Prior-Independent Auctions for Heterogeneous Bidders

Guru Guruganesh[†] Aranyak Mehta[†] Di Wang[†] Kangning Wang[‡]

Abstract

We study the design of prior-independent auctions in a setting with heterogeneous bidders. In particular, we consider the setting of selling to n bidders whose values are drawn from n independent but not necessarily identical distributions. We work in the robust auction design regime, where we assume the seller has no knowledge of the bidders' value distributions and must design a mechanism that is prior-independent. While there have been many strong results on prior-independent auction design in the i.i.d. setting, not much is known for the heterogeneous setting, even though the latter is of significant practical importance. Unfortunately, no prior-independent mechanism can hope to always guarantee any approximation to Myerson's revenue in the heterogeneous setting; similarly, no prior-independent mechanism can consistently do better than the second-price auction. In light of this, we design a family of (parametrized) randomized auctions which approximates at least one of these benchmarks: For heterogeneous bidders with regular value distributions, our mechanisms either achieve a good approximation of the expected revenue of an optimal mechanism (which knows the bidders' distributions) or exceeds that of the second-price auction by a certain multiplicative factor. The factor in the latter case naturally trades off with the approximation ratio of the former case. We show that our mechanism is optimal for such a trade-off between the two cases by establishing a matching lower bound. Our result extends to selling k identical items to heterogeneous bidders with an additional $O(\ln^2 k)$-factor in our trade-off between the two cases.

[†]Google Research. Email: {gurug,aranyak,wadi}@google.com.
[‡]Stanford University. Email: knwang@stanford.edu.

1 Introduction

Auctions are a fundamental component of online commerce and are a suitable mechanism for many different applications. The classic study of auctions focuses on two fundamental objectives: revenue maximization and welfare maximization. In the latter case, there is a long established theory starting with the seminal work of [Vic61] which established the optimal auction for maximizing the welfare of the agents. The auction is relatively simple when we are selling a single item as it coincides with the second-price auction. It is desirable as it is very easy to implement in practice and it requires no assumption on the value distributions of the buyers.

The main drawback of the classical second-price auction is that it achieves no guarantees in general with respect to the optimal revenue that a seller can gain. The optimal revenue is obtained by the celebrated Myerson's auction [Mye81] which uses a more intricate mechanism that utilizes knowledge of a prior distribution on the private values of the buyers. However, such distributional information is often hard to come by and therefore not always practically justified. While one could argue that the seller can learn the distribution over repeated auctions, in practice, distributions need not remain static over time. Furthermore, any such learning mechanism introduces new incentives for the bidders to manipulate the learning itself, making such mechanisms hard to analyze. All this makes the use of Myerson's auction impractical and the second-price auction more appealing despite its poor guarantees on revenue. The need to find mechanisms that do not rely on priors is often referred to as the "Wilson doctrine" or the "Wilson critique" [Wil89] and has been discussed in several works related to ours, e.g., [AB20, DRY15].

A long line of work has tried to bridge the gap between these two different goals of maximizing revenue and not relying on prior information. Towards this, a canonical definition is that of *prior-independent* auctions first introduced by [DRY15]. Here, the goal is to find a mechanism – typically dominant-strategy incentive-compatible (DSIC) – that has no prior information but, no matter what the underlying value distribution is, has revenue competitive with the revenue-optimal auction tailored for that distribution. Specifically, in an i.i.d. setting, once the auction is chosen, an adversary may choose any distribution (from a restricted class) and give each of the buyers a value drawn from this distribution. The mechanism is evaluated in expectation against the revenue achievable by the optimal mechanism that knows the distribution beforehand. In a series of results, beginning with the early work of [Nee03], it is shown that if the distribution comes from a well-formed family such as a *monotone-hazard-rate* (MHR) or *regular* distribution, then it is possible to design such prior-independent auctions in the i.i.d. setting. (We give a more detailed survey in Section 1.2.) The work of [AB20] provided a characterization of optimal prior-independent mechanisms in the i.i.d. setting, gave close upper and lower bounds for two-bidder regular distributions, and proved optimality of the second-price auction for two-bidder MHR distributions. [HJL20] finally closed the gap by giving the tight bounds for two-bidder i.i.d. regular distributions.

