

Can Buyers Reveal for a Better Deal?

Daniel Halpern, Gregory Kehne, Jamie Tucker-Foltz*

Harvard University

{dhalpern, gkehne}@g.harvard.edu, jtuckerfoltz@gmail.com

Abstract

We study 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 of one buyer and one good, showing that such communication can simultaneously improve social welfare and ex ante buyer utility. However, 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 surprising in contrast with 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 well known that Bayesian-optimal mechanisms for revenue maximization may lead to inefficient outcomes. A seller may rationally refuse to sell to buyers 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 possible 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. The increasing accessibility and quality of buyer data have made 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 [Bergemann *et al.*, 2015; Ali *et al.*, 2020]. 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 recent work, Ali, Lewis, and Vasserman [2020] 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. When there is one buyer, one seller, and one good, they demonstrate that there always exists a disclosure strategy for the buyer such that

- the buyer has no 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

*Contact Author

with a *partitional equilibrium*, in which there is some partition \mathcal{P} of the types and every buyer reveals the set in \mathcal{P} to which their type belongs.

Ali *et al.* [2020] 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 sellers 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 we depart.

In Section 3 we investigate the effects of disclosure when there are *two* i.i.d., uniform $[0, 1]$ buyers instead of one. Surprisingly, we find that the natural, buyer-symmetric analogues of the optimal one-buyer equilibria from Ali *et al.* [2020] no longer yield buyer welfare improvements. Perhaps even more surprisingly, it is possible to improve the expected buyer surplus (the sum of the 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 [Ali *et al.*, 2020], we show an extreme impossibility result (Theorem 3.3): in any equilibrium where the good is always allocated to the highest bidder, both buyers must always receive utility zero. Since the model assumes the seller has no cost to sell the goods, this shows that social efficiency is incompatible with maximizing buyer welfare, which lies in stark contrast to the one-buyer, one-good case.

In Section 4 we further generalize these impossibility results to settings with richer disclosure capabilities and arbitrary priors. 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 4.1). This holds even with the restrictions that buyer valuations are additive and independent across goods, as well as independent across buyers.

Finally, in Section 5 we study the problem of maximizing consumer welfare through disclosure schemes from a computational perspective. 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 [Ali *et al.*, 2020] where disclosure messages must be “connected” (Theorem 5.1), the more general problem is (weakly) NP-hard (Theorem 5.2), and is inapproximable by connected equilibria (Proposition 5.3).

1.1 Related Work

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 [2017] on *Bayesian persuasion*. On a more conceptual level, it contributes to a growing literature on the connection between privacy and information in markets. See Acquisti, Taylor, and Wagman [2016] and Bergemann and Bonatti [2019] for relevant surveys.

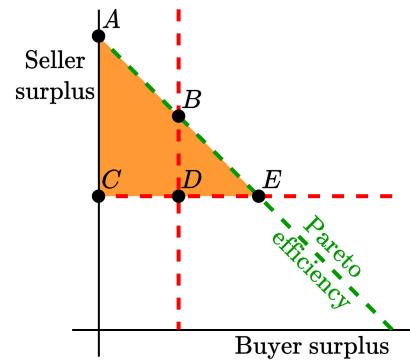


Figure 1: Feasible welfare outcomes from Bergemann *et al.* [2015]

The signals which we consider are the *verifiable disclosures* of [Ali *et al.*, 2020], in which buyers send public signals about their types from a set of possible signals which are demonstrably truthful. The study of verifiable disclosure dates back to classic economic works by Grossman [1981] and Milgrom [1981], and has recently seen a resurgence in the economics literature, both in market settings and more general games [Kartik and Tercieux, 2012; Ben-Porath and Lipman, 2012; Hagenbach *et al.*, 2014; Hart *et al.*, 2017; Ben-Porath *et al.*, 2018; Ben-Porath *et al.*, 2019; Koessler and Skreta, 2019]. The original game focused on a seller being able to verifiably disclose information about the item they are selling, rather than the buyer about their type. A closely related work by Sher and Vohra [2015] considers the 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 Ali *et al.* [2020] is 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 there is some feasible point on line segment BE when there is only one buyer and one good. We show that this 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 4).

