

Can Buyers Reveal for a Better Deal?

Daniel Halpern
Harvard University
dhalpern@g.harvard.edu

Gregory Kehne
Harvard University
gkehne@g.harvard.edu

Jamie Tucker-Foltz
Harvard University
jtuckerfoltz@gmail.com

June 25, 2021

Abstract

We study small-scale market interactions in which buyers are allowed to credibly reveal partial information about their types to the seller. Previous recent work has studied the special case where there is one buyer and one good, showing that such communication can simultaneously improve social welfare and ex ante buyer utility. With multiple buyers, we find that the buyer-optimal signalling schemes from the one-buyer case are actually harmful to buyer welfare. Moreover, we prove several impossibility results showing that, with either multiple i.i.d. buyers or multiple i.i.d. goods, maximizing buyer utility can be at odds with social efficiency, which is a surprising contrast to the one-buyer, one-good case. Finally, we investigate the computational tractability of implementing desirable equilibrium outcomes. We find that, even with one buyer and one good, optimizing buyer utility is generally NP-hard, but tractable in a practical restricted setting.

1 Introduction

It is common knowledge that Bayesian-optimal mechanisms for revenue maximization lend themselves to inefficiency. A seller may rationally refuse to sell to buyers who are unwilling to pay a high price, even if there is an acceptable lower price at which the seller could still make a substantial profit. But what if a buyer is able to *prove* to the seller that they are unwilling to pay the high price? Upon receiving such a proof, the only rational course of action is for the seller to offer a lower price. As a result, both the buyer and the seller will see a welfare improvement.

However, the possibility of such communication will undoubtedly give rise to secondary market effects. Will the seller infer that a buyer has a higher valuation simply because they do not choose to disclose such a proof, and if so, should the seller raise their price even higher? And if there are multiple buyers competing for a single item, how will the disclosures of one buyer affect the ultimate welfare of another? To realistically discuss the overall welfare implications, it is thus necessary to investigate not just specific one-shot interactions, but the equilibria of the game played between the seller and the buyer(s).

This inquiry is inspired by the realm of online commerce, where the increasing accessibility and quality of buyer data makes the personalized pricing of goods an ever more attractive prospect, and has served as the motivation of previous work studying the impact of information signalling on buyer welfare in auctions [3, 1]. Motivated by the prospect of a future in which consumers are able to exert precise control over their online data (and a perhaps more immediate future in which sellers implement personalized pricing), we aim to answer the question,

“Can consumers benefit from the ability to share their private data, and if so, how?”

In a recent work, Ali, Lewis, and Vasserman [1] initiate the study of how such *voluntary disclosure* capabilities can improve welfare, considering a handful of special cases. In their model, a prospective buyer is allowed to credibly disclose to the seller a set of possible types containing their true type; the seller then sets prices based on this information. They report overwhelmingly positive news for consumers. On the one hand, when there is one buyer, one seller, and one good, they demonstrate that there always exists a disclosure strategy for the buyer such that

- no buyer has an incentive to deviate from the strategy after learning their type (technically, the strategies are part of a sequentially-rational Bayes-Nash equilibrium),
- the good is always sold,
- the seller is weakly better off than they would be without disclosure, and
- every interim buyer type is weakly better off than they would be without disclosure.

For a parameterized family of canonical probability distributions over the buyer’s value for the good (including the uniform distribution on $[0, 1]$), they show that it is possible to strictly increase ex ante buyer utility as well. Furthermore, there is an intuitive characterization of the buyer-optimal equilibrium, determined by the limit of a greedy algorithm that iteratively constructs better and better equilibria by having all buyer types who are not sold the good declare to the seller that they are of such a type. In the end we are left with a *partitional equilibrium*, in which there is some partition \mathcal{P} of the type space, and every buyer reveals the set in \mathcal{P} to which their type belongs.

Ali, Lewis, and Vasserman also study a case of multiple sellers; in particular a setting in which there are two sellers, and the lone buyer has strong but private preferences over from whom to buy. Here they show that the buyer may leverage selective (seller-personalized) disclosure in order to play them off of one another and again compel their personalized prices to increase the buyer’s own expected utility.

However, the settings of these results differ markedly from most online commerce, and it is in the direction of these differences which we depart.

In Section 3 we show that, with either **multiple buyers** or **multiple goods**, maximizing buyer surplus may require the seller to sometimes *not* sell all of the goods (Theorem 3.1). Since the model assumes the seller has no cost to sell the goods, this shows that social efficiency is not always compatible with maximizing buyer welfare, which lies in stark contrast to the one-buyer, one-good case. Our result holds even with the restrictions that buyer valuations are additive and independent across goods, as well as independent across buyers.

Section 4 presents a case study of the setting with one good and two i.i.d., uniform $[0, 1]$ buyers, in which many of the phenomena from Section 3 can be better understood visually. We first investigate whether the analogues of the optimal one-buyer equilibria still result in buyer welfare improvements. Surprisingly, we find that the answer is negative for buyer-symmetric disclosure strategies. Perhaps even more surprisingly, it is possible to improve the expected buyer surplus (sum of both buyers’ utilities) by having only one buyer disclose information about their type (though this harms the utility of the other buyer). As for the question of social efficiency, with a few additional assumptions in the spirit of [1], we show an extreme impossibility result (Theorem 4.4), again in contrast to the one-buyer case: in any equilibrium where the good is always allocated to the highest bidder, both buyers must always receive utility zero.

Beyond these results, there is not much more that can be said about the existence of desirable equilibria in general. The usefulness of voluntary disclosure depends heavily on the structure of the common priors over buyer valuations. Thus, a natural followup question concerns the computation of optimal disclosure strategies. We model this problem by approximating arbitrary priors by

discrete probability distributions with finite support, which are encoded as part of the input. We show that, while it is possible to efficiently compute the buyer-optimal equilibrium in the restricted setting from [1] where disclosure messages must be “connected” (Theorem 5.1), the more general problem is (weakly) NP-hard (Theorem 5.2), and is not even well-approximated by the algorithm for the connected case (Proposition 5.3).

1.1 Related literature

This work falls within a larger body of literature on the implications of information signalling in markets, and how strategic disclosure affects equilibria, as in the work of Gentzkow and Kamenica [4] on *Bayesian persuasion*. The signals which we consider are the *verifiable disclosures* of [1], in which buyers send public signals about their types from a set of possible signals which is demonstrably truthful.

There is a parallel line of work aimed at improving buyer welfare via an intermediary that observes buyer valuations and communicates a signal to the seller based on some prearranged signalling scheme. Bergemann, Brooks, and Morris [3] address the case of one buyer and one seller. They show that, if point *D* in Figure 1 represents the welfare of the buyer and seller in the absence of any communication by the intermediary, then any combination of welfares in the shaded triangle is attainable via randomized signalling schemes. Alijani, Banerjee, Munagala, and Wang [2] generalize this to the setting with multiple buyers who share an intermediary, showing that the line segment *BE* is unattainable, though it is possible to approximate the maximum possible buyer welfare when the buyers’ types are independent and identically distributed.

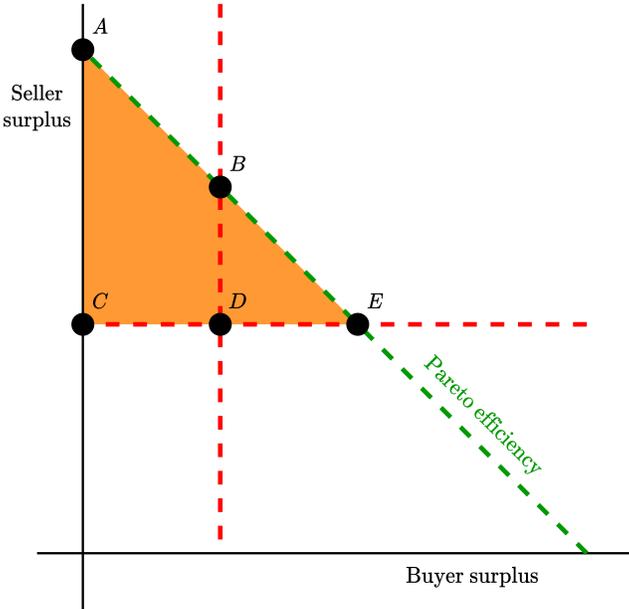


Figure 1: Feasible buyer/seller welfare outcomes from Bergemann, Brooks, and Morris [3].

Similar to our setting is that of Shen, Tang, and Zeng [7], which generalizes [3] to the case of multiple buyers. For them, the buyers each simultaneously signal according to pre-committed signalling schemes to the seller, who then conducts a Myerson auction based on their updated priors. They study the equilibria of this game and determine buyer-optimal equilibria for certain

classes of buyer type distributions, but the buyer disclosures are neither deterministic nor verifiable. Sher and Vohra [8] consider a setting where the disclosures *are* verifiable, and there is commitment from both the buyer and the seller.

In terms of Figure 1, in our setting *with* verifiability but *without* commitment to a signalling scheme, the main result of [1] is that there is some feasible point on line segment BE when there is only one buyer and one good, which we show is false for either multiple buyers or multiple goods. In fact, for multiple goods, we give an example where the only feasible point on segment AE is at A , where the buyer receives no surplus at all (see Section 3.2).

