (α-Majorization) *A signaling scheme $\mathcal{Z}_1$ is* α-majorized *by another signaling scheme $\mathcal{Z}_2$ if $\forall\, m \in (0,1]$, we have:*

$$\alpha \cdot \textsc{SPrefix}(s_{\mathcal{Z}_1}, m) \geq \textsc{SPrefix}(s_{\mathcal{Z}_2}, m).$$

Further, we say a signaling scheme $\mathcal{Z}$ is α-majorized *if it is α-majorized by every other signaling scheme $\mathcal{Z}'$.*

The following established fact [27, 23] is crucial in our universal maximization of well-behaved welfare functions. We include a proof for completeness.

Proposition 2.2. (Proved in Appendix A) *Any α-majorized signaling scheme $\mathcal{Z}$ gives an α-approximation to the welfare under any signaling scheme, as long as the welfare function is symmetric, weakly increasing, concave, and normalized (or non-negative). Conversely, if a signaling scheme $\mathcal{Z}$ gives an α-approximation to all such welfare functions, it must be α-majorized.*

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

3 Finding an 8-Majorized Signaling Scheme

In this section, we construct an 8-majorized signaling scheme. In more detail, in Section 3.1, we present our Split-and-Match Algorithm (Algorithm 1) that given any empirical distribution $\mathcal{D}$ constructs a signaling scheme $\mathcal{Z}_0$ that approximates the m-integration prefix sum of any other signaling scheme:

LEMMA 3.1. *Given $\mathcal{D}$, let $\mathcal{Z}_0$ denote the signaling scheme returned by Split-and-Match Algorithm (Algorithm 1). Then, for any signaling scheme $\mathcal{Z}'$ and any $m \in (0,1]$, we have*

$$4 \cdot \text{IPREFIX}(s_{\mathcal{Z}_0}, m) \geq \text{IPREFIX}(s_{\mathcal{Z}'}, m).$$

However, in order to achieve α-majorization, we need to approximate the optimal *sorted* prefix sum (rather than the optimal integration prefix sum). In Section 3.3, we show an ironing process that transforms the surplus-mass function $s_{\mathcal{Z}_0}(x)$ into a monotonically increasing step function $\tilde{s}(x)$, while preserving the integration prefix sum of $s_{\mathcal{Z}_0}$ at any point of discontinuity. Based on $\mathcal{Z}_0$, we then construct a monotone signaling scheme $\mathcal{Z}_1$ such that the surplus-mass function of $\mathcal{Z}_1$ is exactly half of $\tilde{s}$:

LEMMA 3.2. $\forall x \in (0,1], s_{\mathcal{Z}_1}(x) = \frac{1}{2} \cdot \tilde{s}(x)$.

Combining Lemmas 3.1 and 3.2 leads to our main result:

THEOREM 3.1. *$\mathcal{Z}_1$ is efficient, monotone, and 8-majorized.*

3.1 Construction of $\mathcal{Z}_0$ In this section, we construct the signaling scheme $\mathcal{Z}_0$. The main idea is to decompose any given $\mathcal{D}$ into only two types of posterior distributions, which we refer to as *singleton* and *equal-revenue binary* signals, while achieving the property mentioned in Lemma 3.1. Such decomposition helps simplify the structure of the signaling scheme (since each signal has support size of at most 2), making it easier for us to conduct further manipulations on the signaling scheme.

DEFINITION 3.1. (SINGLETON SIGNAL) *A signal $\mathcal{S} = \mathcal{S}^{\mathbb{1}}_{v_i}$ is said to be a* singleton signal *on v_i if it satisfies $f_{\mathcal{S}}(v) = \mathbb{1}[v = v_i]$.*

DEFINITION 3.2. (EQUAL-REVENUE BINARY SIGNAL) *A signal $\mathcal{S} = \mathcal{S}^{E}_{v_i,v_j}$ is said to be an* equal-revenue binary signal *on $v_i < v_j$ if it satisfies:*

$$f_{\mathcal{S}}(v) = \begin{cases} 1 - \frac{v_i}{v_j} & v = v_i; \\ \frac{v_i}{v_j} & v = v_j; \\ 0 & v \notin \{v_i, v_j\}. \end{cases}$$

Note that if the seller receives $\mathcal{S}^{E}_{v_i,v_j}$, then posting a price of either v_i and v_j leads to the same revenue (hence "equal-revenue"). We assume that the equal-revenue binary signals are indexed from 1 to Q_1, where for each $q \in [Q_1]$, we have the signal $\mathcal{S}_q = \mathcal{S}^{\text{E}}_{v_q^1, v_q^2}$. We call the higher value v_q^2 *taker* and the lower value v_q^1 *giver*. The masses on them (i.e. $\gamma_q f_{\mathcal{S}_q}(v_q^2)$ and $\gamma_q f_{\mathcal{S}_q}(v_q^1)$ respectively) are called *taker mass* and *giver mass* respectively. Using these definitions, we can describe our Split-and-Match Algorithm in Algorithm 1.

To understand this construction, first note that by the Bayes plausibility of a signaling scheme, we have the following set of linear constraints on any equal-revenue binary signal:

$$\forall i \in [n], \sum_{q \in Q_1 : v_q^1 = v_i} \left(1 - \frac{v_i}{v_j}\right) \cdot \gamma_q + \sum_{q \in Q_1 : v_q^2 = v_i} \left(\frac{v_i}{v_j}\right) \cdot \gamma_q \leq f_{\mathcal{D}}(v_i). \tag{3.3}$$

In the construction of $\mathcal{Z}_0$, we strengthen these constraints into the following:

$$\forall i \in [n], \sum_{q \in Q_1 : v_q^1 = v_i} \left(1 - \frac{v_i}{v_j}\right) \cdot \gamma_q \leq \frac{1}{2} \cdot f_{\mathcal{D}}(v_i); \tag{3.4}$$

$$\forall i \in [n], \sum_{q \in Q_1 : v_q^2 = v_i} \left(\frac{v_i}{v_j}\right) \cdot \gamma_q \leq \frac{1}{2} \cdot f_{\mathcal{D}}(v_i). \tag{3.5}$$

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Algorithm 1 Split-and-Match Algorithm

Input: Distribution $\mathcal{D} = \{v_i, f_{\mathcal{D}}(v_i)\}_{i\in[n]}$.
Output: A signaling scheme $\mathcal{Z}_0 = \{\mathcal{S}_q, \gamma_q\}_{q\in[Q]}$.

Initialize: $\mathcal{Z}_0 \leftarrow \varnothing$; $q \leftarrow 0$; $(m_i^{\mathrm{G}}, m_i^{\mathrm{T}}) \leftarrow \left(\frac{1}{2}\cdot f_{\mathcal{D}}(v_i), \frac{1}{2}\cdot f_{\mathcal{D}}(v_i)\right)\ \forall\, i \in [n]$;
repeat
  $q \leftarrow q + 1$;
  Find the smallest $s \in [n]$ such that $m_s^{\mathrm{G}} > 0$;
  Find the smallest $\ell > s$ such that $m_\ell^{\mathrm{T}} > 0$;
  Set $\mathcal{S}_q = \mathcal{S}^{\mathrm{E}}_{v_s, v_\ell}$ and $\gamma_q = \min\{m_s^{\mathrm{G}}/\left(1 - \frac{v_s}{v_\ell}\right), m_\ell^{\mathrm{T}}/\left(\frac{v_s}{v_\ell}\right)\}$;
  Add signal $(\mathcal{S}_q, \gamma_q)$ to $\mathcal{Z}_0$;
  Update $m_s^{\mathrm{G}} \leftarrow m_s^{\mathrm{G}} - \gamma_q \cdot (1 - \frac{v_s}{v_\ell})$; $m_\ell^{\mathrm{T}} \leftarrow m_\ell^{\mathrm{T}} - \gamma_q \cdot \frac{v_s}{v_\ell}$;
until no such (s, l) exists;
Cover all remaining masses using singleton signals and add them to $\mathcal{Z}_0$.

We conduct a greedy process to find a solution satisfying the strengthened constraints. We iteratively find the smallest index s such that Eq. (3.4) is slack for s and the smallest index ℓ ($\ell > s$) such that Eq. (3.5) is slack for ℓ. We add a maximal equal-revenue binary signal with supports v_s and v_ℓ so that one of the two constraints becomes tight. We iterate until no such pair of (i_s, i_ℓ) exists. Fig. 3 illustrates the construction of our signaling scheme $\mathcal{Z}_0$ obtained by running Algorithm 1 on some given $\mathcal{D}$. The intuition of this greedy-style matching comes from the following observation: Consider an equal-revenue binary signal $\mathcal{S}$ with giver a and taker b. The probability on b is thus $f_{\mathcal{S}}(b) = \frac{a}{b}$, and we also have $\frac{\mathcal{CS}_b(\mathcal{S})}{f_{\mathcal{S}}(b)} = b - a$. For a fixed b, on the one hand, we want to make a as small as possible to have the most surplus per unit mass usage of b; on the other hand, setting a too small will lead to a large $\frac{b}{a}$ and hence will exhaust the mass of a too quickly. We want to balance this trade-off between gaining surplus and using low-value giver masses. This leads us to the idea of first splitting each probability mass into two halves and then choosing the smallest a and the smallest b to match.

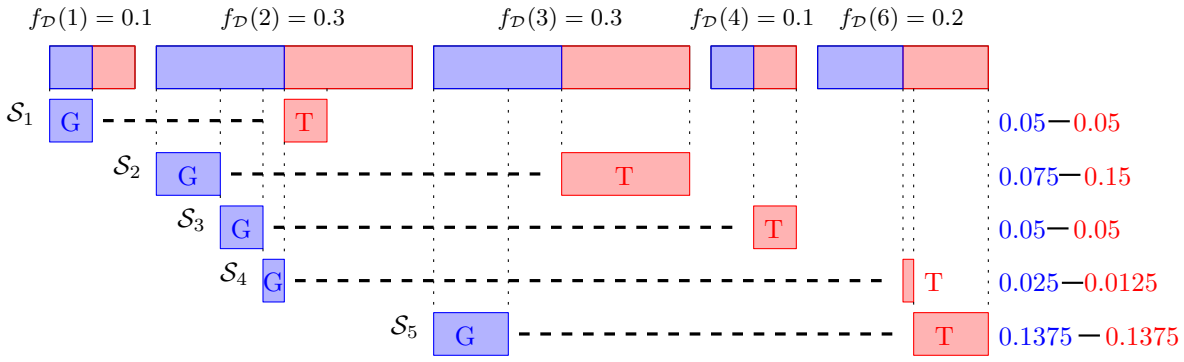

Figure 3: *Illustrating the construction of $\mathcal{Z}_0$ using Algorithm 1: The buyer values are $\langle 1, 2, 3, 4, 6\rangle$ with distribution $\mathcal{D} = \langle 0.1, 0.3, 0.3, 0.1, 0.2\rangle$, indicated using scaled rectangles in the first row. Each subsequent row corresponds to an equal-revenue binary signal $\mathcal{S}_q = \mathcal{S}^E_{v_i, v_j}$ with weight γ_q. The letters* G *and* T *together with the blue and red rectangles represent the giver and the taker corresponding to each signal, while the blue and red numbers on the far right of each row are the giver and taker masses $(\gamma_q f_{\mathcal{S}_q}(v_i), \gamma_q f_{\mathcal{S}_q}(v_j))$; these are also illustrated by the lengths of the rectangles in that row. (Again, we can recover each signal and its weight by normalizing the numbers. For example, signal $\mathcal{S}_1$ satisfies $f_{\mathcal{S}_1}(1) = f_{\mathcal{S}_1}(2) = 0.5$ and has weight $\gamma_1 = 0.1$.)*

3.2 Approximating IPREFIX via Split-and-Match (Proof of Lemma 3.1) Recall for any k, $F_{\mathcal{D}}(v_k) = \sum_{i=1}^{k} f_{\mathcal{D}}(v_i)$ denotes the total population of buyers with value at most v_k. We now prove Lemma 3.1 in two steps: First, in Section 3.2.1, we show an upper bound on $\textsc{IPrefix}\left(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k)\right)$ for any signal $\mathcal{Z}'$. This generalizes the corresponding bound of [4] to a sub-population. Next, in Section 3.2.2, we show that the IPREFIX values of $\mathcal{Z}_0$ approximately achieve this upper bound.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

3.2.1 Bounding IPrefix via the Surplus of Truncated Distributions As a thought experiment, we restrict our attention to the subset of buyers with the smallest k values. What is the maximum possible consumer surplus on this sub-population? In Lemma 3.3, we show that it is upper bounded by the total values in the sub-population, minus the revenue extractable from this sub-population without signaling. We need the following definitions to present the proof.