While the i.i.d. model for the priors is amenable to clean results and provides nice mathematical intuition, in many practical applications, we do not expect to see identical bidders. In this paper, we investigate the potential of prior-independent auction design for revenue maximization for independent but non-identical (i.e. heterogeneous) bidders. This is especially important in online ad auctions where in the same auction we see advertisers of different scales, or with different goals such as brand advertising, targeted advertising and performance-driven advertising (see, e.g., [GLMN21]). It is therefore imperative to examine what kind of results are achievable in the prior-independent setting with non-identical bidders.

1.1 Our Results and Techniques

Metric and Benchmarks

Extending these results to the heterogeneous[1] setting immediately runs into a challenge: the lack of an existing benchmark that is effective at distinguishing the performance of mechanisms in our context. Indeed, as long as the family of distributions being considered is reasonably general (e.g. MHR or regular distributions), no prior-independent mechanism can achieve anything non-trivial with respect to the canonical benchmarks.

- **No mechanism can guarantee any $\varepsilon > 0$ approximation of the Myerson revenue.** Consider a simple example of two buyers with deterministic (but unknown) values $v_1 = 1$ and $v_2 = x > 1$, where the Myerson revenue is x. Informally speaking,[2] if for any $\varepsilon > 0$ a DSIC prior-independent mechanism M can guarantee at least $\varepsilon \cdot x$ for all x's, then the probability of allocating to buyer 2 must strictly increase with x (by DSIC) and eventually has to go above 1, which gives a contradiction.

  Noting that revenue maximization is trivial for multiple i.i.d. point distributions (via e.g. a second-price auction), we can see that this example already highlights the comparative difficulty in the non-i.i.d. setting. This emphasizes that the prior-independent results in the i.i.d. setting cited above heavily leverage the identical nature of the distributions, originating from the intuition from the work of [BK96] that an additional i.i.d. bidder's random value draw serves as a reasonably good reserve price (see also [DRY15]).

- **No mechanism can guarantee at least $1 + \varepsilon$ times the second-price for any $\varepsilon > 0$.** With the above observation, it is natural to turn to the canonical benchmark in the prior-free setting, and ask if we can always beat the expected revenue of the second-price auction, say by a $(1+\varepsilon)$ factor for some $\varepsilon > 0$. However, this is not possible even over MHR distributions, since it is shown in [AB20] that the second-price auction is the optimal prior-independent mechanism for i.i.d. MHR distributions with two bidders.[3]

This suggests that in the heterogeneous context, any non-trivial guarantee w.r.t. either benchmark would be infeasible, and thus we need to find a new benchmark to tell apart good and bad mechanisms (in terms of revenue guarantees). For this purpose, it is illustrative to revisit our earlier two-buyer example with the family of point distributions where $v_1 = 1$ and $v_2 = x$ for all $x \geq 1$. Consider any DSIC prior-independent mechanism M, and denote $M(x)$ as the expected revenue of M when $v_2 = x$. Myerson knows the distributions and thus always gets revenue x; the second-price always gets revenue 1; and intuitively $M(x)$ should look like Fig. 1.[4] Firstly, as discussed earlier,

[1] We use the term "heterogeneous" to stand for independent but not necessarily identical distributions.

[2] Theorem 4.5 with $\tau \to +\infty$ gives a rigorous proof. Also note one can turn any deterministic value in our examples into a uniform distribution over a tiny range around that value to have continuous and regular bidder value distribution.

[3] Our impossibility results (Theorem 4.7) also rule out mechanisms with "beating second-price" type guarantees such as getting at least $(1+\varepsilon)$ of the second-price revenue whenever possible, and otherwise (i.e. on distributions where Myerson gets less than $(1+\varepsilon)$ times the second-price revenue) getting the same as second-price revenue.

[4] The figure is only for illustration purposes and not meant to be rigorous. Note that x exactly captures the how non-identical the bidders are in this example, and intuitively one can imagine a similar figure in general where the x-axis captures the level of heterogeneity of the bidders, although coming up with a formal metric of it is beyond the scope of our study.

we know for larger x's eventually we won't be able to compare with x, so we can only look at how much better we do compared to 1, which is the only other quantity in the system. In addition, to achieve higher revenue than SPA for larger x's, a DSIC mechanism must reduce the probability of allocating to buyer 2 for smaller x's, which means larger gap compared to the optimal revenue x for smaller x's (as the price for buyer 1 is at most $1 \leq x$).