2 Model

As in [1], we operate in the context of a verifiable *disclosure game*. We begin by defining the most general, abstract form of the game, with multiple buyers and multiple goods, but still only one seller. Suppose there are m goods, numbered $1, 2, \dots, m$. We are concerned only with additive valuations, so we denote the type space of each buyer by $\mathbb{R}_{\geq 0}^m$, where, for any $v = (v_1, v_2, \dots, v_m) \in \mathbb{R}_{\geq 0}^m$, each v_k denotes the value the buyer has for good k . There is a common prior over the buyers' values, in which the value any fixed buyer has for different goods may be correlated; but in all of the multiple-buyer scenarios we consider in this paper, the values of different buyers are independent. The disclosure game proceeds in two stages:

1. Each buyer simultaneously observes their value $v \in \mathbb{R}_{\geq 0}^m$ and publicly sends a *message* in the form of a set $M \subseteq \mathbb{R}_{\geq 0}^m$ such that $v \in M$.
2. The seller sells the good(s) to the buyer(s) so as to maximize their expected revenue, taking into account the information conferred by M .

Note that the requirement that $v \in M$ is a key feature of this game: buyers cannot lie about their types in the disclosure stage. We study the subgame-perfect pure-strategy Bayes-Nash equilibria of this game, and evaluate them with respect to *ex ante buyer surplus*, defined as the expected sum of all buyer utilities.

With one buyer and one good, step 2 simply involves the seller choosing a posted price and the buyer accepting or rejecting. With one buyer and multiple goods, the seller posts a menu of bundles of goods, each with an associated price, and the buyer may choose one of them. With one good and multiple buyers, the seller runs a Myerson auction. We will introduce new notation to describe these specialized settings as needed.

2.1 Special equilibria

An equilibrium is *efficient* if the good(s) are always sold to the buyer with the highest value. An equilibrium is *partitional* if each buyer's messaging strategy is induced by a partition \mathcal{P} of $\mathbb{R}_{\geq 0}^m$, where the buyer reveals to which set in \mathcal{P} their value belongs. A *connected* partitional equilibrium is one in which the interiors of the convex hulls of the messages in \mathcal{P} are pairwise disjoint, in which case we say the messages are *connected* and that \mathcal{P} is a *connected partition*. As it is shown in [1], to maximize expected buyer utility in the one-buyer one-good case, it is without loss of generality to restrict attention to efficient, partitional equilibria.

Lemma 2.1 (Efficiency Lemma). *Suppose there is one good and one buyer. Given any pure-strategy equilibrium of the disclosure game, there exists a pure-strategy equilibrium that is efficient and results in the same payoff for every buyer type.*

Lemma 2.2 (Partitional Lemma). *Given any pure-strategy equilibrium of the disclosure game, there exists a pure-strategy equilibrium that is partitional and results in the same selling mechanism for every buyer type.*

The Efficiency Lemma is proved by having all buyer types that do not get the good reveal their type to the seller. One can easily check that this does not change seller incentives when faced with one of the other buyer types. The Partitional Lemma is proved by partitioning the buyer types by the message that each type sent in the original equilibrium, thus conveying the same information to the seller.¹ While the partitional lemma continues to hold for multiple buyers and/or multiple goods, we show in the next section that the Efficiency Lemma does not.

3 General impossibility results

In this section, we exhibit a series of counterexamples to prove the following theorem.

Theorem 3.1. *In any setting with multiple goods or multiple buyers, there exist common priors over buyer valuation functions such that:*

- *The buyer valuation functions are pairwise independent across different buyers.*
- *Every buyer’s valuation function is additive and independent across different goods.*
- *In any pure-strategy equilibrium of the disclosure game in which all goods are always sold, buyer surplus is strictly lower than it would be in the absence of disclosure. (Additionally, in the case of multiple goods, this buyer surplus must be zero.)*

This implies that the Efficiency Lemma does not hold beyond the limited context of one good and one buyer. In other words, social efficiency may be incompatible with maximizing expected buyer welfare. We begin by discussing the computational approach used to verify the claims about optimal mechanisms made in this section.

3.1 Computing the optimal mechanism over discrete distributions

Suppose the buyers have sent messages to the seller, and now the seller is deciding how to optimally sell the goods. By the revelation principle, it is without loss of generality to consider selling mechanisms in which each buyer is asked to exactly reveal the values they have for each good (but, unlike in the initial disclosure, now they are allowed to lie, so the mechanism must be incentive compatible). If there are only a finite number of types each buyer can have, then we can write the seller’s optimization problem as a linear program, as follows.

Suppose there are ℓ buyers and m goods. For any $j \in [\ell] := \{1, 2, \dots, \ell\}$, suppose there are n_j possible types that buyer j can have. Let T denote the joint type space of all buyers,

$$T := \prod_{j \in [\ell]} [n_j],$$

where an element $\mathbf{i} \in T$ denotes a vector of types for each buyer, $\mathbf{i} = i_1, i_2, \dots, i_\ell$. Analogously, for any buyer $j \in [\ell]$, let T^{-j} denote the joint type space of all buyers other than j ,

$$T^{-j} := \prod_{j \in [\ell] \setminus \{j\}} [n_j],$$

¹The version of the Partitional Lemma proved in [1] is slightly different, since it is proved in the more complicated setting where seller messages are required to be connected intervals. It only applies to efficient equilibria, and only guarantees the same selling mechanism for *almost* every buyer type.

where an element $-\mathbf{i} \in T^{-j}$ denotes a vector of types for each buyer other than j . For any type vector $\mathbf{i} \in T$, suppose that $p_{\mathbf{i}}$ denotes the probability that \mathbf{i} is realized according to the common prior. Finally, for any buyer $j \in [\ell]$, type $i \in [n_j]$, and good $k \in [m]$, suppose that $v_{i,j,k}$ is the value that type i of buyer j has for good k .

Given this information, a selling mechanism is fully specified the price $r_{\mathbf{i},j}$ that each buyer j pays when types are revealed according to \mathbf{i} , and the quantity $q_{\mathbf{i},j,k}$ of each good k that each buyer j is allocated when types are revealed according to \mathbf{i} . The optimal \mathbf{r} and \mathbf{q} are determined by the following linear program, which maximizes expected revenue subject to individual rationality (IR) and incentive compatibility (IC) constraints.

$$\begin{array}{ll}
\text{maximize} & \sum_{\mathbf{i} \in T} \sum_{j \in [\ell]} p_{\mathbf{i}} r_{\mathbf{i},j} \\
\\
& r_{\mathbf{i},j} \geq 0 & \text{for all } \mathbf{i} \in T, j \in [\ell] \\
& q_{\mathbf{i},j,k} \geq 0 & \text{for all } \mathbf{i} \in T, j \in [\ell], k \in [m] \\
& \sum_{j \in [\ell]} q_{\mathbf{i},j,k} \leq 1 & \text{for all } \mathbf{i} \in T, k \in [m] \\
\text{subject to} & \sum_{k \in [m]} v_{\mathbf{i},j,k} q_{\mathbf{i},j,k} - r_{\mathbf{i},j} \geq 0 & \text{for all } \mathbf{i} \in T, j \in [\ell] \text{ (IR)} \\
& \sum_{k \in [m]} v_{(i,-\mathbf{i}),j,k} q_{(i,-\mathbf{i}),j,k} - r_{(i,-\mathbf{i}),j} \geq \sum_{-\mathbf{i} \in T^{-j}} & \text{for all } j \in [\ell], -\mathbf{i} \in T^{-j}, \\
& p_{(i,-\mathbf{i})} \left(\sum_{k \in [m]} v_{(i',-\mathbf{i}),j,k} q_{(i',-\mathbf{i}),j,k} - r_{(i',-\mathbf{i}),j} \right) & i, i' \in [n_j] \text{ (IC)}
\end{array}$$

Note that the optimal solution may involve non-integral values of $q_{\mathbf{i},j,k}$, which can be interpreted as giving good k to buyer j with probability $q_{\mathbf{i},j,k}$ when types are realized according to \mathbf{i} . Also, a more general version of this LP would allow for ex post violations of IR or IC; the more constrained version above is equivalent in terms of the optimal value and resulting buyer surplus, and it produces more reasonable-looking mechanisms.

We implemented an algorithm to compute the optimal mechanism with respect to any seller belief distribution, using Gurobi [5] to solve the linear program, with the secondary optimization objective of maximizing buyer surplus. For any prior distribution, we can enumerate over all tuples of partitions of buyer types and compute expected buyer/seller utilities in the corresponding LP for each possible vector of disclosure messages, then combine to compute the overall expected utilities of the disclosure game. We ran thousands of trials with valuations and probabilities drawn independently and uniformly at random from $[0, 1]$; the counterexamples in Section 3.2 were discovered in this process, then subsequently simplified.

3.2 Counterexamples with multiple goods

Surprisingly, even for the simple case of **one buyer, two goods**, and two buyer types, it is possible for the partition $\{\{1, 2\}\}$ (in which there is no voluntary disclosure) to be the best partition (and thus the best pure-strategy equilibrium) in terms of ex ante buyer utility, even though it results in the buyer sometimes not getting one of the goods. The following is a simple example in which the valuations of the two goods are correlated.