Definition 3.3. (Truncated Distribution) *Given distribution $\mathcal{D}$ with finite support $S \subset \mathbb{R}_{\geq 0}$ and any $x \in S$, the truncated distribution of $\mathcal{D}$ on x, denoted by $\overline{\mathcal{D}}(x)$, satisfies:*

$$f_{\overline{\mathcal{D}}(x)}(v) = \begin{cases} \frac{f_{\mathcal{D}}(v)}{F_{\mathcal{D}}(x)} & v \in S \text{ and } v \leq x; \\ 0 & \text{Otherwise.} \end{cases}$$

Definition 3.4. (Surplus Prefix Sum) *Given any $k \in [n]$, the k-*surplus prefix sum *of the buyers is*

$$V_k = \sum_{i=1}^{k} v_i \cdot f_{\mathcal{D}}(v_i).$$

Lemma 3.3. *For any $k \in [n]$ and any signaling scheme $\mathcal{Z}'$, we have*

$$\textsc{IPrefix}\big(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k)\big) \leq V_k - \max_{i \in [k]} \left\{ v_i \cdot \sum_{j=i}^{k} f_{\mathcal{D}}(v_j) \right\}.$$

Proof. Suppose that $\mathcal{Z}' = \{(\mathcal{S}_q, \gamma_q)\}_{q \in [Q]}$. We have

$$\textsc{IPrefix}\big(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k)\big) = \sum_{i=1}^{k} f_{\mathcal{D}}(v_i) \cdot \left(\sum_{q=1}^{Q} \frac{\gamma_q \cdot f_{\mathcal{S}_q}(v_i)}{f_{\mathcal{D}}(v_i)} \cdot \mathbb{1}[v_i \geq p^*_{\mathcal{S}_q}] \cdot (v_i - p^*_{\mathcal{S}_q}) \right).$$

Let $\hat{\mathcal{S}}_q = \overline{\mathcal{S}_q}(v_k)$ be the truncated distribution on the buyers with value at most v_k. Denote the optimal price of $\hat{\mathcal{S}}_q$ by $p^*_{\hat{\mathcal{S}}_q}$. We now claim that $p^*_{\hat{\mathcal{S}}_q} \leq p^*_{\mathcal{S}_q}$. First assume $p^*_{\mathcal{S}_q} \leq v_k$, otherwise the statement is trivial. Next, to find the optimal price, we can ignore the scaling factor $1/F_{\mathcal{D}}(v_k)$ (since it scales up the revenue for each price by the same amount). Thus, when we truncate at v_k, we can view it as removing some probability mass μ_k beyond v_k. Note that for any price p, the decrease in revenue is $p \cdot \mu_k$. This means larger prices suffer larger drops in revenue, i.e., the new optimal price $p^*_{\hat{\mathcal{S}}_q}$ cannot be larger.

Moreover, since $(v_i - p)$ is monotonically decreasing as a function of p, substituting $p^*_{\hat{\mathcal{S}}_q}$ for $p^*_{\mathcal{S}_q}$ in the above equality, we have:

$$\begin{aligned} \textsc{IPrefix}\big(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k)\big) &\leq \sum_{q=1}^{Q} \left(\sum_{i \in [k]: v_i \geq p^*_{\hat{\mathcal{S}}_q}} \gamma_q \cdot f_{\mathcal{S}_q}(v_i) \cdot (v_i - p^*_{\hat{\mathcal{S}}_q}) \right) \\ &\leq \sum_{q=1}^{Q} \left(\sum_{i \in [k]} \gamma_q \cdot f_{\mathcal{S}_q}(v_i) \cdot v_i \right) - \sum_{q=1}^{Q} \gamma_q \cdot \left(\sum_{i \in [k]: v_i \geq p^*_{\hat{\mathcal{S}}_q}} f_{\mathcal{S}_q}(v_i) \cdot p^*_{\hat{\mathcal{S}}_q} \right). \end{aligned}$$

Let $i^*_k = \arg\max_{i \in [k]} \left\{ v_i \cdot \sum_{i'=i}^{k} f_{\mathcal{D}}(v_{i'}) \right\}$. Since $p^*_{\hat{\mathcal{S}}_q}$ is the optimal price on $\hat{S}_q$, we have

$$p_{\hat{\mathcal{S}}^*_q} \sum_{i \in [k]: v_i \geq p^*_{\hat{\mathcal{S}}_q}} f_{\mathcal{S}_q}(v_i) \geq v_{i^*_k} \sum_{i \in [k]: v_i \geq v_{i^*_k}} f_{\mathcal{S}_q}(v_i)$$

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

for any $q \in [Q]$. Plugging this into the previous inequality, we finally have:

$$\begin{aligned}\textsc{IPrefix}\big(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k)\big) &\le \sum_{i=1}^{k}\left(v_i \cdot \sum_{q=1}^{Q}\gamma_q \cdot f_{\mathcal{S}_q}(v_i)\right) - \sum_{q=1}^{Q}\gamma_q \cdot \left(\sum_{i\in[k]:v_i\ge v_{i_k^*}} f_{\mathcal{S}_q}(v_i)\cdot v_{i_k^*}\right)\\ &= \sum_{i=1}^{k} v_i \cdot f_{\mathcal{D}}(v_i) - v_{i_k^*}\cdot\sum_{i=i_k^*}^{k} f_{\mathcal{D}}(v_i)\\ &= V_k - \max_{i\in[k]}\left\{v_i\cdot\sum_{i'=i}^{k} f_{\mathcal{D}}(v_{i'})\right\},\end{aligned}$$

completing the proof. ∎

3.2.2 Approximating Prefix Sums via the Split-and-Match Algorithm By Definition 2.5, for any signaling scheme $\mathcal{Z}$ and any $k \in [n]$, $s_{\mathcal{Z}}$ is constant on the interval $(F_{\mathcal{D}}(v_{k-1}), F_{\mathcal{D}}(v_k)]$. Therefore, $\textsc{IPrefix}(\mathcal{Z}, m)$ is a linear function of m on the interval $(F_{\mathcal{D}}(v_{k-1}), F_{\mathcal{D}}(v_k)]$. Therefore, to prove Lemma 3.1, it suffices to show it when $m = F_{\mathcal{D}}(v_k)$ for $k \in [n]$. Moreover, we can further replace $\textsc{IPrefix}(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k))$ with the upper bound we obtained Lemma 3.3 (i.e, with the maximum consumer surplus of truncated distributions). Thus we can obtain Lemma 3.1 as an immediate consequence of the following lemma.

Lemma 3.4. *Let $\mathcal{Z}_0$ be the signaling scheme returned by Algorithm 1 for a given $\mathcal{D}$. Then for any signaling scheme $\mathcal{Z}'$ and any $k \in [n]$, we have*

$$4\cdot\textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) \ge V_k - \max_{i\in[k]}\left\{v_i\cdot\sum_{i'=i}^{k} f_{\mathcal{D}}(v_{i'})\right\}.$$

Consider the first point in time in Algorithm 1 when $\ell = k+1$; if $k = n$, this is the stopping time of the algorithm. Let i^* be the smallest index such that Eq. (3.4) is still slack. This means Eq. (3.4) is tight for all $i \in [i^*-1]$ and Eq. (3.5) is tight for all $i \in [i^*+1, k]$ at this point in time. Note that $i^* = \min\{i \mid m_i^{\mathrm{G}} > 0\}$.

We now establish two lower bounds for $\textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k))$, based on the definition of V_k in Definition 3.4.

Proposition 3.1. *We have the following two inequalities.*

- $2\cdot\textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) \ge V_{i^*-1}$.
- $2\cdot\textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) \ge V_k - V_{i^*-1} - \max_{i\in[k]}\left\{v_i\cdot\sum_{i'=i}^{k} f_{\mathcal{D}}(v_{i'})\right\}$.

Proof. Assume that the j^{th} equal-revenue binary signal added to $\mathcal{Z}_0$ during Algorithm 1 is $\mathcal{S}_j = \mathcal{S}^{\mathrm{E}}_{v_j^{\mathrm{G}}, v_j^{\mathrm{T}}}$ with weight γ_j. Let $t_j^{\mathrm{G}} = f_{\mathcal{S}_j}(v_j^{\mathrm{G}})\cdot\gamma_j$ and $t_j^{\mathrm{T}} = f_{\mathcal{S}_j}(v_j^{\mathrm{T}})\cdot\gamma_j$. By Definition 3.2, we have

$$t_j^{\mathrm{T}}\cdot v_j^{\mathrm{T}} = (t_j^{\mathrm{G}} + t_j^{\mathrm{T}})\cdot v_j^{\mathrm{G}}. \tag{3.6}$$

Therefore,

$$\begin{aligned}&\text{(By definition of buyers' surplus)} \quad 2\cdot\textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) = 2\cdot\sum_{i=1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i}(v_j^{\mathrm{T}} - v_j^{\mathrm{G}})\cdot t_j^{\mathrm{T}}\right)\\ &\text{(By Eq. (3.6))} \quad = 2\cdot\sum_{i=1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i} v_j^{\mathrm{G}}\cdot t_j^{\mathrm{G}}\right).\end{aligned}$$

Since Eq. (3.4) is tight for any $i < i^*$, we have

$$\forall\, i \in [i^*-1], \quad \sum_{j:\, v_j^{\mathrm{G}}=v_i, v_j^{\mathrm{T}}\le v_k} t_j^{\mathrm{G}} = \frac{1}{2}\cdot f_{\mathcal{D}}(v_i).$$