A parallel line of work aims to improve buyer welfare via an intermediary that observes buyer valuations and then sends a signal to the seller based on some prearranged signalling scheme. Bergemann, Brooks, and Morris [2015] address the case of one buyer and one seller. They show that any combination of welfares in the shaded triangle is attainable via randomized signalling schemes. Alijani, Banerjee, Munagala, and Wang [2020] 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.

Similar to our setting is that of Shen, Tang, and Zeng [2019], which generalizes [Bergemann *et al.*, 2015] 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. This setting also differs from ours in that buyer signalling schemes are required to be best responses to each other, a constraint that we do not enforce.

2 Model

As in [Ali *et al.*, 2020], 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 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; however, 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 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 misrepresent 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.¹ We 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

¹See [Myerson, 1981] for a description and discussion of virtual values which we use extensively. In this case, step 2 implicitly involves additional communication from the buyers to the seller, revealing their types as prescribed by the direct revelation mechanism. This communication is different from that of step 1 in that now buyers are allowed to lie about their types.

connected and that \mathcal{P} is a *connected partition*. As it is shown in [Ali *et al.*, 2020], 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 [Ali *et al.*, 2020]). *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 [Ali *et al.*, 2020]). *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.² While the partitional lemma continues to hold for multiple buyers and/or multiple goods, we show in Section 4 that the Efficiency Lemma does not.

3 Disclosure with 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 canonical setting 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 (Lemma 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 1,³ then A is guaranteed not to derive any utility from the resulting Myerson auction, since in order to receive the good they must clear the Myerson reserve of $\bar{v}_A \geq v_A$. 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 .

²The version of the Partitional Lemma proved in [Ali *et al.*, 2020] 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.

³This 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.

3.1 Example: The First Step of Zeno’s Partition

In [Ali *et al.*, 2020], 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. 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 situation with one buyer: 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.

The sale outcomes in the partitioned equilibrium are depicted in Figure 2b. The case where buyer A discloses “low” and buyer B discloses “high” corresponds to the top-left quadrant. The small blue triangle represents the case where the seller sells the good to A because, even though they have a lower value, they have a *higher virtual value*.

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, 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 one 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, as shown in the full version of this paper.⁴ All we have accomplished is a redistribution of welfare

from higher types to lower types, and we have introduced new inefficiencies.

3.2 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 partitioned equilibria $\mathcal{P}_A, \mathcal{P}_B$ which partition $[0, 1]$ into (a reasonable number of) intervals, described in the full version. Our computational results strongly support the following conjecture:

Conjecture 3.1. *For $v_A, v_B \sim U[0, 1]$ the partitioned 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 partitioned equilibria.*

This disclosure profile (with allocations shown in Figure 2d) 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 begins to decrease. In the limit, where A discloses v_A exactly and B discloses nothing, with $\mathbb{E}[U_A] = 0$ and $\mathbb{E}[U_B] = 1/24$.

Experimental evidence also suggests that:

Conjecture 3.2. *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 3.1. In fact, *all* symmetric buyer disclosure profiles with a high/low threshold $t \in [0, \frac{1}{2}]$ realize an expected buyer welfare of $\frac{1}{6}$. (See the full version for a proof.)

By contrast, the equilibrium with symmetric disclosure given by Zeno’s partition (Figure 2c) confers expected buyer utility $\frac{23}{147} < \frac{1}{6}$.

3.3 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. Our first main result is that, if we restrict attention to connected partitioned equilibria, the answer is negative. The proof of this theorem can be found in the full version.