Type	Probability	Value for good 1	Value for good 2
1	1/2	3	4
2	1/2	4	9

In the no-disclosure equilibrium induced by the partition $\{\{1, 2\}\}$, the unique optimal mechanism is for the seller to post the following menu of choices.

Bundle	Price
Only good 1	3
Both goods 1 and 2	12

Type 1 buyers purchase only good 1, and type 2 buyers purchase both goods 1 and 2. Notice that type 2 buyers get utility $1 > 0$. This is because the seller is unable to extract any more utility from them, for otherwise they would opt to only buy good 1 (and if the seller raised the price on good 1, they would completely exclude type 1 buyers, hurting revenue even more). There is some inefficiency though, as good 2 is only sold with ex ante probability $\frac{1}{2}$, even though the buyer always has positive utility for it. The only other partition to consider is $\{\{1\}, \{2\}\}$, in which the buyer exactly reveals their type. While this always yields an efficient outcome, the seller will clearly be able to extract all of the surplus, leaving the buyer with utility zero. Thus, we conclude that the Efficiency Lemma no longer holds when there are 2 goods.

In fact, it turns out that this can happen even when the valuations for the two goods are independent, as shown in the following example (which has been simplified as much as possible).

Type	Probability	Value for good 1	Value for good 2
1	$0.15 \cdot 0.4$	56	38
2	$0.15 \cdot 0.6$	56	69
3	$0.85 \cdot 0.4$	91	38
4	$0.85 \cdot 0.6$	91	69

In this case, in the unique optimal mechanism is the following.

Bundle	Price
Good 2 for sure, and good 1 with probability $31/35$	118.6
Both goods 1 and 2	129

A type 1 buyer will purchase nothing, a type 2 buyer will purchase the randomized bundle, and types 3 and 4 purchase both goods. Note that the $31/35$ is the maximum probability at which types 3 and 4 weakly prefer buying both goods. Only type 4 gets nonzero utility. By exhaustively checking each of the 14 alternative partitions of $\{1, 2, 3, 4\}$, we verified that, in any pure-strategy equilibrium in which both goods always get sold, the ex ante buyer utility is strictly lower.

3.3 Counterexample with multiple buyers

Now suppose there is **one good** and **two buyers**, whose valuations for the good are independent and identically distributed as follows.

Type	Probability	Value for good
1	$1/4$	1
2	$1/4$	2
3	$1/2$	3

In the absence of disclosure, an optimal mechanism for the seller (as verified computationally) is as follows. If at least one buyer bids 3, randomly choose one such buyer and sell them the good for 2.5. If no buyer bids 3, but at least one buyer bids 2, randomly choose one such buyer and sell

them the good for 2. Otherwise, do not sell the good. The expected buyer surplus is $\frac{3}{4} \cdot \frac{1}{2} = \frac{3}{8}$, as there is a $\frac{3}{4}$ chance that there will be some buyer of type 3, in which case the buyer surplus will be $\frac{1}{2}$. However, the no-disclosure equilibrium is clearly not efficient, as the good is not sold $\frac{1}{16}$ of the time (when both buyers have type 1).

An intuitive idea to repair this equilibrium so that it is efficient is to use the partition $\{\{1\}, \{2, 3\}\}$. By separating the lowest-value type from the other types, we can prevent the inefficient outcome, since there is no longer any incentive for the seller to refuse to sell. In fact, in the one-buyer case, this partition is precisely the one constructed in the proof of the Efficiency Lemma. In this case, when the buyer sends the message $\{1\}$, the seller sells the good for a price of 1, and efficiency is regained with no loss to buyer surplus; and when the buyer sends the message $\{2, 3\}$, the seller's incentives are the same as in the original equilibrium, so again, there is no loss to buyer surplus.

However, with two buyers, the seller's incentives actually *do* change, and for this particular example, they change in a way that *harms* buyer surplus. In the partitional equilibrium where both buyers use the partition $\{\{1\}, \{2, 3\}\}$ and both send the message $\{2, 3\}$, the optimal mechanism is the same as in the no-disclosure equilibrium, except that the price the seller charges to a buyer who bid 3 is 2.75 instead of 2.5, and hence the buyer surplus is $\frac{8}{9} \cdot \frac{1}{4} = \frac{2}{9} < \frac{3}{8}$. Even worse, when one buyer sends the message $\{1\}$ and the other buyer sends the message $\{2, 3\}$, having the outside option of selling the good to the first buyer for 1 strictly incentivizes the seller to only sell to the second buyer for a price of 3. This means that the buyer surplus is zero, and the good might be sold to a buyer of value 1 over a buyer of value 2. So overall, using the partition $\{\{1\}, \{2, 3\}\}$ for both buyers results in a lower surplus of $\frac{3}{4} \cdot \frac{3}{4} \cdot \frac{2}{9} = \frac{1}{8}$, yet still does not resolve the inefficiency. In fact, the inefficiency is even greater, as the seller's welfare gains do not offset the buyers' welfare losses.

Having only one buyer use the partition $\{\{1\}, \{2, 3\}\}$ and the other buyer use the no-disclosure partition $\{\{1, 2, 3\}\}$ suffers from similar problems. While the total social welfare is the same, the buyer surplus is lower, and the equilibrium is still inefficient. By enumerating over all pairs of partitions of the set $\{1, 2, 3\}$, we verified that, despite its inefficiency, the no-disclosure equilibrium yields the strictly highest expected buyer surplus.

These negative secondary effects of disclosure are not mere anomalies specific to our discrete counterexample. As we show in the next section, they manifest themselves even in the simple setting where both buyers' valuations are drawn uniformly from the interval $[0, 1]$.

4 The case of two uniform $[0, 1]$ buyers

For any $a \leq b$, let $U[a, b]$ denote the uniform distribution on $[a, b]$. In this section we consider the special case where there are two buyers, A and B , with valuations v_A and v_B for a single good drawn independently from $U[0, 1]$.

By the Partitional Lemma (Theorem 2.2), we may restrict attention to equilibria of the disclosure game for which A and B report messages P_A and P_B from \mathcal{P}_A and \mathcal{P}_B , which are partitions of $[0, 1]$. In this section, we will only be concerned with the special case where \mathcal{P}_A and \mathcal{P}_B are connected, i.e., each element is an interval. Again, as in the one-buyer setting, all pairs of partitions $\mathcal{P}_A, \mathcal{P}_B$ are supportable as equilibria in this game. If the seller holds the off-path belief that, upon receiving from A any message $M \notin \mathcal{P}_A$, the valuation of A is $\bar{v}_A := \max\{M\}$ with probability 1,² then A is guaranteed not to derive any utility from the resulting Myerson auction, since in order

²This maximum may not be well-defined in general, but it is well-defined for connected, partitional equilibria with a $U[0, 1]$ prior, since we may ensure all messages are closed on the top end. This can only improve buyer utilities, and does not change seller incentives since $U[0, 1]$ has no atomic points.

to receive the good they must clear the Myerson reserve of $\overline{v}_A \geq v_A$. Facing a seller with these off-path beliefs, A is therefore incentivized to report the unique $P_A \in \mathcal{P}_A$ for which $v_A \in P_A$, since otherwise A is guaranteed to receive no utility. The same argument of course applies to B .

Since we are concerned with buyers with values drawn from $U[0, 1]$ sending interval messages, we will make extensive use of the following observation:

Lemma 4.1. *A buyer with value $v \sim U[a, b]$ has virtual value $\phi(v) = 2v - b$.*

Proof. For this distribution, $F(v) = \frac{v-a}{b-a}$ and $f(v) = \frac{1}{b-a}$, and so $\phi(v) = v - \frac{1-F(v)}{f(v)} = 2v - b$. \square

4.1 The first step of Zeno's partition

In [1], the optimal equilibrium for the one-buyer $U[0, 1]$ distribution is induced by ‘‘Zeno's partition,’’

$$\mathcal{P}_Z := \{(2^{-k-1}, 2^{-k}] \mid k \in \mathbb{Z}_{\geq 0}\} \cup \{0\}.$$

It is constructed through a sequence of steps from the no-disclosure equilibrium, where in each step, all buyer types who are currently not sold the good are separated into a new element of the partition. In the one-buyer case, each step is a Pareto improvement for all buyer types; let us now consider the two-buyer case and see what happens when we implement the first step of Zeno's partition for both buyers simultaneously.

In the no-disclosure equilibrium (where each buyer always sends the message $[0, 1]$), the seller runs a second-price auction with reserve price $\frac{1}{2}$. A simple calculation shows that the expected buyer surplus is $\frac{1}{6}$. For a plot of the allocation as a function of (v_A, v_B) see Figure 2a.

Now consider what happens when buyers reveal whether their valuation is greater than or less than $\frac{1}{2}$. There are 3 cases to consider. If both buyers have value less than $\frac{1}{2}$, then the seller will run a second-price auction with reserve price $\frac{1}{4}$, so the result will be the same as in the no-disclosure case, but with everything scaled down by a factor of 2. Thus, the expected buyer surplus will be $\frac{1}{2} \cdot \frac{1}{6} = \frac{1}{12}$. If both buyers have value $\geq \frac{1}{2}$, it is not too hard to check that the seller still runs a second-price auction with a reserve price of $\frac{1}{2}$ (or equivalently, no reserve price). Thus we have the same outcome as in the no-disclosure case for an expected buyer surplus of $\frac{1}{6}$.