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Since $\sum_{i=1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i} v_j^{\mathrm{G}}\cdot t_j^{\mathrm{G}}\right) \geq \sum_{i=1}^{i^*-1}\left(v_i\cdot\sum_{j:\, v_j^{\mathrm{G}}=v_i, v_j^{\mathrm{T}}\leq v_k} t_j^{\mathrm{G}}\right)$, we further have

$$2\cdot \textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) \geq \sum_{i=1}^{i^*-1} v_i\cdot f_{\mathcal{D}}(v_i) = V_{i^*-1}. \tag{3.7}$$

This completes the proof of the first inequality. To show the second inequality, we have

$$\begin{aligned}
2\cdot \textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) &\geq 2\cdot\sum_{i=i^*+1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i}(v_j^{\mathrm{T}}-v_j^{\mathrm{G}})\cdot t_j^{\mathrm{T}}\right)\\
\text{(Since all the giver values are at most } v_{i^*}) \qquad &\geq 2\cdot\sum_{i=i^*+1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i}(v_j^{\mathrm{T}}-v_{i^*})\cdot t_j^{\mathrm{T}}\right)\\
&= 2\cdot\sum_{i=i^*+1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i} v_j^{\mathrm{T}}\cdot t_j^{\mathrm{T}}\right) - 2\cdot v_{i^*}\cdot\sum_{i=i^*+1}^{k}\left(\sum_{j:\, v_j^{\mathrm{T}}=v_i} t_j^{\mathrm{T}}\right).
\end{aligned}$$

Since Eq. (3.5) is tight for any $i\in[i^*+1,\ k]$, this means

$$\forall i\in[i^*+1,\ k],\quad \sum_{j:\, v_j^{\mathrm{T}}=v_i} t_j^{\mathrm{T}} = \frac{1}{2}\cdot f_{\mathcal{D}}(v_i).$$

Using this in the above derivation, we have:

$$\begin{aligned}
2\cdot \textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) &\geq \sum_{i=i^*+1}^{k} f_{\mathcal{D}}(v_i)\cdot v_i - v_{i^*}\cdot\sum_{i=i^*+1}^{k} f_{\mathcal{D}}(v_i)\\
&= \sum_{i=i^*}^{k} f_{\mathcal{D}}(v_i)\cdot v_i - v_{i^*}\cdot\sum_{i=i^*}^{k} f_{\mathcal{D}}(v_i)\\
&\geq V_k - V_{i^*-1} - \max_{i\in[k]}\left\{v_i\cdot\sum_{i'=i}^{k} f_{\mathcal{D}}(v_{i'})\right\}.
\end{aligned}$$

This completes the proof of the second inequality. ∎

Proof. [Proof of Lemma 3.4] Adding the inequalities in the proposition above, we have

$$\begin{aligned}
4\cdot \textsc{IPrefix}(s_{\mathcal{Z}_0}, F_{\mathcal{D}}(v_k)) &\geq V_{i^*-1} + V_k - V_{i^*-1} - v_{i^*}\cdot\sum_{i=i^*}^{k} f_{\mathcal{D}}(v_i)\\
&\geq V_k - \max_{i\in[k]} v_i\cdot\sum_{i'=i}^{k} f_{\mathcal{D}}(v_{i'})\\
\text{(By Lemma 3.3)} \qquad &\geq \textsc{IPrefix}(s_{\mathcal{Z}'}, F_{\mathcal{D}}(v_k)),
\end{aligned}$$

completing the proof of Lemma 3.4 and hence that of Lemma 3.1. ∎

3.3 Extending to Sorted Prefix Sums via Ironing and Smoothing We will now prove Lemma 3.2. In Section 3.3.1, we introduce the (classical) ironing process that transforms the surplus-mass function $s_{\mathcal{Z}_0}(x)$ into $\tilde{s}(x)$. After that, in Section 3.3.2, we conduct a smoothing process to obtain a signaling scheme $\mathcal{Z}_0'$ whose induced surplus-mass function is at least half of $\tilde{s}(x)$ (see Lemma 3.7). The ironing process can be intuitively viewed as moving surplus from some high-surplus but lower-value buyers to some high-value but low-surplus buyers, so that the ironed function is monotone. It makes the surplus-mass function more "even". The smoothing process describes what specific modification we should operate on the signaling scheme (or signals) to achieve the ironing purpose on the surplus-mass function. This will show Lemma 3.2. Denote the final signaling scheme after decomposition by $\mathcal{Z}_1$. In Section 3.3.4, we show that $\mathcal{Z}_1$ simultaneously guarantees 8-majorization and monotonicity, thus completing the proof of Theorem 3.1.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

3.3.1 Ironing So far, we have approximated the optimal integration prefix sum of consumer surplus. We need to transform the approximation on integration prefix sum into the approximation on *sorted* prefix sum, which will yield the bound on approximate majorization. Our first step is to process the surplus-mass function via ironing. Ironing is a standard process on functions to achieve monotonicity. It is first applied in the auction scenario by [41]. For the completeness of our paper, we also include a description of the ironing process in this section.

Consider the surplus-mass function $s_{\mathcal{Z}_0}(x)$. We operate the ironing process as follows:

1. Compute the integral of $s_{\mathcal{Z}_0}(x)$ as $F(x)$.
2. Compute the lower convex envelope of $F(x)$, denoted as $\tilde{F}(x)$.
3. Compute the derivative of $\tilde{F}(x)$ as the ironed function $\tilde{s}(x)$. Define the value at any point of discontinuity as its left limit.

We have the following properties of $\tilde{s}(x)$.

LEMMA 3.5. *The ironed function $\tilde{s}(x)$ satisfies:*

1. *$\tilde{s}(x)$ is weakly increasing;*
2. *For any $m \in (0, 1]$, we have $\text{IPREFIX}(\tilde{s}, m) \leq \text{IPREFIX}(s_{\mathcal{Z}_0}, m)$.*

Proof. Since $\tilde{F}$ is convex, we have its derivative $\tilde{s}$ is weakly increasing. Moreover, since the convex envelope has property that $\tilde{F}(x) \leq F(x)$ and the IPREFIX is defined by the integration from 0, we have $\forall\, m \in (0, 1]$, $\text{IPREFIX}(\tilde{s}, m) \leq \text{IPREFIX}(s_{\mathcal{Z}_0}, m)$. ∎

Since both $\tilde{s}(x)$ and $s_{\mathcal{Z}_0}(x)$ are step functions, the range of x where $F(x)$ and $\tilde{F}(x)$ are different consists of a collection of open intervals. On the graph depicting $F(x)$ and $\tilde{F}(x)$, each interval represents a region of x such that $\tilde{F}(x)$ falls below $F(x)$.

We call these open intervals "ironing intervals" and denote them by $I_1, I_2, \ldots, I_T$. Within each interval I_t, $\tilde{F}(x)$ is a linear function, and thus $\tilde{s}(x)$ is a constant. We denote this constant by $\tilde{s}_t$.

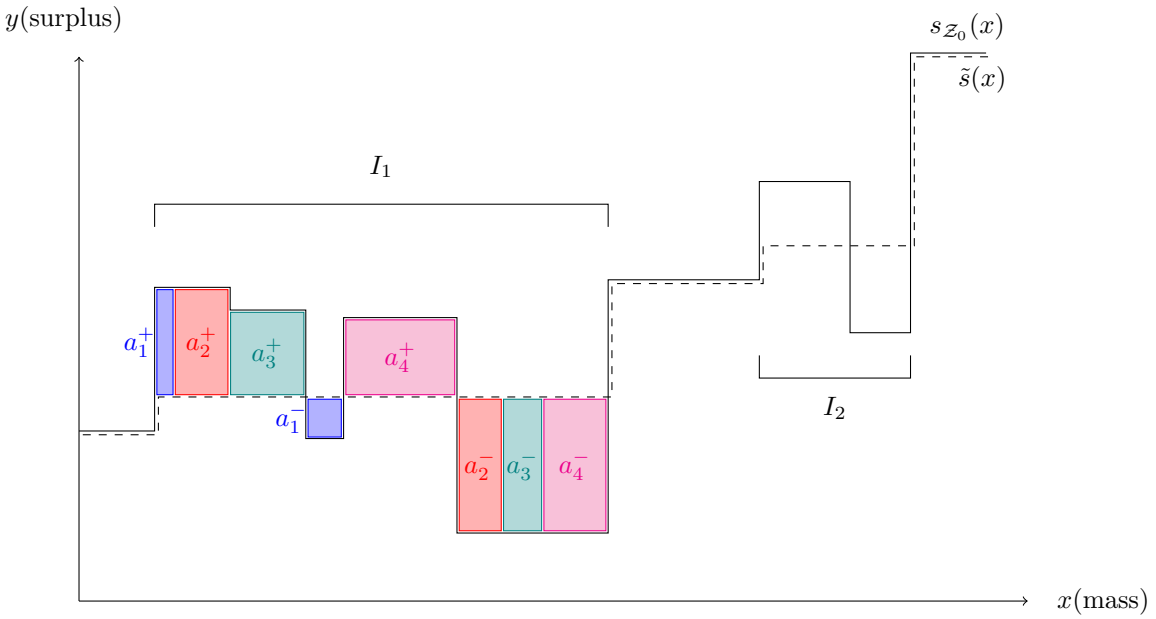

Figure 4: The solid line denotes the function $s_{\mathcal{Z}_0}(x)$. After ironing, the dashed line denotes the function $\tilde{s}(x)$. The masses are sorted in ascending order of the values. The four pairs of rectangles with the same colour share the same area respectively. The first two ironing intervals are shown as I_1 and I_2.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Fix an ironing interval $I_t = (\ell_t, r_t)$. By the definition of ironing intervals, we have that for any $x_0 \in \{\ell_t, r_t\}$, $F(x_0) = \int_0^{x_0} s_{\mathcal{Z}_0}(x)\,\mathrm{d}x = \int_0^{x_0} \tilde{s}(x)\,\mathrm{d}x = \tilde{F}(x_0)$, thus $\int_{\ell_t}^{r_t} s_{\mathcal{Z}_0}(x)\,\mathrm{d}x = \int_{\ell_t}^{r_t} \tilde{s}(x)\,\mathrm{d}x$. Equivalently, we have

$$\int_{x\in I_t:\tilde{s}(x)>s_{\mathcal{Z}_0}(x)} (\tilde{s}(x) - s_{\mathcal{Z}_0}(x))\,\mathrm{d}x = \int_{x\in I_t:\tilde{s}(x)\le s_{\mathcal{Z}_0}(x)} (s_{\mathcal{Z}_0}(x) - \tilde{s}(x))\,\mathrm{d}x.$$

This means $\int_{x\in I_t}[s_{\mathcal{Z}_0}(x) - \tilde{s}_t]^+\,\mathrm{d}x = \int_{x\in I_t}[\tilde{s}_t - s_{\mathcal{Z}_0}(x)]^+\,\mathrm{d}x$. Moreover, by the second property in Lemma 3.5, we have

$$\forall x_0 \in (\ell_t, r_t],\ \int_{\ell_t}^{x_0} [s_{\mathcal{Z}_0}(x) - \tilde{s}_t]^+\,\mathrm{d}x \ge \int_{\ell_t}^{x_0} [\tilde{s}_t - s_{\mathcal{Z}_0}(x)]^+\,\mathrm{d}x.$$

Based on these two observations, we can split the area above $y = \tilde{s}_t$ while below $s_{\mathcal{Z}_0}(x)$ into Y_t rectangles $\{a_1^+, a_2^+, \ldots, a_{Y_t}^+\}$, as well as the area below $y = \tilde{s}_t$ while above $s_{\mathcal{Z}_0}(x)$ into the same number of rectangles $\{a_1^-, a_2^-, \ldots, a_{Y_t}^-\}$. Algorithm 2 describes the process of constructing such pairs of rectangles. For any $y \in [Y_t]$, the pair of rectangles (a_y^+, a_y^-) satisfies the following conditions:

1. They have the same area;
2. a_y^+ is on the left of a_y^-;
3. Each rectangle a_y^+ (resp. a_y^-) corresponds to a single buyer value v_y^+ (resp. v_y^-).