⁴Full version available at: <https://arxiv.org/abs/2106.13882>

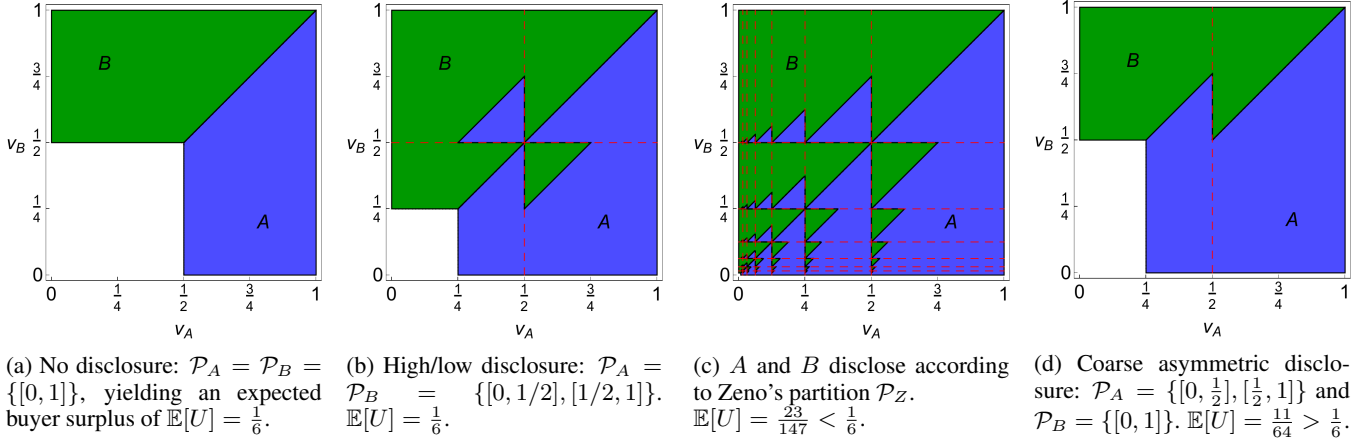


Figure 2: Allocations of the good in different partitional equilibria. Blue denotes that A receives the good, green denotes that B receives the good, and white denotes no sale. Expected buyer surplus $\mathbb{E}[U]$ is given for each case, and the dashed red line indicates the boundary between parts in the partitions.

Theorem 3.3. 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.

4 General Impossibility Results

In this section we briefly discuss our second main result, which, unlike Theorem 3.3, holds with respect to *arbitrary* pure-strategy equilibria, not just those induced by connected partitions.

Theorem 4.1. 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.

The full proof, which relies heavily on computational search/verification, is deferred to the full version. Here we provide a relatively simple example illustrating one special case of the theorem, when there is one buyer and two goods with correlated (but still additive) valuations.

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 is clearly 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.

5 Complexity of Welfare Maximization

In settings where disclosure schemes *can* improve buyer welfare, there arises the computational question of how to find such schemes. In this section we study the complexity of finding a buyer-optimal pure-strategy equilibrium over a discrete distribution. We assume that there is one good and one buyer with 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 [Ali *et al.*, 2020], in which buyer messages must be connected, and thus the search problem is restricted to connected partitions. Note that Example 1 of [Ali *et al.*, 2020] illustrates 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; see the full version for the details.

Theorem 5.1. There is a polynomial-time algorithm to compute a connected partition that induces an equilibrium maximizing ex ante buyer utility.

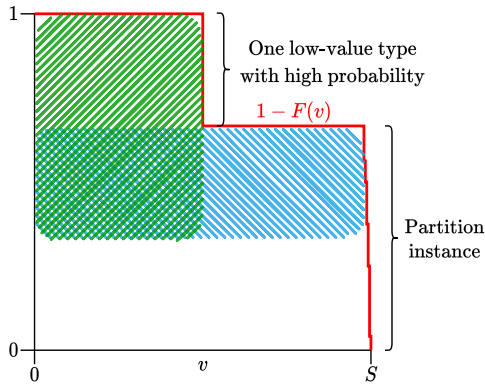


Figure 3: The reduction from PARTITION to BUYER-OPT.