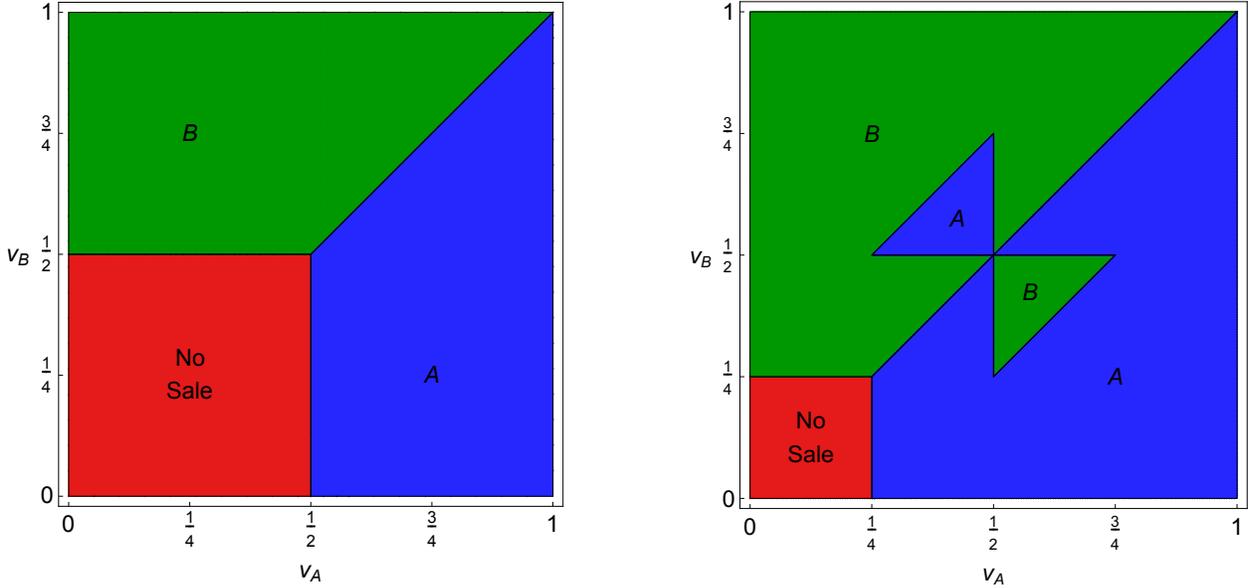
So far, in the first two cases, everything is analogous to the 1-buyer case: if the low-value types disclose that they have low value, we see an improvement in buyer welfare, whereas if there is no disclosure, the seller incentives remain the same, so the buyers achieve the same welfare. However, in the third case, where one buyer discloses they have a low value and the other does not, something quite different happens: we see competition between the two buyers that ultimately reduces the welfare of the high-value buyer.

Suppose buyer A reveals that $v_A \in [0, \frac{1}{2}]$ and buyer B reveals that $v_B \in [\frac{1}{2}, 1]$. Applying Theorem 4.1, the virtual values of the two buyers and their inverses are as follows:

$$\phi_A(v_A) = 2v_A - \frac{1}{2} \quad \phi_B(v_B) = 2v_B - 1 \quad \phi_A^{-1}(x) = \frac{x}{2} + \frac{1}{4} \quad \phi_B^{-1}(x) = \frac{x}{2} + \frac{1}{2}.$$

Note that, for $v_B \in [\frac{1}{2}, 1]$, $\phi_B(v_B) \geq 0$, so the good is always sold under the Myerson auction. However, for $v_B < v_A + \frac{1}{4}$, the good will be sold to buyer A since $\phi_B(v_B) < \phi_A(v_A)$, even though they have a lower valuation, just as in Section 3.3. Pictorially, the sale outcomes for this case are as shown in the top-left quadrant of Figure 2b, where the small blue triangle represents the case where the seller sells the good to A because $v_B < v_A + \frac{1}{4}$.

We can immediately deduce that, for small $\varepsilon > 0$, a buyer with value $\frac{1}{2} + \varepsilon$ gets a lower interim expected utility compared to the no-disclosure equilibrium. For in the no-disclosure equilibrium,



(a) No disclosure: $\mathcal{P}_A = \mathcal{P}_B = \{[0, 1]\}$, yielding an expected buyer surplus of $\mathbb{E}[U] = \frac{1}{6}$.

(b) High/low disclosure: $\mathcal{P}_A = \mathcal{P}_B = \{[0, 1/2], [1/2, 1]\}$, with $\mathbb{E}[U] = \frac{1}{6}$.

Figure 2: Allocations of the good under different partitional disclosure strategies for A and B . Green denotes that A receives the good, blue denotes that B receives the good; red denotes no sale. Expected buyer surplus $\mathbb{E}[U]$ is given for each case.

they are sold the good with probability roughly $\frac{1}{2}$, and when they do win, they gain an expected utility of roughly ε . However, in this new equilibrium, they are only sold the good with probability roughly $\frac{1}{4}$, yet still only gain an expected utility of roughly ε when they win. Thus, in contrast to the 1-buyer case, the new equilibrium is not a Pareto improvement across buyer types.

In fact, it turns out that it gives *exactly the same* ex ante utility for the buyers. To compute the expected buyer surplus in the case where $v_A \in [0, \frac{1}{2}]$ and $v_B \in [\frac{1}{2}, 1]$, we break our computation up into 3 disjoint sub-cases:

- A has value less than $1/4$ so B wins with utility $v_B - \phi_B^{-1}(0)$.
- A has value greater than $1/4$, but $\phi_B(v_B) \geq \phi_A(v_A)$, so B still wins with utility $v_B - \phi_B^{-1}(\phi_A(v_A))$.
- $\phi_B(v_B) < \phi_A(v_A)$, so A wins with utility $v_A - \phi_A^{-1}(\phi_B(v_B))$.

We compute the expected buyer surplus by computing the corresponding 3 integrals with Mathematica. The expected buyer surplus from the case where $v_A \in [0, \frac{1}{2}]$ and $v_B \in [\frac{1}{2}, 1]$ is

$$\begin{aligned} \mathbb{E}[U] &= \int_0^{1/4} \int_{1/2}^1 (v_B - \phi_B^{-1}(0)) dv_B dv_A \\ &\quad + \int_{1/4}^{1/2} \int_{v_A+1/4}^1 (v_B - \phi_B^{-1}(\phi_A(v_A))) dv_B dv_A \\ &\quad + \int_{1/4}^{1/2} \int_{1/2}^{v_A+1/4} (v_A - \phi_A^{-1}(\phi_B(v_B))) dv_B dv_A = \frac{5}{96}. \end{aligned}$$

The case where $v_B \in [0, \frac{1}{2}]$ and $v_A \in [\frac{1}{2}, 1]$ is symmetric, so the expected buyer surplus coming this case is also $\frac{5}{96}$.

Adding up the expected buyer surplus in each of the four cases, we have that the overall expected buyer surplus is

$$\frac{1}{4} \cdot \frac{1}{12} + \frac{1}{4} \cdot \frac{1}{6} + 2 \cdot \frac{5}{96} = \frac{1}{48} + \frac{2}{48} + \frac{5}{48} = \frac{8}{48} = \frac{1}{6}.$$

Thus, we have the same expected buyer surplus as in the no-disclosure equilibrium. All we have accomplished is a redistribution of welfare from higher types to lower types, and we have introduced new inefficiencies.

4.2 The search for better equilibria

In an effort to better understand the equilibria of this disclosure game when $v_A, v_B \sim U[0, 1]$, we conducted a computer search over the space of partitional equilibria $\mathcal{P}_A, \mathcal{P}_B$ which partition $[0, 1]$ into (a reasonable number of) intervals. We write partitions with overlapping endpoints for convenience; this has no bearing on the computations.

We implemented the search by identifying an efficient means for computing the expected payoff for buyer A under the Myerson auction, when the messages P_A and P_B have informed the seller that $v_A \in [a, b]$ and $v_B \in [c, d]$ for arbitrary intervals $[a, b], [c, d] \subseteq [0, 1]$. This amounts to an automation of the calculations described in Section 4.1.