Figure 4 shows an example of ironing and pairing of rectangles within the interval I_1. The four rectangles above $y = \tilde{s}_t$ (i.e. $\{a_1^+, a_2^+, a_3^+, a_4^+\}$) have the same areas as the four rectangles below $y = \tilde{s}_t$ (i.e. $\{a_1^-, a_2^-, a_3^-, a_4^-\}$) respectively.

Algorithm 2 Construct a rectangle pairing on I_t

Input: $s_{\mathcal{Z}_0}(x)$, $I_t = (\ell_t, r_t]$, $\tilde{s}_t$ (ironed surplus on I_t).
Output: A set of rectangle pairs $\{(a_y^+, a_y^-)\}_{y\in[Y_t]}$.

Let $v_{i_1^+} < \cdots < v_{i_g^+}$ be the buyer values whose occupied interval in I_t satisfies $s_{\mathcal{Z}_0}(x) > \tilde{s}_t$;
Let $v_{i_1^-} < \cdots < v_{i_h^-}$ be the buyer values whose occupied interval in I_t satisfies $s_{\mathcal{Z}_0}(x) < \tilde{s}_t$;
$x^+ \leftarrow F_{\mathcal{D}}(v_{i_1^+ - 1})$; $x^- \leftarrow F_{\mathcal{D}}(v_{i_1^- - 1})$;
$y \leftarrow 0$; $R \leftarrow \varnothing$;
repeat
  Find the smallest $\ell \in [g]$ such that $F_{\mathcal{D}}(v_{i_\ell^+}) > x^+$; $x^+ = \max\{x^+, F_{\mathcal{D}}(v_{i_\ell^+ - 1})\}$;
  Find the smallest $r \in [h]$ such that $F_{\mathcal{D}}(v_{i_r^-}) > x^-$; $x^- = \max\{x^-, F_{\mathcal{D}}(v_{i_r^- - 1})\}$;
  $A \leftarrow \min\left\{(F_{\mathcal{D}}(v_{i_\ell^+}) - x^+)\cdot(s_{\mathcal{Z}_0}(v_{i_\ell^+}) - \tilde{s}_t),\ (F_{\mathcal{D}}(v_{i_r^-}) - x^-)\cdot(\tilde{s}_t - s_{\mathcal{Z}_0}(v_{i_r^-}))\right\}$;
  $y \leftarrow y + 1$;
  Let a_y^+ be the rectangle with bottom-left corner coordinate at $(x^+, \tilde{s}_t)$ with width $w_y^+ = \frac{A}{v_{i_\ell^+} - \tilde{s}_t}$ and height $h_y^+ = v_{i_\ell^+} - \tilde{s}_t$;
  Let a_y^- be the rectangle with top-left corner coordinate at $(x^-, \tilde{s}_t)$ with width $w_y^- = \frac{A}{\tilde{s}_t - v_{i_r^-}}$ and height $h_y^- = \tilde{s}_t - v_{i_r^-}$;
  Add the rectangle pair (a_y^+, a_y^-) to R;
  $x^+ \leftarrow x^+ + w_y^+$; $x^- \leftarrow x^- + w_y^-$;
until no such (ℓ, r) exists.

3.3.2 Smoothing We now present the smoothing process. Consider the ironed interval I_t. Our goal is to make all buyers with expected consumer surplus less than $\tilde{s}_t/2$ in $\mathcal{Z}_0$ (call them *poor* buyers) have expected consumer surplus at least $\tilde{s}_t/2$ after smoothing. We do so by collecting a portion of giver masses from all the equal-revenue binary signals that contribute to the surplus of carefully chosen high-surplus (or *rich*) buyers. We use the collected masses as giver masses of these signals to construct new equal-revenue binary signals with the poor buyers' values as the taker value, hence bringing their expected surplus to at least $\tilde{s}_t/2$.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

We now describe the process in more detail. By applying Algorithm 2, we obtain a set of pairs of rectangles $\{(a_y^+, a_y^-)\}_{y\in[Y_t]}$. Suppose that the x-coordinates of a_y^+ and a_y^- correspond to buyer values v_y^+ and v_y^- respectively. Denote the widths of a_y^+ and a_y^- by w_y^+ and w_y^-, and the heights by h_y^+ and h_y^- respectively. Note that for a_y^-, the expected consumer surplus is $\tilde{s}_t - h_y^-$, while for a_y^+ it is $\tilde{s}_t + h_y^+$.

For each pair of rectangles (a_y^+, a_y^-), if $h_y^- > \tilde{s}_t/2$ (i.e. a buyer with value v_y^- is poor), we apply the following three steps on $\mathcal{Z}_0$:

1. Remove $\frac{w_y^-}{f_{\mathcal{D}}(v_y^-)}$ fraction of the weight from all the singleton signals on v_y^- and all the equal-revenue binary signals where v_y^- is the taker. For this removed weight, collect their taker masses, and discard their giver masses. (Intuitively, we collect the rectangle a_y^-.)

2. Remove $\frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+}$ fraction of weight from all equal-revenue binary signals where v_y^+ is the taker. For the removed weight, collect their giver masses and discard their taker masses. (Intuitively, we collect the givers to the rectangle a_y^+.)

3. Build equal-revenue binary signals using the masses collected in Step 2 as giver masses and the masses collected in Step 1 as taker masses.

Algorithm 3 Smoothing Algorithm

Input: $\mathcal{Z}_0 = \{(\mathcal{S}_q, \gamma_q)\}_{q\in[Q]}$, ironed surplus-mass function $\tilde{s}(x)$, ironing intervals $\{I_t\}_{t\in[T]}$.
Output: Smoothed signaling scheme $\mathcal{Z}_0' = \{(\mathcal{S}_{q'}', \gamma_{q'}')\}_{q'\in[Q']}$.

$\mathcal{Z}_0' \leftarrow \mathcal{Z}_0$; $q' \leftarrow Q$; ▷ Initialize $\mathcal{Z}_0'$ to $\mathcal{Z}_0$. Denote current $\mathcal{Z}_0'$ by $\{\mathcal{S}_q', \gamma_q'\}_{q\in[Q]}$;
for $t = 1$ to T **do**
  Apply Algorithm 2 to find rectangle pairs $R_t = \{(a_y^+, a_y^-)\}_{y\in[Y_t]}$ on ironing interval I_t;
  Width and height of a_y^+ (resp. a_y^-) are w_y^+ and h_y^+ (resp. w_y^- and h_y^-); a_y^+ (resp. a_y^-) are occupied by buyers with value v_y^+ (resp. v_y^-);
  for each rectangle pair (a_y^+, a_y^-) in R_t with $h_y^- > \tilde{s}_t/2$ **do**
    for each $\mathcal{S}_q$ in $\mathcal{Z}_0$ that is an equal-revenue binary signal with taker v_y^-, or a singleton signal on v_y^- **do**
      $\gamma_q' \leftarrow \gamma_q' - \gamma_q \cdot \frac{w_y^-}{f_{\mathcal{D}}(v_y^-)}$; ▷ Step 1: Collect $\gamma_q \cdot \frac{w_y^-}{f_{\mathcal{D}}(v_y^-)} \cdot f_{\mathcal{S}_q}(v_y^-)$ taker mass from value v_y^-
    end for
    for each $\mathcal{S}_q$ in $\mathcal{Z}_0$ that is an equal-revenue binary signal with taker v_y^+ and some giver v' **do**
      $\gamma_q' \leftarrow \gamma_q' - \gamma_q \cdot \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+}$;
      ▷ Step 2: Collect $\gamma_q \cdot \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot f_{\mathcal{S}_q}(v')$ giver mass from value v'
      $q' \leftarrow q' + 1$;
      $\gamma_{q'}' = \gamma_q \cdot \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot \big(1 - \frac{v'}{v_y^+}\big)/\big(1 - \frac{v'}{v_y^-}\big)$;
      Add equal-revenue binary signal $\mathcal{S}_{q'}' = \mathcal{S}^{\mathrm{E}}_{v', v_y^-}$ with weight $\gamma_{q'}'$ to $\mathcal{Z}_0'$;
      ▷ Step 3: Build equal-revenue binary signal on (v', v_y^-) with collected mass
    end for
  end for
end for
Make all the remaining masses singleton signals and add them to $\mathcal{Z}_0'$.

We formally present the smoothing process in Algorithm 3. Denote its output by $\mathcal{Z}_0' = \{(\mathcal{S}_{q'}', \gamma_{q'}')\}_{q'\in[Q']}$.

3.3.3 Analysis: Proof of Lemma 3.2 We show that $\mathcal{Z}_0'$ is feasible in Lemma 3.6. In other words, we prove that the outcome of the algorithm satisfies Bayes plausibility as in Eq. (3.3).

LEMMA 3.6. $\forall i \in [n]$, $\sum_{q'\in[Q']} f_{\mathcal{S}_{q'}'}(v_i) \cdot \gamma_{q'}' = f_{\mathcal{D}}(v_i)$.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Proof. Since Algorithm 3 makes all the remaining masses into singleton signals and adds them to $\mathcal{Z}_0'$ at the end, we simply need to show that the sum of masses in the equal-revenue binary signals does not exceed the mass of the prior for each value. Since we only add new signals in Line 13, it suffices to argue that there is always enough mass on v' and v_y^- to build equal-revenue binary signals.

We first argue that the total mass at v' is preserved. Consider a single run of Line 10 to Line 13 with fixed (t, y, q). In Line 10, the mass on v' is reduced by

$$f_{\mathcal{S}_q}(v') \cdot \gamma_q \cdot \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} = \left(1 - \frac{v'}{v_y^+}\right) \cdot \gamma_q \cdot \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+}.$$

This is exactly the mass added to v' on the newly constructed signal (i.e. $\gamma'_{q'} \cdot \left(1 - \frac{v'}{v_y^-}\right)$), thus the mass on v' is preserved.