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, and so the buyer could potentially do even better. Thus, suppose now that the 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-optimal disclosure scheme. Further, the optimal surplus using connected messages does not even provide a constant-factor approximation to the unconstrained optimal buyer surplus.

Formally, consider the following 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.*

As might be expected, we reduce from the PARTITION problem. The main idea 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 3. 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 3 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 the full version for the full proof.

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 the full version for the proof.

6 Conclusion

The availability of voluntary disclosure technology is a double-edged sword. While it may be helpful for some buyer types, we have demonstrated how, in complex markets with multiple goods or buyers, it may inevitably reduce the seller’s uncertainty to the point where buyers are harmed.

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, even in the more complicated settings we consider.⁵ We argue only that the compatibility of buyer and social welfare and the buyer advantages of disclosure are no longer guaranteed in larger, realistic market settings. Depending on the setting, instances with multiple buyers/goods where disclosure is useful may well 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 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 yield a higher expected buyer surplus than would be possible with those 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 subgame they are in, standard tools from auction theory do not apply; e.g., even for the setting from Section 3 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 well-approximated in polynomial time. This leaves open the possibility of a pseudo-polynomial time algorithm or a polynomial time approximation scheme for solving this problem.

⁵We present one such example in the full version.

Acknowledgments

We are deeply grateful to Brendan Lucier and Nicole Immorlica for their generous advice and feedback on earlier versions of this paper. We would also like to thank our anonymous IJCAI reviewers for many useful suggestions on improving clarity.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE1745303. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

- [Acquisti *et al.*, 2016] Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–92, 2016.
- [Ali *et al.*, 2020] 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.
- [Alijani *et al.*, 2020] Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. The limits of an information intermediary in auction design. *CoRR*, abs/2009.11841, 2020.
- [Ben-Porath and Lipman, 2012] Elchanan Ben-Porath and Barton L. Lipman. Implementation with partial provability. *Journal of Economic Theory*, 147(5):1689–1724, 2012.
- [Ben-Porath *et al.*, 2018] Elchanan Ben-Porath, Eddie Dekel, and Barton L. Lipman. Disclosure and choice. *The Review of Economic Studies*, 85(3):1471–1501, 2018.
- [Ben-Porath *et al.*, 2019] Elchanan Ben-Porath, Eddie Dekel, and Barton L. Lipman. Mechanisms with evidence: Commitment and robustness. *Econometrica*, 87(2):529–566, 2019.
- [Bergemann and Bonatti, 2019] Dirk Bergemann and Alessandro Bonatti. Markets for information: An introduction. *Annual Review of Economics*, 11:85–107, 2019.
- [Bergemann *et al.*, 2015] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.
- [Gentzkow and Kamenica, 2017] Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.
- [Grossman, 1981] Sanford J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.
- [Hagenbach *et al.*, 2014] Jeanne Hagenbach, Frédéric Koessler, and Eduardo Perez-Richet. Certifiable pre-play communication: Full disclosure. *Econometrica*, 82(3):1093–1131, 2014.
- [Hart *et al.*, 2017] Sergiu Hart, Ilan Kremer, and Motty Perry. Evidence games: Truth and commitment. *American Economic Review*, 107(3):690–713, 2017.
- [Kartik and Tercieux, 2012] Navin Kartik and Olivier Tercieux. Implementation with evidence. *Theoretical Economics*, 7(2):323–355, 2012.
- [Koessler and Skreta, 2019] Frédéric Koessler and Vasiliki Skreta. Selling with evidence. *Theoretical Economics*, 14(2):345–371, 2019.
- [Milgrom, 1981] Paul R. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.
- [Myerson, 1981] Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.
- [Shen *et al.*, 2019] 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.
- [Sher and Vohra, 2015] Itai Sher and Rakesh Vohra. Price discrimination through communication. *Theoretical Economics*, 10(2):597–648, 2015.