This proceeded starts from the observation that, for given v_B , the expectation of the payment p_A charged to A is

$$\mathbb{E}_{v_A \sim U[a, b]}[p_A | v_B] = \mathbb{E}_{v_A}[p_A] = \mathbb{E}_{v_A}[x(v_A, v_B) \cdot \phi_A(v_A)],$$

and so the expected payoff u_A of A over $v_A \sim U[a, b]$ and $v_B \sim U[c, d]$ is

$$\mathbb{E}_{v_B \sim U[c, d]} \mathbb{E}_{v_A \sim U[a, b]}[u_A] = \mathbb{E}_{v_B \sim U[c, d]} \mathbb{E}_{v_A \sim U[a, b]}[x(v_A, v_B) \cdot (v_A - \phi_A(v_A))].$$

Applying Theorem 4.1 to $v_A \sim U[a, b]$ and simplifying then lets us write the expected payoff of buyer A as

$$\mathbb{E}_{v_B \sim U[c, d]} \mathbb{E}_{v_A \sim U[a, b]}[u_A] = \frac{1}{(b-a)(d-c)} \int_R b - v_A \, dv_A \, dv_B,$$

where R is the region of (v_A, v_B) given by the inequalities

$$a \leq v_A \leq b, \quad c \leq v_B \leq d, \quad v_A \geq b/2, \quad v_A \geq v_B + \frac{b-d}{2}.$$

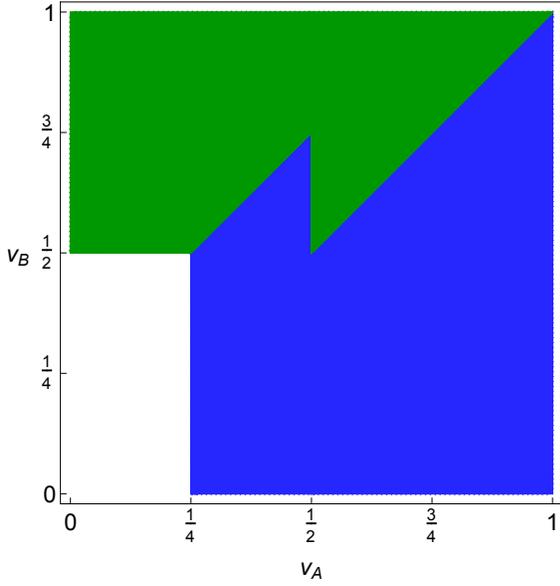
Finally, the buyer surplus is just $U = u_A + u_B$, so for a given region it is simply

$$\mathbb{E}_{v_B \sim U[c, d]} \mathbb{E}_{v_A \sim U[a, b]}[U] = \mathbb{E}_{v_B \sim U[c, d]} \mathbb{E}_{v_A \sim U[a, b]}[u_A] + \mathbb{E}_{v_B \sim U[c, d]} \mathbb{E}_{v_A \sim U[a, b]}[u_B].$$

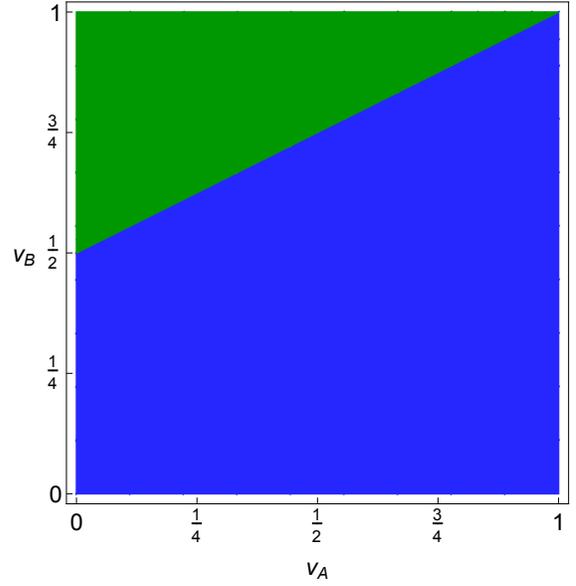
Once the boundaries of R are determined, this integral can be computed symbolically. Therefore computing the expected buyer surplus for a given pair of partitions $\mathcal{P}_A, \mathcal{P}_B$ simply entails evaluating the above expression for each pair of possible intervals (P_A, P_B) .

Using this approach, our search over pairs of partitions into intervals (with representative allocations shown in Figure 3) strongly suggests the following conjecture:

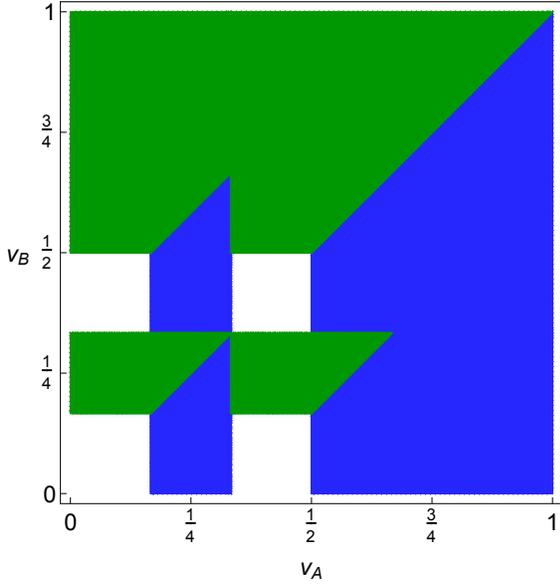
Conjecture 4.2. *For $v_A, v_B \sim U[0, 1]$ the partitional equilibrium given by $\mathcal{P}_A = \{[0, \frac{1}{2}], [\frac{1}{2}, 1]\}$ and $\mathcal{P}_B = \{[0, 1]\}$ maximizes expected buyer welfare for the class of partitional equilibria.*



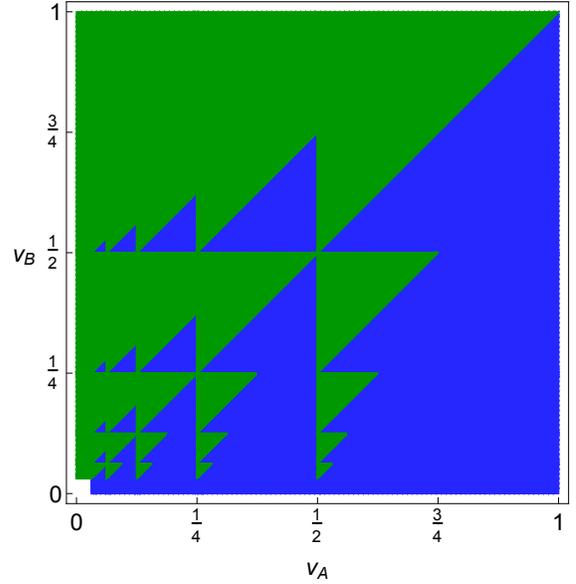
(a) Coarse asymmetric disclosure: $\mathcal{P}_A = \{[0, \frac{1}{2}], [\frac{1}{2}, 1]\}$, and $\mathcal{P}_B = \{[0, 1]\}$, and $\mathbb{E}[U] = \frac{11}{64} > \frac{1}{6}$.



(b) Maximally asymmetric disclosure: $\mathcal{P}_A = \{x : x \in [0, 1]\}$ and $\mathcal{P}_B = \{[0, 1]\}$, yielding $\mathbb{E}[U] = \frac{1}{24}$.



(c) A and B disclose according to $\mathcal{P}_A = \mathcal{P}_B = \{[0, \frac{1}{3}], [\frac{1}{3}, \frac{2}{3}], [\frac{2}{3}, 1]\}$, $\mathbb{E}[U] = \frac{1}{6}$.



(d) A and B disclose according to Zeno's partition \mathcal{P}_Z . Here $\mathbb{E}[U] = \frac{23}{147} < \frac{1}{6}$.

Figure 3: Allocations of the good in different partitional equilibria. As in Figure 2, blue denotes that A receives the good, green denotes that B receives the good; in this case unshaded denotes no sale. Expected buyer utility $\mathbb{E}[U]$ is given for each case.

This disclosure profile (with allocations shown in Figure 3a) is notable in that it is asymmetric, and in that it appears to strictly outperform all symmetric profiles in terms of expected buyer surplus. But most surprisingly, the buyer who discloses (in this case, buyer A) has increased expected utility, while the buyer who does not disclose suffers a strict decrease in expected utility as compared to the symmetric no-disclosure equilibrium (Figure 2a); A receives expected utility $\frac{13}{128} \approx 0.102$, while B receives $\frac{9}{128} \approx 0.070$. The total expected surplus is $\frac{11}{64} > \frac{1}{6}$.

At the same time, in this asymmetric regime there is a tradeoff between disclosure detail and likelihood of receiving the good on the one hand, and expected surplus on the other. But as A discloses more information and B does not disclose, A 's expected utility and relative surplus decreases. In the limit, where A discloses v_A exactly and B discloses nothing, the allocation is shown in Figure 3b. Here $\mathbb{E}[U_A] = 0$, while $\mathbb{E}[U_B] = 1/24$.

Experimental evidence also suggests that:

Conjecture 4.3. *For $v_A, v_B \sim U[0, 1]$ nondisclosure maximizes expected buyer welfare for the class of all symmetric message sets which are interval partitions.*

This no-disclosure equilibrium yields an expected buyer surplus of $\frac{1}{6}$, the same as the equilibrium analyzed in Section 4.1. In fact, this has the same expected buyer surplus as the symmetric profile shown in Figure 3c, and indeed it appears that *all* symmetric buyer disclosure profiles with a high/low threshold $t \in [0, \frac{1}{2}]$ realize an expected buyer welfare of $\frac{1}{6}$.

By contrast, the equilibrium with symmetric disclosure given by the Zeno partition (Figure 3d) confers expected buyer utility $\frac{23}{147} < \frac{1}{6}$.

4.3 Efficiency versus buyer welfare

If both buyers fully disclose their types, then the seller will sell to the buyer with a higher value at that value, so we will have a socially efficient outcome, but neither buyer will receive any utility. On the other hand, in all of the partial disclosure equilibria considered thus far, including no-disclosure, some buyer gets nonzero utility, but the outcome is not guaranteed to be efficient (the good may be sold to the buyer with a lower value, or not sold at all). A natural question is whether it is possible to simultaneously achieve efficiency and nonzero buyer surplus. If we restrict attention to connected partitional equilibria, the answer is negative.