We next argue the mass collected from v_y^- is enough for the new equal-revenue binary signals. Consider a single run of Line 6 to Line 14 with fixed (t, y). In $\mathcal{Z}_0$, at least $\frac{1}{2} \cdot f_{\mathcal{D}}(v_y^-)$ mass from value v_y^- is devoted to singletons signals on v_y^- or binary signals where v_y^- is a taker. Therefore, the mass collected from v_y^- in Line 7 is at least

$$\frac{1}{2} \cdot f_{\mathcal{D}}(v_y^-) \cdot \frac{w_y^-}{f_{\mathcal{D}}(v_y^-)} = w_y^-/2.$$

Let $r(t, y) = \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+}$. Consider the new equal-revenue signals added in Line 13 with v_y^- as taker. We assume that the j^{th} equal-revenue binary signal added to $\mathcal{Z}_0$ during Algorithm 1 is $\mathcal{S}_j = \mathcal{S}^{\text{E}}_{v_j^{\text{G}}, v_j^{\text{T}}}$ with weight γ_j. Let $t_j^{\text{G}} = f_{\mathcal{S}_j}(v_j^{\text{G}}) \cdot \gamma_j$ and $t_j^{\text{T}} = f_{\mathcal{S}_j}(v_j^{\text{T}}) \cdot \gamma_j$. Denote the sum of their taker masses by m_y^-. We have:

$$m_y^- = \sum_{j:\, v_j^{\text{T}} = v_y^+} \gamma_j \cdot r(t, y) \cdot \frac{v_j^{\text{G}}}{v_y^-} \cdot \left(1 - \frac{v_j^{\text{G}}}{v_j^{\text{T}}}\right) \Big/ \left(1 - \frac{v_j^{\text{G}}}{v_y^-}\right)$$

(Since $\mathcal{S}_j$ is an equal-revenue binary signal, $f_{\mathcal{S}_j}(v_j^{\text{T}}) = \frac{v_j^{\text{G}}}{v_j^{\text{T}}}$) $= \sum_{j:\, v_j^{\text{T}} = v_y^+} r(t, y) \cdot t_j^{\text{T}} \cdot \frac{v_j^{\text{T}} - v_j^{\text{G}}}{v_y^- - v_j^{\text{G}}}.$

Since $v_j^{\text{T}} = v_y^+ < v_y^-$ and since at most $\frac{1}{2} \cdot f_{\mathcal{D}}(v_y^+)$ mass from v_y^+ is used as taker mass in equal-revenue binary signals in $\mathcal{Z}_0$, the above simplifies to:

$$m_y^- \le \sum_{j:\, v_j^{\text{T}} = v_y^+} r(t, y) \cdot t_j^{\text{T}} \le \frac{f_{\mathcal{D}}(v_y^+)}{2} \cdot \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \le \frac{w_y^+}{2} \cdot \frac{h_y^+}{\tilde{s}_t}.$$

Since we have $w_y^+ \cdot h_y^+ = w_y^- \cdot h_y^-$ (paired rectangles have the same area) and since $h_y^- > \tilde{s}_t/2$, it follows that m_y^- is upper bounded by $w_y^-/2$. Therefore, the mass on v_y^- collected in Line 7 is enough for constructing all the new binary signals. The outcome of Algorithm 3 is therefore a feasible signaling scheme. ∎

We now show the following lemma that lower bounds each buyer's expected consumer surplus in $\mathcal{Z}_0'$ by half of the ironed expected consumer surplus:

LEMMA 3.7. *For any $x \in (0, 1]$, $s_{\mathcal{Z}_0'}(x) \ge \tilde{s}(x)/2$.*

Proof. Consider a single iteration of the loop Line 6 to Line 14, with fixed (t, y). For each equal-revenue binary signal $\mathcal{S}_j$ in $\mathcal{Z}_0$ with v_y^+ as taker, we have collected $\frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+}$ fraction of the giver mass from it in Line 10. When we combine this giver mass (denoted by M_j^{G}) with taker mass on value v_y^- to form an equal-revenue binary signal $\mathcal{S}'_{q'} = \mathcal{S}^{\text{E}}_{v_j^{\text{G}}, v_y^-}$, the taker mass in this binary signal (denoted by M_j^{T}) is thus

$$M_j^{\text{T}} = M_j^{\text{G}} \cdot \frac{v_j^{\text{G}}}{v_y^- - v_j^{\text{G}}} = \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot t_j^{\text{T}} \cdot \frac{v_j^{\text{T}} - v_j^{\text{G}}}{v_j^{\text{G}}} \cdot \frac{v_j^{\text{G}}}{v_y^- - v_j^{\text{G}}}$$
$$= \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot t_j^{\text{T}} \cdot \frac{v_j^{\text{T}} - v_j^{\text{G}}}{v_y^- - v_j^{\text{G}}}.$$

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Since buyers with value v_y^- gain surplus $(v_y^- - v_j^{\mathrm{G}})$ in this binary signal, the contribution from this signal to the total surplus of the buyers with value v_y^- is

$$M_j^{\mathrm{T}} \cdot (v_y^- - v_j^{\mathrm{G}}) = \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot t_j^{\mathrm{T}} \cdot (v_j^{\mathrm{T}} - v_j^{\mathrm{G}}).$$

Again assume the j^{th} equal-revenue binary signal added to $\mathcal{Z}_0$ during Algorithm 1 is $\mathcal{S}_j = \mathcal{S}^{\mathrm{E}}_{v_j^{\mathrm{G}}, v_j^{\mathrm{T}}}$ with weight γ_j. Let $t_j^{\mathrm{G}} = f_{\mathcal{S}_j}(v_j^{\mathrm{G}}) \cdot \gamma_j$ and $t_j^{\mathrm{T}} = f_{\mathcal{S}_j}(v_j^{\mathrm{T}}) \cdot \gamma_j$. Thus the total surplus of buyers with value v_y^- in the newly constructed signals $(\mathcal{S}'_{q'})$ is:

$$\sum_{j:\, v_j^{\mathrm{T}} = v_y^+} \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot t_j^{\mathrm{T}} \cdot (v_j^{\mathrm{T}} - v_j^{\mathrm{G}}) = \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot \sum_{j:\, v_j^{\mathrm{T}} = v_y^+} t_j^{\mathrm{T}} \cdot (v_j^{\mathrm{T}} - v_j^{\mathrm{G}})$$

(By the definition of buyer surplus) $$= \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} \cdot (\tilde{s}_t + h_y^+) \cdot f_{\mathcal{D}}(v_y^+)$$

(Since paired rectangles have the same area) $$= w_y^- \cdot h_y^-.$$

We now argue that, after the smoothing process, the expected consumer surplus of buyers with any value in the ironing interval I_t is at least $\tilde{s}_t/2$. We consider three cases based on the expected surplus value of the buyer in $\mathcal{Z}_0$:

1. Suppose that before the smoothing process, buyers with value v_i have surplus strictly less than $\tilde{s}_t/2$. The corresponding area below $\tilde{s}_t$ (on the mass coordinates of these buyers) consists of rectangles $a_{y_1}^-$ to $a_{y_\psi}^-$. Since all these rectangles participate in the smoothing process, we can sum up the surplus after smoothing from $y = y_1$ to $y = y_\psi$ as

   (Since $\sum_{y=y_1}^{y_\psi} w_y^- = m_i$ and $h_y^- > \tilde{s}_t/2$) $$\sum_{y=y_1}^{y_\psi} w_y^- \cdot h_y^- = m_i \cdot h_y^- > m_i \cdot \tilde{s}_t/2.$$

   Dividing by total mass m_i, the expected surplus of these buyers is at least $\tilde{s}_t/2$.

2. On the buyers with value v_i who originally have surplus less than $\tilde{s}_t$ but at least $\tilde{s}_t/2$, we have not changed any equal-revenue signal which contains v_i as taker value in the process. Their expected surplus does not change and is at least $\tilde{s}_t/2$.

3. On the buyers with value v_i who originally gain surplus more than $\tilde{s}_t$. Suppose their corresponding area above $\tilde{s}_t$ is split into rectangles $a_{y_1}^+$ to $a_{y_\xi}^+$. We have extracted at most $\sum_{y=y_1}^{y_\xi} \frac{w_y^+}{f_{\mathcal{D}}(v_y^+)} \cdot \frac{h_y^+}{\tilde{s}_t + h_y^+} = \frac{h_{y_1}^+}{\tilde{s}_t + h_{y_1}^+}$ fraction of the weight of their original signals. After the smoothing process, their expected surplus is still at least $\frac{\tilde{s}_t}{\tilde{s}_t + h_{y_1}^+} \cdot (\tilde{s}_t + h_{y_1}^+) = \tilde{s}_t$.

Combining the above three cases, on each ironing interval I_t, in all buyers have surplus at least $\tilde{s}_t/2$ in $\mathcal{Z}'_0$. This completes the proof. ∎

Finally, we make the scheme monotone. Suppose that $s_{\mathcal{Z}'_0}(x) > \tilde{s}(x)/2$ for some x corresponding to buyers of value v. For each equal-revenue binary signal in $\mathcal{Z}'_0$ where v is the taker value, we remove $\frac{2s_{\mathcal{Z}'_0}(x) - \tilde{s}(x)}{2s_{\mathcal{Z}'_0}(x)}$ fraction of the weight of the signal into two singleton signals. Denote the signaling scheme after the decomposition by $\mathcal{Z}_1$. Since the singleton signals do not provide buyer surplus, we have

$$\mathcal{CS}_v(\mathcal{Z}_1) = s_{\mathcal{Z}'_0}(x) \cdot \left(1 - \frac{2s_{\mathcal{Z}'_0}(x) - \tilde{s}(x)}{2s_{\mathcal{Z}'_0}(x)}\right) = \tilde{s}(x)/2.$$

The surplus-mass function of buyers in $\mathcal{Z}_1$ is exactly $\tilde{s}(x)/2$, completing the proof of Lemma 3.2.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

3.3.4 Proof of Theorem 3.1 Based on the construction of $\mathcal{Z}_0, \mathcal{Z}_0'$, and $\mathcal{Z}_1$, we prove our main theorem: $\mathcal{Z}_1$ is efficient, monotone, and 8-majorized.

Proof. (Proof of Theorem 3.1) In the whole construction process of $\mathcal{Z}_0$ and $\mathcal{Z}_1$, we only include equal-revenue binary signals (signals where the induced posterior only has two supports and they yield the same revenue for the seller, see Definition 3.2 as a formal definition) or singleton signals (signals where the induced posterior only has one support). The seller will always select the lowest support in the posterior as the price in these two types of signals. Therefore, the item is always sold, and $\mathcal{Z}_1$ is efficient. Further, the surplus-mass function of $\mathcal{Z}_1$ is $\tilde{s}/2$, which is monotonically increasing.

Finally, we show that $\mathcal{Z}$ is 8-majorized. Consider an arbitrary scheme $\mathcal{Z}'$ and any value $m \in (0, 1]$. Note that $\tilde{s}$ is a step function and it preserves the prefix sum of $s_{\mathcal{Z}_0}$ at break-point coordinates $m_0, m_1, \ldots, m_\tau$ (i.e. $\forall i \in [\tau]$, $\textsc{IPrefix}(\tilde{s}, m_i) = \textsc{IPrefix}(s_{\mathcal{Z}_0}, m_i)$) in ascending order. We have $m_0 = 0$ and $m_\tau = 1$. Since $0 < m \leq 1$, there exists a unique $k \in [\tau + 1]$ such that

$$m_{k-1} < m \leq m_k.$$

By setting $\lambda = \frac{m - m_{k-1}}{m_k - m_{k-1}}$, we have

$$m = (1 - \lambda) \cdot m_{k-1} + \lambda \cdot m_k.$$

Consider the sorted m-prefix sum of $s_{\mathcal{Z}'}$, we have

(By convexity of SPrefix as a function of m)
$$\begin{aligned}
\textsc{SPrefix}(s_{\mathcal{Z}'}, m) &\leq \textsc{SPrefix}(s_{\mathcal{Z}'}, m_{k-1}) + \lambda \cdot (\textsc{SPrefix}(s_{\mathcal{Z}'}, m_k) - \textsc{SPrefix}(s_{\mathcal{Z}'}, m_{k-1})) \\
&= (1 - \lambda) \cdot \textsc{SPrefix}(s_{\mathcal{Z}'}, m_{k-1}) + \lambda \cdot \textsc{SPrefix}(s_{\mathcal{Z}'}, m_k)
\end{aligned}$$
(Prefix sum is at least sorted prefix sum)
$$\leq (1 - \lambda) \cdot \textsc{IPrefix}(s_{\mathcal{Z}'}, m_{k-1}) + \lambda \cdot \textsc{IPrefix}(s_{\mathcal{Z}'}, m_k)$$
(By Lemma 3.1)
$$\leq 4 \cdot [(1 - \lambda) \cdot \textsc{IPrefix}(s_{\mathcal{Z}_0}, m_{k-1}) + \lambda \cdot \textsc{IPrefix}(s_{\mathcal{Z}_0}, m_k)]$$
(Since $\tilde{s}$ preserves the prefix sum of $s_{\mathcal{Z}_0}$ at m_{k-1} and m_k)
$$= 4 \cdot [(1 - \lambda) \cdot \textsc{IPrefix}(\tilde{s}, m_{k-1}) + \lambda \cdot \textsc{IPrefix}(\tilde{s}, m_k)]$$
(Since $\tilde{s}$ is monotone)
$$= 4 \cdot [(1 - \lambda) \cdot \textsc{SPrefix}(\tilde{s}, m_{k-1}) + \lambda \cdot \textsc{SPrefix}(\tilde{s}, m_k)]$$
(Since $\tilde{s}$ is constant on $(m_{k-1}, m_k]$)
$$\leq 4 \cdot \textsc{SPrefix}(\tilde{s}, m)$$
(By Lemma 3.2)
$$= 8 \cdot \textsc{SPrefix}(s_{\mathcal{Z}_1}, m).$$

Therefore, by Definition 2.9, $\mathcal{Z}_1$ is 8-majorized. ∎

4 Lower Bounds

Finally, we complement our 8-majorized signaling scheme with two lower bounds for finding α-majorized signaling schemes. Our first bound shows the impossibility of α-majorization for any constant α if we restrict to *buyer-optimal* signaling schemes, and the second shows the impossibility of α-majorization for $\alpha < 1.5$. (Recall that α-majorization is *equivalent* to simultaneous α-approximation to all symmetric, weakly increasing, concave, and normalized welfare functions (Proposition 2.2).) Both our hard instances are, in a sense, the simplest possible – they involve distributions $\mathcal{D}$ over only three different values. (For distributions over two values, there is an exactly majorized scheme, since agents with the lower value always get surplus 0.) Additionally, our lower bounds hold even for just two welfare functions: One is the utilitarian welfare that maximizes the sum of buyers' utilities, and the other is the "max-min" welfare that maximizes the minimum utility after excluding the buyers who cannot get positive utilities in any signaling scheme.

4.1 Incompatibility of Approximate Majorization and Buyer Optimality We now show that no buyer-optimal scheme (i.e., one that maximizes $\mathcal{CS}(\mathcal{Z})$) can yield our guarantees; in particular, no buyer-optimal scheme can be α-majorized, for any constant $\alpha \in \mathbb{R}^+$. This also shows that exact majorization is impossible since exact majorization implies buyer optimality. This motivates the need for approximations and the need for looking beyond buyer-optimal schemes.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Recall that a signaling scheme $\mathcal{Z}$ is *buyer-optimal* if $\sum_{i=1}^{n} f_{\mathcal{D}}(v_i) \cdot \mathcal{CS}_{v_i}(\mathcal{Z}) = \sum_{i=1}^{n} f_{\mathcal{D}}(v_i) \cdot v_i - \mathcal{R}^{\mathsf{My}}(\mathcal{D})$, where $\mathcal{R}^{\mathsf{My}}(\mathcal{D}) = \max_v v \cdot G_{\mathcal{D}}(v)$ is the optimal (Myerson) revenue under $\mathcal{D}$ (i.e. without signaling). The following observation is immediate.

LEMMA 4.1. *If a scheme is buyer-optimal, any optimal price for the original distribution will remain optimal in any signal of the scheme.*

Proof. Since the scheme is buyer-optimal, the item is always sold. The buyers' total surplus is thus the expected value of the buyer minus the seller's revenue. Suppose that p^{My} is an optimal price in the original distribution $\mathcal{D}$. The seller gets revenue $R^{\mathsf{My}} = p^{\mathsf{My}} \cdot G_{\mathcal{D}}(p^{\mathsf{My}})$ without signaling.

Note that the seller can still gain R^{My} in the scheme if they nevertheless post p^{My} for all the signals. Suppose for contradiction that in some signal, p^{My} is not an optimal price, the seller must gain strictly more revenue from posting price p' than that from posting p^{My}. Since in all other signals, the seller gains revenue at least as much as by posting price p^{My}, the seller's overall revenue is strictly greater than R^{My}. Therefore, $\mathcal{Z}_0$ cannot be buyer-optimal, leading to a contradiction. ∎

To prove the lower bounds for any signaling scheme, we need the following lemma to narrow the space of signaling schemes we are considering. The lemma says that, for any signaling scheme, we can always transform it into an equivalent scheme (by which we mean that the expected consumer surplus of any buyer remains unchanged) with a simple form. []DBLP:conf/soda/CummingsD0W20 have a similar observation, but we need the following lemma that provides finer structural characterizations.

LEMMA 4.2. *For any signaling scheme $\mathcal{Z}$, there exists an* efficient *signaling scheme $\mathcal{Z}^{\mathrm{E}}$, such that:*

1. *$\mathcal{Z}^{\mathrm{E}}$ generates the same expected consumer surplus as $\mathcal{Z}$ for any buyer;*
2. *$\mathcal{Z}^{\mathrm{E}}$ includes at most n signals, all of which have different lowest-supports (i.e. smallest value with non-zero mass).*

Proof. We conduct two operations on $\mathcal{Z}$ to construct $\mathcal{Z}'$. First, for each signal in $\mathcal{Z}$, we discard all masses on values strictly less than the optimal price. Then we put all the discarded masses into at most n singleton signals. Each singleton signal includes all the discarded masses on a different value v_i ($i \in [n]$). Consider any signal in $\mathcal{Z}$ before this operation. Since any buyer with the discarded values does not gain surplus from this signal and after discarding the optimal price remains the same, the expected consumer surplus of any buyer does not change.

Second, we combine all groups of signals with the same smallest value in support into one signal by adding up the masses on each value. Let $\mathcal{Z}'$ be the signaling scheme with all the signals after the combinations. After each combination operation, the optimal price remains the same at the smallest support. Therefore, this combination process again does not change the expected consumer surplus of any buyer. Since there are in total n different values, there are at most n signals in $\mathcal{Z}'$, each corresponding to a different smallest support.

Since the optimal price for each signal remains the lowest support after combination, the item is always sold. Therefore, $\mathcal{Z}'$ is efficient. ∎

THEOREM 4.1. *For any given $\alpha \geq 1$, there exist instances under which no buyer-optimal scheme is α-majorized.*

Proof. Consider the instance where there are three buyer values: $v_1 = 1$, $v_2 = N$, $v_3 = N+1$, and the probability masses of these values are

$$f_{\mathcal{D}}(v_1) = \frac{N^2-1}{N^3+2N^2+N},\ f_{\mathcal{D}}(v_2) = \frac{N^2+1}{N^3+2N^2+N},\ f_{\mathcal{D}}(v_3) = \frac{N^3+N}{N^3+2N^2+N}.$$

We will show that no buyer-optimal scheme can be α-majorized on this instance for $\alpha < N$. Suppose for the purpose of contradiction that $\mathcal{Z}$ is α-majorized where $\alpha < N$.

Notice that v_2 and v_3 are both optimal prices for the seller. By Lemma 4.2, we can transform $\mathcal{Z}$ into a signaling scheme $\mathcal{Z}'$ with at most three signals, each with a different smallest support. Since the transformation preserves any buyer's expected consumer surplus, $\mathcal{Z}'$ is still buyer-optimal.

By Lemma 4.1, in $\mathcal{Z}'$, we must include both v_2 and v_3 as the optimal prices. Therefore, there is no signal with v_3 as the smallest support. Since the signaling scheme is efficient (by buyer optimality), the smallest support in any signal must also be the optimal price. Therefore, we conclude that in $\mathcal{Z}'$, there are only two signals $\mathcal{S}_1$ and $\mathcal{S}_2$:

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

1. The first signal $\mathcal{S}_1$ includes all three values. Each value is an optimal price on the signal. Since each value as price provides the same revenue to the seller, the probability masses of $\mathcal{S}_1$ on the three values must be

$$f_{\mathcal{S}_1}(v_1) = \frac{N^2-1}{N^2+N};\ f_{\mathcal{S}_1}(v_2) = \frac{1}{N^2+N};\ f_{\mathcal{S}_1}(v_3) = \frac{N}{N^2+N}.$$

   Since only $\mathcal{S}_1$ include v_1, we have $\gamma_1 = f_{\mathcal{D}}(v_1)/\left(\frac{N^2-1}{N^2+N}\right) = \frac{1}{N+1}$.

2. The second signal includes only two values v_2 and v_3. Both of them are optimal prices. It is also an equal-revenue signal. The probability masses of $\mathcal{S}_2$ on the three values are

$$f_{\mathcal{S}_1}(v_1) = 0;\ f_{\mathcal{S}_1}(v_2) = \frac{1}{N+1};\ f_{\mathcal{S}_1}(v_3) = \frac{N}{N+1}.$$

   Since $\gamma_2 + \gamma_1 = 1$, we have $\gamma_2 = \frac{N}{N+1}$.