Theorem 4.4. *For $v_A, v_B \sim U[0, 1]$, there are no partitions $\mathcal{P}_A, \mathcal{P}_B$ of $[0, 1]$ into intervals inducing an efficient equilibrium that confers nonzero expected buyer surplus.*

Proof. We assume that the equilibrium induced by \mathcal{P}_A and \mathcal{P}_B confers nonzero buyer surplus, and show that it must be inefficient for some $v_A, v_B \in [0, 1]$. The only property we actually need to assume is that there exists some interval $P_A \in \mathcal{P}_A$ of measure greater than zero. This follows without loss of generality, for if all intervals in both partitions have measure zero, i.e., are singletons, then both buyers always exactly disclose their values, and hence get zero utility. Let $0 \leq a < b \leq 1$ be such that the closure of P_A is $\overline{P_A} = [a, b]$. Let

$$x := \frac{\max\{0, 2a - b\} + a}{2}.$$

Note that x is the average of a and some point weakly less than a , so $x \leq a$ (and $x = 0$ precisely when $a = 0$; otherwise x is in the open interval $(0, a)$). Let P_B be the unique element of \mathcal{P}_B containing x , and let $0 \leq c \leq d \leq 1$ be such that $\overline{P_B} = [c, d]$. There are five cases to consider.

Case 1: $x = 0$. In this case, it must be that $a = c = 0$. Let $v_A := \frac{b}{4} \in P_A$ and $v_B := 0$. After disclosing that their type is in $[a, b]$, by Theorem 4.1, the virtual value of buyer A is

$$\phi_A = 2v_A - b = -\frac{b}{2} < 0,$$

so they will not get the good. This is inefficient since $v_A > v_B = 0$.

Case 2: $x > 0$ and $d \leq a$. For small $\varepsilon > 0$, let $v_A := a + \varepsilon \in P_A$ and $v_B := \frac{x+d}{2} \in P_B$. Then by Theorem 4.1 the virtual value of buyer A is

$$\phi_A = 2v_A - b = 2a - b + 2\varepsilon$$

and the virtual value of buyer B is

$$\phi_B = 2v_B - d = x \geq \frac{3a - b}{2} > \frac{4a - 2b}{2} = 2a - b = \phi_A - 2\varepsilon.$$

Therefore, for sufficiently small ε , buyer B will get the good, which is inefficient since

$$v_B = \frac{x + d}{2} \leq \frac{x + a}{2} \leq \frac{a + a}{2} = a = v_A - \varepsilon < v_A.$$

Case 3: $a < d < b$. For small $\varepsilon \geq 0$, let $v_A := d + \frac{b-d}{4} \in P_A$ and $v_B := d - \varepsilon \in P_B$. (Note that the ε is needed to ensure we have $v_B \in P_B$. Choosing $\varepsilon > 0$ is necessary if $d \notin P_B$; choosing $\varepsilon = 0$ is necessary if $P_B = \{c\} = \{d\}$.) Then the virtual value of buyer A is

$$\phi_A = 2v_A - b = \frac{3d - b}{2}$$

and the virtual value of buyer B is

$$\phi_B = 2v_B - d = d - 2\varepsilon = \frac{2d}{2} - 2\varepsilon > \frac{3d - b}{2} - 2\varepsilon = \phi_A - 2\varepsilon.$$

Therefore, for sufficiently small ε , buyer B will get the good, which is inefficient since

$$v_B = d - \varepsilon < d + \frac{b - d}{4} = v_A.$$

Case 4: $d > b$. Observe that $c \leq x \leq a < b$. Thus, we have $c < b < d$, so this is the same as Case 3, with the two roles of the two buyers reversed.

Case 5: $x > 0$ and $d = b$. Let x' be any point in the open interval (x, a) . Let P'_A be the unique element of \mathcal{P}_A containing x' , and let $0 \leq a' \leq b' \leq a$ be such that $\overline{P'_A} = [a', b']$. Observe that $c \leq x < x' \leq b' \leq a < b = d$. Since $c < b' < d$, we are again in Case 3, with the roles of the two buyers reversed, and substituting P'_A for P_A . \square

5 Complexity

In this section, we study the complexity of finding a buyer-optimal pure-strategy equilibrium over a discrete distribution. In this setting, a buyer will have possible types $N = \{1, \dots, n\}$, each with value $v_1 < \dots < v_n$ and occurring with probability p_1, \dots, p_n respectively.

First, we analyze the restricted setting introduced by [1], in which buyer messages must be connected, and thus the search problem is restricted to connected partitions. Note that in Example

1 of [1], they show that their greedy algorithm is unable to compute the buyer-optimal disclosure strategy in this setting. Using dynamic programming, we give a polynomial time algorithm to solve this problem.

Next, we ask whether the buyer could do even better if they were allowed to report from a broader class of messages. Specifically, we turn our attention to a more general setting in which a buyer may report any arbitrary subset of their values $\{v_1, \dots, v_n\}$ (or, equivalently, an arbitrary subset of N). Here, things are not quite as easy. As it turns out, it is weakly NP-hard to compute the buyer's optimal strategy. Further, the optimal connected strategy does not even provide a constant factor approximation to the optimal buyer surplus.

5.1 Connected partitions

Algorithm 1: Optimizes buyer welfare over connected partitions.

Input: Values $v_1 < \dots < v_n$ and probabilities p_1, \dots, p_n

Output: A connected partition inducing an equilibrium with optimal ex ante buyer utility

```

1 Initialize array  $Q[0 \dots n]$ ;
2  $Q[0].\text{utility} \leftarrow 0$ ;
3  $Q[0].\text{partition} \leftarrow \emptyset$ ;
4 for  $i \leftarrow 1 \dots n$  do
5    $j^* = \operatorname{argmax}_{j \in \{0, \dots, i-1\}} \{Q[j].\text{utility} + \text{buyerUtility}(\{j+1, \dots, i\})\}$ ;
6    $Q[i].\text{partition} = Q[j^*].\text{partition} \cup \{j+1, \dots, i\}$ ;
7    $Q[i].\text{utility} = Q[j^*].\text{utility} + \text{buyerUtility}(\{j+1, \dots, i\})$ ;
8 end
9 return  $Q[n].\text{partition}$ ;
```

Let us first consider the case of connected messages. By the Partitional Lemma, it suffices to restrict attention to equilibria induced by some partition \mathcal{P} of the set of buyer types $N = \{1, \dots, n\}$. We show the following.

Theorem 5.1. *Algorithm 1 computes a connected partition that induces an equilibrium which maximizes ex ante buyer utility in polynomial time.*

Proof. First we will prove the correctness of Algorithm 1. It uses a dynamic programming approach with array Q . Each $Q[i]$ will store the best partition of $\{1, \dots, i\}$ as well as the expected buyer utility when using this partition, scaled down by $\sum_{j=1}^i p_j$. Note that when we are considering the next element i , conditioned on it being in some partition element $\{j+1, \dots, i\}$, the remaining elements of the partition will be the optimal partition of $\{1, \dots, j\}$. Hence, by choosing the best possible j , that will allow us to find the optimal partition for $Q[i]$. The `buyerUtility` function, given a possible message of the buyer, gives the highest utility the buyer can receive when the seller best responds to this message.

To show the runtime is polynomial, note that the initialization can be done in $O(n)$ time. The for loop is runs for n iterations. In each iteration, note that `buyerUtility` can be computed by checking each seller price in $\{v_j, \dots, v_i\}$, of which there are at most n . For each of these, we must check the probability of sale and expected buyer utility. This can be done in time linear in n by iterating through the values in order, and keeping track of the total probability the buyer has value above this value. In total, this loop takes $O(n^3)$ time. \square

5.2 Arbitrary partitions

While message connectivity might reflect realistic practical constraints on the set of feasible messages, it is nevertheless interesting to consider settings where this constraint is not present. Unfortunately, it turns out that, in the general case, computing a pure-strategy equilibrium that maximizes ex ante buyer utility is NP-hard. More formally, we consider the following decision problem.

BUYER-OPT: *Given a sequence of probabilities p_1, p_2, \dots, p_n (where $\sum_{i \in [n]} p_i = 1$), corresponding distinct positive valuations v_1, v_2, \dots, v_n , and a positive number U , determine whether there exists a pure-strategy equilibrium with expected buyer utility U in the disclosure game with one buyer, one seller, and one good, where the buyer has value v_i for the good with probability p_i .*

Theorem 5.2. *BUYER-OPT is weakly NP-complete.*

We show BUYER-OPT is NP-hard by reducing from the following problem.