The expected consumer surplus from $\mathcal{Z}'$ of any buyer with value $v_2 = N$ is

$$\mathcal{CS}_{v_2}(\mathcal{Z}') = \frac{\gamma_1 \cdot f_{\mathcal{S}_1}(v_2) \cdot (N-1)}{f_{\mathcal{D}}(v_2)} = \frac{N-1}{N^2+1}.$$

The expected consumer surplus from $\mathcal{Z}'$ of any buyer with value $v_3 = N+1$ is

$$\mathcal{CS}_{v_3}(\mathcal{Z}') = \frac{\gamma_1 \cdot f_{\mathcal{S}_1}(v_3) \cdot N + \gamma_2 \cdot f_{\mathcal{S}_2}(v_3) \cdot 1}{f_{\mathcal{D}}(v_3)} = \frac{N+N^2}{N^2+1}.$$

We consider another scheme $\mathcal{Z}_2$ consisting of three signals $\mathcal{S}_1'$, $\mathcal{S}_2'$ and $\mathcal{S}_3'$:

$$f_{\mathcal{S}_1}(v_1) = \frac{N^2-1}{N^2+N};\ f_{\mathcal{S}_1}(v_2) = \frac{N+1}{N^2+N};\ f_{\mathcal{S}_1}(v_3) = 0;\ \gamma_1 = \frac{1}{N+1};$$
$$f_{\mathcal{S}_2}(v_1) = 0;\ f_{\mathcal{S}_2}(v_2) = \frac{N^2-N}{N^3-N};\ f_{\mathcal{S}_2}(v_3) = \frac{N^3-N^2}{N^3-N};\ \gamma_2 = \frac{N-1}{N+1};$$
$$f_{\mathcal{S}_3}(v_1) = 0;\ f_{\mathcal{S}_3}(v_2) = 0;\ f_{\mathcal{S}_3}(v_3) = 1;\ \gamma_3 = \frac{1}{N+1}.$$

In $\mathcal{Z}_2$, the expected consumer surplus of any buyer with value $v_2 = N$ is

$$\mathcal{CS}_{v_2}(\mathcal{Z}_2) = \frac{(N+1)\cdot(N-1)}{N^2+1} = \frac{N^2-1}{N^2+1};$$

The expected consumer surplus of any buyer with value $v_3 = N+1$ is

$$\mathcal{CS}_{v_3}(\mathcal{Z}_2) = \frac{(N^3-N^2)\cdot 1}{N^3+N} = \frac{N^2-N}{N^2+1};$$

Therefore, the smallest non-zero expected consumer surplus of a buyer in any buyer-optimal scheme is $\frac{N-1}{N^2+1}$ and the smallest non-zero expected consumer surplus of a buyer in $\mathcal{Z}_2$ is $\frac{N^2-N}{N^2+1}$. Since we have $\left(\frac{N^2-N}{N^2+1}\right)/\left(\frac{N-1}{N^2+1}\right) = N > \alpha$, $\mathcal{Z}$ is not α-majorized by $\mathcal{Z}_2$, leading to a contradiction. ∎

4.2 Lower Bound for Approximate Majorization Finally we provide a lower bound on approximate majorization under general signaling schemes. The following theorem shows that no scheme can be better than 1.5-majorized. This complements our upper bound of 8-majorization.

THEOREM 4.2. *For any $\alpha < 1.5$, there exist instances where no signaling scheme is α-majorized.*

Proof. Suppose that there exists a signaling scheme $\mathcal{Z}$ that is α-majorized. Let $\alpha = 1.5 - \delta$, where $\delta > 0$. Set ε such that $0 < \varepsilon \ll \delta$. There are three buyer values $\langle v_1, v_2, v_3\rangle = \langle 1, 1+\varepsilon, 2+\varepsilon\rangle$. The probability masses of $\mathcal{D}$ on

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

these values are $f_{\mathcal{D}}(v_1) = \varepsilon^2 + 2\varepsilon$, $f_{\mathcal{D}}(v_2) = 1 + (1+\varepsilon)^2$, $f_{\mathcal{D}}(v_3) = (1+\varepsilon) + (1+\varepsilon)^3$. For simplicity, we omit the multiplicative factor of normalizing the whole population to 1.

First, we compute the largest expected consumer surplus of the smallest-surplus buyer. By Lemma 4.2, we can transform $\mathcal{Z}$ into $\mathcal{Z}'$ without changing any buyer's expected consumer surplus. Assume $\mathcal{Z}'$ consists of the following three signals, each represented by the mass vector $\gamma_q \cdot \mathcal{S}_q$:

$$\gamma_1 \cdot \mathcal{S}_1 = \langle x, y, z\rangle;\ \gamma_2 \cdot \mathcal{S}_2 = \langle 0, y', z'\rangle;\ \gamma_3 \cdot \mathcal{S}_3 = \langle 0, 0, z''\rangle.$$

The expected consumer surplus of a buyer with value v_2 is $\frac{y\cdot\varepsilon}{f_{\mathcal{D}}(v_2)}$. The expected consumer surplus of a buyer with value v_3 is $\frac{z\cdot(1+\varepsilon)+z'\cdot 1}{f_{\mathcal{D}}(v_3)}$. The maximum of the smallest expected consumer surplus of a buyer with value v_2 or v_3 can be solved by the following LP:

$$\text{Maximize } s_{\min},\ \text{s.t.}$$

$$\frac{y\cdot\varepsilon}{f_{\mathcal{D}}(v_2)} \geq s_{\min}; \tag{4.8}$$

$$\frac{z\cdot(1+\varepsilon)+z'\cdot 1}{f_{\mathcal{D}}(v_3)} \geq s_{\min}; \tag{4.9}$$

$$1\cdot(x+y+z) \geq (1+\varepsilon)\cdot(y+z); \tag{4.10}$$

$$1\cdot(x+y+z) \geq (2+\varepsilon)\cdot z; \tag{4.11}$$

$$(1+\varepsilon)\cdot(y'+z') \geq (2+\varepsilon)\cdot z'; \tag{4.12}$$

$$0 \leq x \leq f_{\mathcal{D}}(v_1); \tag{4.13}$$

$$y + y' \leq f_{\mathcal{D}}(v_2); \tag{4.14}$$

$$z + z' + z'' \leq f_{\mathcal{D}}(v_3); \tag{4.15}$$

$$y, z, y', z', z'' \geq 0.$$

Eq. (4.8) and Eq. (4.9) mean that the minimum expected consumer surplus of any buyer with value v_2 or v_3 is at least the objective function. Eq. (4.10) and Eq. (4.11) are constraints on the masses of the first signal so that 1 is the optimal price. Similarly, Eq. (4.12) ensures that $v_2 = 1+\epsilon$ is the optimal price in the second signal. Eqs. (4.13) to (4.15) ensures that for any $i \in [3]$, the sum of the masses on a value v_i in all three signals does not exceed the total mass m_i.

One feasible solution of the LP has value

$$s^*_{\min} = \frac{4+3\varepsilon+\varepsilon^2}{2+\varepsilon}\cdot\frac{\varepsilon}{f_{\mathcal{D}}(v_2)} > \frac{2\varepsilon}{f_{\mathcal{D}}(v_2)},$$

and is obtained when

$$x = f_{\mathcal{D}}(v_1) = \varepsilon^2 + 2\varepsilon;$$
$$y = \frac{4+3\varepsilon+\varepsilon^2}{2+\varepsilon};$$
$$z = \frac{\varepsilon}{2+\varepsilon};$$
$$y' = f_{\mathcal{D}}(v_2) - y = \frac{\varepsilon^3+3\varepsilon^2+3\varepsilon}{2+\varepsilon};$$
$$z' = y'\cdot(1+\varepsilon) = \frac{(1+\varepsilon)\cdot(\varepsilon^3+3\varepsilon^2+3\varepsilon)}{2+\varepsilon}.$$

If $\mathcal{Z}'$ is α-majorized (and recall that $\alpha = 1.5-\delta$), the smallest positive expected consumer surplus of any buyer is at least

$$\frac{s^*_{\min}}{\alpha} = s^*_{\min}\cdot\frac{1}{1.5-\delta} > \frac{2\varepsilon}{f_{\mathcal{D}}(v_2)}\cdot\frac{1}{1.5-\delta} \geq \frac{2\varepsilon}{f_{\mathcal{D}}(v_2)}\cdot\left(\frac{2}{3}+\frac{4}{9}\delta\right).$$

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Therefore, we have that the expected consumer surplus of any buyer with value v_2 from $\mathcal{Z}'$ is at least $\frac{s^*_{\min}}{\alpha}$:

$$\mathcal{CS}_{v_2}(\mathcal{Z}') = \frac{y \cdot \varepsilon}{f_{\mathcal{D}}(v_2)} \geq \frac{s^*_{\min}}{\alpha} \geq \frac{2\varepsilon}{f_{\mathcal{D}}(v_2)} \cdot \left(\frac{2}{3} + \frac{4}{9}\delta\right),$$

and thus,

$$y \geq \frac{4}{3} + \frac{8}{9}\delta. \tag{4.16}$$

By Eq. (4.14): $y + y' \leq f_{\mathcal{D}}(v_2) = 2 + 2\varepsilon + \varepsilon^2$, and we have

$$y' \leq 2 + 2\varepsilon + \varepsilon^2 - \frac{4}{3} - \frac{8}{9}\delta = \frac{2}{3} + 2\varepsilon + \varepsilon^2 - \frac{8}{9}\delta. \tag{4.17}$$

By Eq. (4.10): $x+y+z \geq (1+\varepsilon)\cdot(y+z)$, we have $x \geq \varepsilon \cdot (y+z)$. Since $x \leq f_{\mathcal{D}}(v_1) = \varepsilon^2 + 2\varepsilon$, we have $y+z \leq 2+\varepsilon$ and thus

$$z \leq 2 + \varepsilon - y \leq 2 + \varepsilon - \left(\frac{4}{3} + \frac{8}{9}\delta\right) = \frac{2}{3} + \varepsilon - \frac{8}{9}\delta.$$

By Eq. (4.17) and Eq. (4.12), we have

$$z' \leq (1+\varepsilon) \cdot y' \leq (1+\varepsilon) \cdot \left(\frac{2}{3} + 2\varepsilon + \varepsilon^2 - \frac{8}{9}\delta\right).$$

Consider the overall surplus of all buyers from $\mathcal{Z}'$ (i.e. $\textsc{SPrefix}(\mathcal{Z}', \sum_{i\in[3]} f_{\mathcal{D}}(v_i))$. We have:

$$\begin{aligned}
\textsc{SPrefix}\left(\mathcal{Z}', \sum_{i\in[3]} f_{\mathcal{D}}(v_i)\right) &= \mathcal{CS}_{v_2}(\mathcal{Z}') \cdot f_{\mathcal{D}}(v_2) + \mathcal{CS}_{v_3}(\mathcal{Z}') \cdot f_{\mathcal{D}}(v_3) \\
&= y \cdot \varepsilon + z \cdot (1+\varepsilon) + z' \cdot 1 \\
&\leq (2 + 2\varepsilon + \varepsilon^2) \cdot \varepsilon + \left(\frac{2}{3} + \varepsilon - \frac{8}{9}\delta\right) \cdot (1+\varepsilon) + (1+\varepsilon) \cdot \left(\frac{2}{3} + 2\varepsilon + \varepsilon^2 - \frac{8}{9}\delta\right) \\
&= \frac{4}{3} - \frac{16}{9}\delta + o(\delta) \qquad \text{(Since } \varepsilon \ll \delta\text{)} \\
&< \frac{4}{3}.
\end{aligned}$$

Since v_3 is the optimal price on the original distribution, the optimal overall consumer surplus is:

$$\begin{aligned}
f_{\mathcal{D}}(v_1) \cdot v_1 + f_{\mathcal{D}}(v_2) \cdot v_2 + f_{\mathcal{D}}(v_3) \cdot v_3 - f_{\mathcal{D}}(v_3) \cdot v_3 &= f_{\mathcal{D}}(v_1) + (1+\varepsilon) \cdot f_{\mathcal{D}}(v_2) \\
&= (\varepsilon^2 + 2\varepsilon) + (1+\varepsilon)(2 + 2\varepsilon + \varepsilon^2) \\
&> 2.
\end{aligned}$$

Since $\left(\frac{4}{3}\right)/2 < \frac{1}{1.5-\delta}$, $\mathcal{Z}'$ cannot achieve α-approximation on the optimal overall consumer surplus. Therefore, $\mathcal{Z}$ is not α-majorized, leading to a contradiction. ∎

5 Conclusion

In our work, we investigate fairness in a third-degree price discrimination model, focusing on achieving an equitable distribution of consumer surplus among buyers. We introduce a novel signaling scheme that simultaneously approximates all non-negative, monotonically increasing, symmetric, and concave welfare functions by a factor of 8. It is an open direction to future improve our approximation upper bound of 8 or our lower bound of 1.5.