PARTITION: *Given a finite sequence of positive integers s_1, s_2, \dots, s_m , adding up to some even positive integer S , determine whether there exists $I \subseteq [m]$ such that*

$$\sum_{i \in I} s_i = \frac{S}{2}.$$

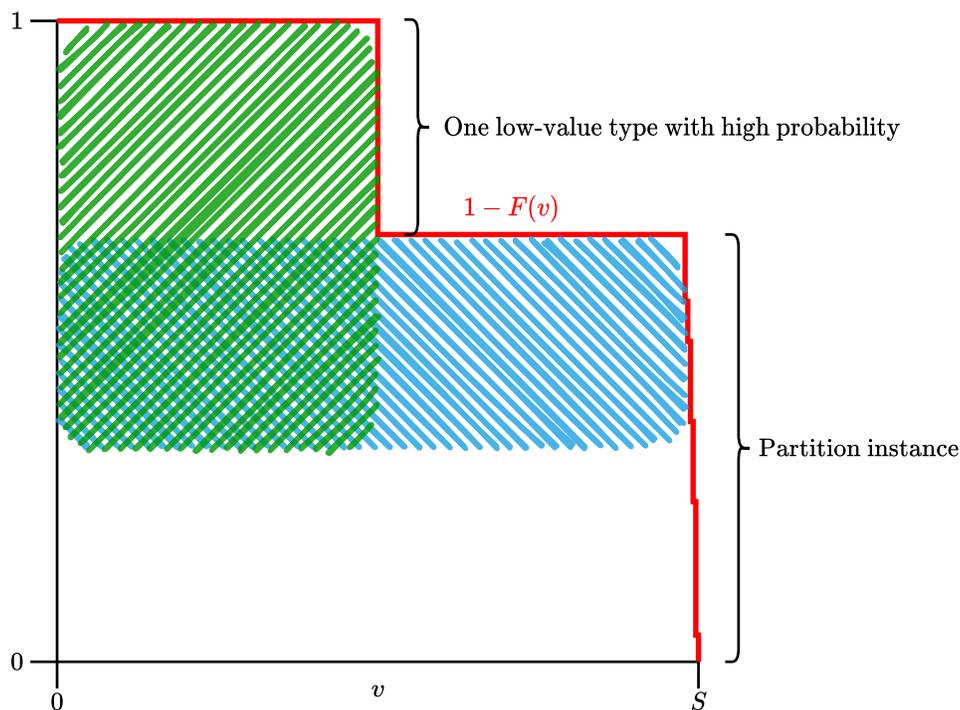


Figure 4: The reduction from PARTITION to BUYER-OPT.

The main idea is to behind the reduction is that there will be one low-value type with high probability and many high value types with low probability, as shown in Figure 4. The zigzagging red line shows the probability of sale given a posted price according to the instance defined by the

reduction. Ideally, to maximize buyer welfare, we would like to pool some of the high-value types together with the low-value type, so that all types in the pool send the same message, and the seller will set their price equal to the value of the low-value type, giving the high-value types positive utility. The seller maximizes revenue by choosing the largest rectangle under the curve, ignoring the types that are not in this pool. Thus, in order to incentivize the seller to choose the low price, we need the vertical green rectangle in Figure 4 to be larger than the horizontal blue rectangle, which is only possible if we pool at most half of the high-value types (weighted by probability mass) with the low-value type. Optimally, we would like to get exactly half of the probability mass, which requires solving PARTITION. See Appendix A for the full proof.

We now ask whether our polynomial time algorithm from Section 5.1 could be used to at least approximate optimal ex ante buyer utility in this more general setting. Unfortunately, we show here that this is not the case.

Proposition 5.3. *No equilibrium for the general setting induced by a connected partition can have ex ante buyer utility that approximates the optimal within a constant factor.*

See Appendix B for the proof.

6 Conclusion

Can buyers reveal for a better deal, and if so, how? In the restricted setting with one buyer and one good, Ali, Lewis, and Vasserman [1] answer the former question with an emphatic “Yes!” Utility improving disclosure schemes are generally feasible, always socially efficient, and in the presence of connectivity constraints on the disclosure technology, can easily be exactly optimized for the buyer. Unfortunately, it appears that these desirable properties are extremely non-robust.

We believe that our negative results are not mere peculiarities of the model, but are instead due to fundamental market forces. With either multiple goods or multiple buyers, the buyer(s) benefit substantially from the uncertainty the seller faces. While, at the surface, it still seems plausible that a buyer should only gain utility from credibly revealing to the seller that they would not be willing to pay the seller’s optimal reserve price, such disclosures may have secondary effects, dramatically reducing the seller’s uncertainty to the point where the buyer(s) are ultimately harmed in expectation.

We do note that, despite all of this negative evidence on the benefit of disclosure, there are in fact instances where it can provably help.³ The main point of this paper is not that buyer welfare and social welfare are always incompatible, or that the ability of buyers to disclose information about their type is useless. Rather, instances with multiple buyers where disclosure is useful seem to be the exception, rather than the norm.

There are still many dimensions to the voluntary disclosure model that have yet to be explored. Most notable is the possibility of mixed strategies. A tantalizing open question, left unaddressed

³As an example, consider the following alternative prior distribution in which buyers have i.i.d. values, symmetric strategies, and both Pareto improve in terms of interim utility by disclosing some information. The value distribution will be discrete; with probability $\frac{99}{100}$ they will have value 1000 and with probability $\frac{1}{200}$ each, they will have value 1 and 2. With no disclosure, the revenue-optimal mechanism run by the seller will only sell at a price of 1000. This means that buyers never receive positive utility as they cannot purchase the good for a price strictly less than their value. On the other hand, suppose buyers disclose whether they are of “low” type with value in $\{1, 2\}$ or “high” type with value 1000. In the scenario where both buyers reveal a low type, from the seller’s perspective, each buyer has value 1 or 2 with probability $\frac{1}{2}$, so the optimal selling mechanism is a second price auction with reserve price 1. Hence, in the case where one buyer has value 2 and the other has value 1, the former will receive positive utility, strictly improving on the no-disclosure setting.

by this as well as prior work, is whether there exist mixed-strategy equilibria in which the buyer(s) obtain strictly higher ex ante utility than in any pure-strategy equilibrium—one could imagine a scenario where some buyer types are indifferent between multiple messages, some seller types are indifferent between multiple prices, and somehow the complex belief distributions generated yields a higher expected buyer surplus than would be possible with the simpler belief distributions generated by pure messaging strategies. Even with the restriction that there is one buyer and one good, we do not see an obvious way to rule out this possibility (even for simple distributions like $U[0, 1]$), and while such a scenario may sound absurd in the one-buyer, one-good setting, it seems entirely plausible with multiple buyers or goods.

An orthogonal extension of the model would be to allow the buyers to *privately* disclose information to the seller, without other buyers observing the value. While possibly more applicable in some scenarios, this model appears to be less tractable, as the seller’s revenue maximization problem cannot be solved independently for every possible tuple of buyer messages. Since the buyers face uncertainty over which of these subgames they are in, standard tools from auction theory do not apply; e.g., even for the setting from Section 4 of two uniform $[0, 1]$ buyers sending connected messages, a Myerson auction is no longer guaranteed to be optimal or incentive compatible.

A final direction for future work concerns the complexity of the BUYER-OPT decision problem. Our reduction is from PARTITION, which is weakly NP-Hard, and it is clear that the optimal values of the instances produced by the reduction can be approximated to arbitrary accuracy in polynomial time. This leaves open the possibility of a pseudo-polynomial time algorithm for solving this problem, or potentially some polynomial time approximation scheme.

References

- [1] S Nageeb Ali, Greg Lewis, and Shoshana Vasserman. Voluntary disclosure and personalized pricing. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 537–538, 2020.
- [2] Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. The limits of an information intermediary in auction design. *CoRR*, abs/2009.11841, 2020.
- [3] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.
- [4] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.
- [5] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2021. URL <https://www.gurobi.com>.
- [6] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.
- [7] Weiran Shen, Pingzhong Tang, and Yulong Zeng. Buyer signaling games in auctions. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1591–1599, 2019.
- [8] Itai Sher and Rakesh Vohra. Price discrimination through communication. *Theoretical Economics*, 10(2):597–648, 2015.

Appendix

A Proof of Theorem 5.2

To see that BUYER-OPT is in NP, observe that the Partitional Lemma implies there is a buyer-optimal equilibrium induced by a partition of N . Given such a partition \mathcal{P} , for every $P \in \mathcal{P}$, there are at most $|P|$ prices the seller may choose to set (a price not equal to v_i for some $i \in P$ cannot be optimal). Therefore, we can efficiently compute the optimal selling mechanism and its expected buyer utility, then verify that the utility is at least U .