In our definition, a welfare function computes an aggregation of the expected consumer surplus of each consumer. It is intriguing to explore other fairness notions in this setting. For example, one can aim for fairness on *ex-post* surplus allocations. We leave these questions for future work.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

A Omitted Proofs

Proof. [Proof of Proposition 2.2] Let $\vec{u}$ be the vector of expected utilities of $\mathcal{Z}$. Let $\vec{w}$ be that vector of any signaling scheme. We wish to show $\mathcal{W}(\vec{u}) \geq \frac{1}{\alpha} \cdot \mathcal{W}(\vec{w})$.

By normalization (or non-negativity) and concavity, $\mathcal{W}(\vec{u}) \geq \frac{1}{\alpha} \cdot \mathcal{W}(\alpha\vec{u})$. There exists a $\vec{u}'$ (by starting from $\alpha\vec{u}$ and gradually decreasing its largest elements) such that $\vec{u}' \leq \alpha\vec{u}$, $\|\vec{u}'\|_1 = \|\vec{w}\|_1$, and $\vec{u}'$ is majorized by $\vec{w}$. Therefore, $\mathcal{W}(\vec{u}) \geq \frac{1}{\alpha} \cdot \mathcal{W}(\alpha\vec{u}) \geq \frac{1}{\alpha} \cdot \mathcal{W}(\vec{u}') \geq \frac{1}{\alpha} \cdot \mathcal{W}(\vec{w})$, where the penultimate inequality is because the welfare function is weakly increasing, and the last inequality is from Schur-concavity (see e.g. [43, p. 258]) of $\mathcal{W}$, implied by symmetry and concavity.

For the converse, notice that for any $m \in [0, 1]$, the sorted m-prefix sum is welfare function that is symmetric, non-decreasing, concave and normalized (or non-negative). ∎

References

[1] R. Alijani, S. Banerjee, K. Munagala, and K. Wang, *The limits of an information intermediary in auction design*, in Proceedings of the 23rd ACM Conference on Economics and Computation (EC), ACM, 2022, pp. 849–868.

[2] Y. Babichenko, I. Talgam-Cohen, H. Xu, and K. Zabarnyi, *Regret-minimizing bayesian persuasion*, in Proceedings of the 22nd ACM Conference on Economics and Computation (EC), ACM, 2021, p. 128.

[3] M. Bagnoli and T. Bergstrom, *Log-concave probability and its applications*, Economic Theory, 26 (2005), pp. 445–469.

[4] D. Bergemann, B. Brooks, and S. Morris, *The limits of price discrimination*, American Economic Review, 105 (2015), pp. 921–957.

[5] D. Bergemann, B. Brooks, and S. Morris, *First-price auctions with general information structures: Implications for bidding and revenue*, Econometrica, 85 (2017), pp. 107–143.

[6] D. Bergemann, P. Dütting, R. Paes Leme, and S. Zuo, *Calibrated click-through auctions*, in Proceedings of the 2022 ACM Web Conference (TheWebConf), ACM, 2022, pp. 47–57.

[7] D. Bergemann and S. Morris, *Information design: A unified perspective*, Journal of Economic Literature, 57 (2019), pp. 44–95.

[8] Y. Cai, F. Echenique, H. Fu, K. Ligett, A. Wierman, and J. Ziani, *Third-party data providers ruin simple mechanisms*, Proc. ACM Meas. Anal. Comput. Syst., 4 (2020), pp. 12:1–12:31.

[9] D. Chakrabarty and C. Swamy, *Approximation algorithms for minimum norm and ordered optimization problems*, in Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC), ACM, 2019, pp. 126–137.

[10] A. Chakraborty and R. Harbaugh, *Persuasive puffery*, Mark. Sci., 33 (2014), pp. 382–400.

[11] S. Chawla and M. Jagadeesan, *Individual fairness in advertising auctions through inverse proportionality*, in Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 42:1–42:21.

[12] S. Chawla, R. Rezvan, and N. Sauerberg, *Individually-fair auctions for multi-slot sponsored search*, in 3rd Symposium on Foundations of Responsible Computing (FORC), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 4:1–4:22.

[13] Y. Che and I. L. Gale, *The optimal mechanism for selling to a budget-constrained buyer*, Journal of Economic Theory, 92 (2000), pp. 198–233.

[14] R. Cummings, N. R. Devanur, Z. Huang, and X. Wang, *Algorithmic price discrimination*, in Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2020, pp. 2432–2451.

[15] DoubleClick, *Ad Exchange*. https://www.google.com/adexchange/AdExchangeOverview.pdf.

[16] S. Dughmi, *Algorithmic information structure design: a survey*, SIGecom Exch., 15 (2017), pp. 2–24.

[17] S. Dughmi, D. Kempe, and R. Qiang, *Persuasion with limited communication*, in Proceedings of the 17th ACM Conference on Economics and Computation (EC), ACM, 2016, pp. 663–680.

[18] S. Dughmi and H. Xu, *Algorithmic bayesian persuasion*, SIAM J. Comput., 50 (2021).

[19] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. S. Zemel, *Fairness through awareness*, in Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS), ACM, 2012, pp. 214–226.

[20] M. Essaidi and S. M. Weinberg, *On symmetries in multi-dimensional mechanism design*, in Proceedings of the 17th International Conference on Web and Internet Economics (WINE), Springer, 2021, pp. 59–75.

[21] V. Gkatzelis, D. Halpern, and N. Shah, *Resolving the optimal metric distortion conjecture*, in Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2020, pp. 1427–1438.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited

Downloaded 07/31/24 to 128.84.124.83 . Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

[22] A. Goel, A. K. Krishnaswamy, and K. Munagala, *Metric distortion of social choice rules: Lower bounds and fairness properties*, in Proceedings of the 18th ACM Conference on Economics and Computation (EC), ACM, 2017, pp. 287–304.
[23] A. Goel and A. Meyerson, *Simultaneous optimization via approximate majorization for concave profits or convex costs*, Algorithmica, 44 (2006), pp. 301–323.
[24] A. Goel, A. Meyerson, and S. A. Plotkin, *Approximate majorization and fair online load balancing*, ACM Trans. Algorithms, 1 (2005), pp. 338–349.
[25] S. Gupta, J. Moondra, and M. Singh, *Which lp norm is the fairest? approximations for fair facility location across all "p"""*, in Proceedings of the 24th ACM Conference on Economics and Computation (EC), ACM, 2023, p. 817.
[26] N. Haghpanah and R. Siegel, *Pareto-improving segmentation of multiproduct markets*, Journal of Political Economy, 131 (2023), pp. 1546–1575.
[27] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, Cambridge Mathematical Library, Cambridge University Press, 1952.
[28] J. D. Hartline, *Bayesian mechanism design*, Foundations and Trends in Theoretical Computer Science, 8 (2013), pp. 143–263.
[29] Ú. Hébert-Johnson, M. P. Kim, O. Reingold, and G. N. Rothblum, *Multicalibration: Calibration for the (computationally-identifiable) masses*, in Proceedings of the 35th International Conference on Machine Learning (ICML), PMLR, 2018, pp. 1944–1953.
[30] C. Ilvento, M. Jagadeesan, and S. Chawla, *Multi-category fairness in sponsored search auctions*, in Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT*), ACM, 2020, pp. 348–358.
[31] E. Kamenica and M. Gentzkow, *Bayesian persuasion*, American Economic Review, 101 (2011), pp. 2590–2615.
[32] J. Karamata, *Sur une inégalité relative aux fonctions convexes*, Publications de l'Institut mathematique, 1 (1932), pp. 145–147.
[33] M. J. Kearns, S. Neel, A. Roth, and Z. S. Wu, *Preventing fairness gerrymandering: Auditing and learning for subgroup fairness*, in Proceedings of the 35th International Conference on Machine Learning (ICML), PMLR, 2018, pp. 2569–2577.
[34] ———, *An empirical study of rich subgroup fairness for machine learning*, in Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency (FAT*), ACM, 2019, pp. 100–109.
[35] S. Ko and K. Munagala, *Optimal price discrimination for randomized mechanisms*, in Proceedings of the 23rd ACM Conference on Economics and Computation (EC), ACM, 2022, pp. 477–496.
[36] A. K. Krishnaswamy, Z. Jiang, K. Wang, Y. Cheng, and K. Munagala, *Fair for all: Best-effort fairness guarantees for classification*, in Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR, 2021, pp. 3259–3267.
[37] A. Kumar and J. M. Kleinberg, *Fairness measures for resource allocation*, SIAM J. Comput., 36 (2006), pp. 657–680.
[38] J.-J. Laffont and J. Robert, *Optimal auction with financially constrained buyers*, Economics Letters, 52 (1996), pp. 181–186.
[39] J. Mao, R. Paes Leme, and K. Wang, *Interactive communication in bilateral trade*, in Proceedings of the 13th Innovations in Theoretical Computer Science Conference (ITCS), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 105:1–105:21.
[40] Microsoft Advertising, *Ad Exchange*. https://ads.microsoft.com/.
[41] R. B. Myerson, *Optimal auction design*, Math. Oper. Res., 6 (1981), pp. 58–73.
[42] Pubmatic, *Ad Exchange*. https://pubmatic.com/.
[43] A. W. Roberts and D. E. Varberg, *Convex Functions*, Academic Press, 1973.
[44] A. Rubinstein and S. M. Weinberg, *Simple mechanisms for a subadditive buyer and applications to revenue monotonicity*, ACM Trans. Economics and Comput., 6 (2018), pp. 19:1–19:25.
[45] W. Shen, P. Tang, and Y. Zeng, *A closed-form characterization of buyer signaling schemes in monopoly pricing*, in Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS), IFAAMAS / ACM, 2018, pp. 1531–1539.
[46] A. Watts, *Fairness and efficiency in online advertising mechanisms*, Games, 12 (2021), p. 36.
[47] H. Xu, Z. Rabinovich, S. Dughmi, and M. Tambe, *Exploring information asymmetry in two-stage security games*, in Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI), AAAI Press, 2015, pp. 1057–1063.
[48] R. Xu, X. Zhang, P. Cui, B. Li, Z. Shen, and J. Xu, *Regulatory instruments for fair personalized pricing*, in Proceedings of the 2022 ACM Web Conference (TheWebConf), ACM, 2022, pp. 4–15.
[49] Yahoo Ads, *Ad Exchange*. https://policies.yahoo.com/us/en/yahoo/terms/yahooadexchange/index.htm.

Copyright © 2024 by SIAM
Unauthorized reproduction of this article is prohibited