To prove hardness, we reduce from PARTITION, which is known to be weakly NP-hard [6]. In terms of the notation defining BUYER-OPT and PARTITION, our reduction is defined as follows:

$$\begin{aligned}
 n &:= m + 1 \\
 p_i &:= \frac{2s_i}{3S} && \text{for } i \leq m \\
 v_i &:= S - \frac{1}{4+i} && \text{for } i \leq m \\
 p_n &:= \frac{1}{3} \\
 v_n &:= \frac{S}{2} \\
 U &:= \frac{S}{6} - \frac{1}{12}
 \end{aligned}$$

For the forward direction, suppose that I is a solution to the given PARTITION instance. Then consider the partitional equilibrium induced by the partition

$$\{I \cup \{n\}, [m] \setminus I\}.$$

To compute the buyer surplus, there are three cases to consider:

1. The buyer's type is n .
2. The buyer's type is in I .
3. The buyer's type is in $[m] \setminus I$.

Observe that the first case occurs with probability $\frac{1}{3}$, and the second case occurs with probability

$$\sum_{i \in I} p_i = \sum_{i \in I} \frac{2s_i}{3S} = \frac{2 \sum_{i \in I} s_i}{3S} = \frac{2 \cdot \frac{S}{2}}{3S} = \frac{1}{3},$$

so, in fact, all three cases occur with probability $\frac{1}{3}$. Consider the utility of the buyer in case (2), in which the buyer reveals that their type is in $I \cup \{n\}$, so the seller knows they are in case (1) or (2). Given this information, with probability $\frac{1}{2}$ we are in case (1), so the buyer's value is $\frac{S}{2}$, and with probability $\frac{1}{2}$, we are in case (2), so the buyer's value is in the range $(S - \frac{1}{4}, S)$. If the seller sets a price in the range $(\frac{S}{2}, S)$, they will exclude type n and only make a sale with probability $\frac{1}{2}$, so their expected revenue is strictly less than $\frac{S}{2}$. If the seller sets a price of S or greater, they will definitely not make a sale, so their expected revenue is zero. In either case, their expected revenue is less than $\frac{S}{2}$, which is what they could get by setting a price of $\frac{S}{2}$ and guaranteeing sale. Therefore, the

seller's optimal price is at most $\frac{S}{2}$. Since we are in case (2), the buyer has value at least $S - \frac{1}{4}$, so the expected buyer surplus will be at least

$$S - \frac{1}{4} - \frac{S}{2} = \frac{S}{2} - \frac{1}{4}.$$

Since case (2) occurs with probability $\frac{1}{3}$, the overall expected buyer surplus is at least

$$\frac{1}{3} \left(\frac{S}{2} - \frac{1}{4} \right) = U.$$

For the backward direction, suppose there is a pure-strategy equilibrium giving the buyer ex ante expected utility at least U . By the Partitional Lemma, it is without loss of generality to assume the equilibrium is induced by some partition \mathcal{P} of $[n]$. Let $I \subseteq [m]$ be the unique set of types such that $I \cup \{n\} \in \mathcal{P}$. We claim that I is a solution to the PARTITION instance, i.e.,

$$\sum_{i \in I} s_i = \frac{S}{2}.$$

For any $i \in [n]$, let r_i be the price the seller charges type i at equilibrium. Then

$$\begin{aligned} \frac{S}{6} - \frac{1}{12} &= U \\ &\leq \sum_{i \in [n]} p_i(v_i - r_i) \quad (\text{by the Efficiency Lemma and the definition of } U) \\ &= \sum_{P \in \mathcal{P} \setminus \{I \cup \{n\}\}} \sum_{i \in P} p_i(v_i - r_i) + \sum_{i \in I} p_i(v_i - r_i) + p_n(v_n - r_n) \\ &= \sum_{P \in \mathcal{P} \setminus \{I \cup \{n\}\}} \sum_{i \in P} p_i \left(v_i - \min_{j \in P} v_j \right) + \sum_{i \in I} p_i(v_i - v_n) + p_n(v_n - v_n) \\ &\quad (\text{by the Efficiency Lemma}) \\ &< \sum_{P \in \mathcal{P} \setminus \{I \cup \{n\}\}} \sum_{i \in P} p_i \left(S - \left(S - \frac{1}{4} \right) \right) + \sum_{i \in I} p_i(S - v_n) \\ &\quad (\text{since, for all } i \in [m], v_i \in (S - \frac{1}{4}, S)) \\ &= \sum_{P \in \mathcal{P} \setminus \{I \cup \{n\}\}} \sum_{i \in P} p_i \left(\frac{1}{4} \right) + \sum_{i \in I} p_i \left(S - \frac{S}{2} \right) \\ &= \frac{1}{4} \sum_{P \in \mathcal{P} \setminus \{I \cup \{n\}\}} \sum_{i \in P} p_i + \sum_{i \in I} \frac{2s_i}{3S} \cdot \frac{S}{2} \\ &= \frac{1}{4} \sum_{P \in \mathcal{P} \setminus \{I \cup \{n\}\}} \sum_{i \in P} p_i + \frac{1}{3} \sum_{i \in I} s_i \\ &\leq \frac{1}{4} \cdot \frac{2}{3} + \frac{1}{3} \sum_{i \in I} s_i \\ &= \frac{1}{6} + \frac{1}{3} \sum_{i \in I} s_i \end{aligned}$$

Multiplying by 3 and rearranging, we have

$$\sum_{i \in I} s_i > \frac{S}{2} - \frac{3}{4}.$$

Since $\sum_{i \in I} s_i$ and $\frac{S}{2}$ are integers, it follows that

$$\sum_{i \in I} s_i \geq \frac{S}{2}.$$

To prove the opposite inequality, suppose that the buyer reveals that their type is in $I \cup \{n\}$. If the seller sets a price of $S - \frac{1}{4}$, they exclude type n , and thus sell with probability

$$\frac{\sum_{i \in I} p_i}{\frac{1}{3} + \sum_{i \in I} p_i},$$

so their expected revenue is

$$\left(S - \frac{1}{4}\right) \left(\frac{\sum_{i \in I} p_i}{\frac{1}{3} + \sum_{i \in I} p_i}\right).$$

Since the Efficiency Lemma implies that the seller instead chooses to sell at a price of $\frac{S}{2}$ for an expected revenue of $\frac{S}{2}$, we must have that

$$\left(S - \frac{1}{4}\right) \left(\frac{\sum_{i \in I} p_i}{\frac{1}{3} + \sum_{i \in I} p_i}\right) \leq \frac{S}{2},$$

or, equivalently,

$$\left(S - \frac{1}{4}\right) \sum_{i \in I} p_i \leq \frac{S}{2} \left(\frac{1}{3} + \sum_{i \in I} p_i\right).$$

Rearranging again, we have

$$\sum_{i \in I} p_i \leq \frac{\frac{S}{2} \cdot \frac{1}{3}}{S - \frac{1}{4} - \frac{S}{2}} = \frac{\frac{S}{6}}{\frac{S}{2} - \frac{1}{4}} = \frac{1}{6S - 3} + \frac{1}{3}.$$

Plugging in the definition of p_i , this becomes

$$\frac{2}{3S} \sum_{i \in I} s_i \leq \frac{1}{6S - 3} + \frac{1}{3},$$

which simplifies to

$$\sum_{i \in I} s_i \leq \frac{3S}{2(6S - 3)} + \frac{3S}{6} = \frac{S}{4S - 1} + \frac{S}{2}$$

Since $\sum_{i \in I} s_i$ and $\frac{S}{2}$ are integers and $\frac{S}{4S - 1} < 1$ for any positive integer S , it follows that

$$\sum_{i \in I} s_i \leq \frac{S}{2}.$$

Thus, we have equality, so I is a solution to the PARTITION instance. □

B Proof of Theorem 5.3

Consider a setting in which a buyer has possible values $1, 2, 2 + \delta$ for some $0 < \delta < 1$ each with probabilities $\frac{1}{3}, \frac{5}{9}, \frac{1}{9}$ respectively. We will show that every equilibrium induced by a connected partition gives the buyer expected utility at most $\frac{\delta}{9}$ while an arbitrary partition can give expected utility at least $\frac{1+\delta}{9}$. This implies no approximation ratio better than $\frac{1+\delta}{\delta}$ is achievable. By letting δ approach 0, this lower bound grows arbitrarily large.

Let us first show that when a buyer is allowed to report arbitrary subsets, there are equilibria in which the buyer ex ante utility is at least $\frac{1+\delta}{9}$. Consider the equilibrium induced by the partition $\{\{1, 2 + \delta\}, \{2\}\}$. Now, when the buyer reports $\{1, 2 + \delta\}$, the seller will prefer to set a price of 1, which will result in expected revenue 1, over $2 + \delta$, which (after updating their posterior,) will only sell with probability $\frac{1}{4}$, resulting in expected revenue $\frac{2+\delta}{4} < 1$. Hence, by setting a price of 1, the buyer receives utility $1 + \delta$ when they have value $2 + \delta$, so their ex ante utility is at least $\frac{1+\delta}{9}$.

Next, we will upper bound the ex ante buyer utility of a connected partition strategy. First, note that a seller will always choose a posted price $p \in \{1, 2, 2 + \delta\}$. This implies that whenever a buyer has value 1, they cannot receive any utility, as the price will always be at least 1. Further, we will show that when a buyer has value 2, the seller will always set a posted price of at least 2. Indeed, regardless of what the buyer reports, since messages must contain the buyer's true value, the seller knows that 2 is a possibility. Hence, if the buyer posts a price of 1, they will get expected revenue 1, while if they post a price of 2, they will get expected revenue at least $\frac{5}{9} \cdot 2 > 1$. This implies in equilibrium they will never set a price of 1, so they buyer cannot receive any utility. Finally, let us consider what happens when the buyer has true utility $2 + \delta$. In this case, if they report $\{2 + \delta\}$, the seller will set the price to $2 + \delta$ and they will receive zero utility. On the other hand, any other report will include 2 in the message, and by previous argument, the seller will not set a price of 1, so the most utility they can receive is $2 + \delta - 2 = \delta$. Hence, the buyer's ex ante utility is at most $\frac{\delta}{9}$. \square