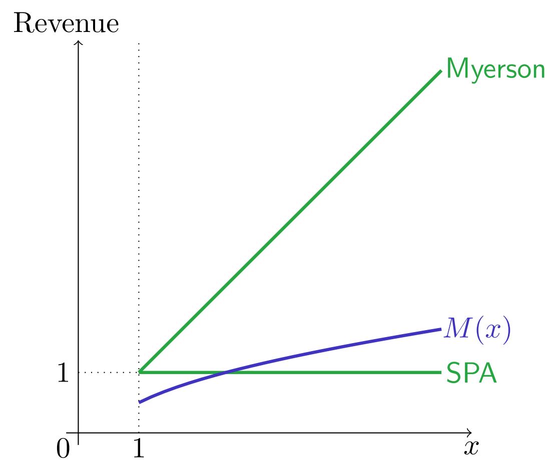

Figure 1: revenue curves for different point distributions in the illustrative example

It is thus clear in this example that when considering worst-case guarantees (i.e. hold for all x), it is a necessary trade-off for *any* prior-independent mechanism between how much advantage compared to 1 for larger x's and how much loss compared to x for smaller x's. A meaningful benchmark in more general cases must also capture this similar phenomenon, i.e. achieving higher revenue (compared to second-price) in high heterogeneity cases must sacrifice revenue (compared to Myerson) when heterogeneity is low (e.g., i.i.d. case). This motivates a natural either-or type benchmark which integrates the two canonical benchmarks: on any bidder distributions (from some family of distributions) the expected revenue of the mechanism can either beat the second-price auction by some multiplicative factor α, or is a β-approximation of the Myerson revenue. Our benchmark takes the either-or form because we do not have a good metric to capture the heterogeneity level in general (and not aim to come up with one in this work), and want a worst-case benchmark (i.e. holds for any level of heterogeneity). Informally α captures the revenue guarantee for higher heterogeneity (i.e. larger x's in our example) and β captures the revenue guarantee on the other end (i.e. smaller x's).

Main results

Our main result (Section 4) is that one can design mechanisms that achieve either a constant fraction of the optimal Myerson's revenue or beat the revenue of the second-price auction by a constant factor.

Theorem (Informal)**.** For any parameter $\tau > e$ there exists a prior-independent randomized mechanism such that for n buyers with private values drawn from any independent (but not necessarily identical) regular distributions, the expected revenue of the mechanism is at least

$$\min\left(\Omega\Big(\frac{1}{\ln\tau}\Big)\cdot\mathsf{Myerson},\ \tau\cdot\mathsf{SPA}\right),$$

where $\mathsf{Myerson}$ is the expected revenue of Myerson's auction for the given distributions, and SPA is that of the second-price auction.

We emphasize that the auctioneer can choose the value of the mechanism's parameter τ, and the resulting revenue guarantee will hold all distributions no matter the level of heterogeneity. One can interpret our benchmark as a *robustness guarantee* in the sense that when SPA is guaranteed to be at least a (good) constant fraction of Myerson (e.g., in the i.i.d. or nearly-i.i.d. cases), the $\tau\cdot\mathsf{SPA}$ part in the benchmark is a good approximation of Myerson, so a guarantee against the min still approximates the optimal revenue well. On the other hand, the benchmark still recovers elegantly in the cases when approximating Myerson is not possible: the expected revenue in such cases is not only bounded from below, but is in fact a τ-factor better than that of SPA.

We complement our main result with an almost matching upper bound.

Theorem (Informal)**.** For 2 bidders and any $\tau \geq 3$, there is no prior-independent mechanism with expected revenue greater than

$$\min\left(\frac{2.5}{\ln\tau}\cdot\mathsf{Myerson},\ \tau\cdot\mathsf{SPA}\right),$$

for all pairs of independent value distributions.

The upper bound above holds even over a very restricted class of distributions – point distributions, i.e., when the two bidders values are deterministic.[5] Note that these distributions are also MHR, thus showing that the main result above cannot be improved significantly for the class of MHR distributions.

To complete the picture, we show (in Section 3) that the main result is tight along different dimensions: One cannot achieve even a minor approximation of this form if either (a) the distributions are allowed to be general (non-regular) independent ones, or (b) the class of distributions is regular, but the mechanism is deterministic. Specifically, in these settings one cannot have a mechanism that achieves either a $(1+\varepsilon)$-multiple of the revenue of the second-price auction or an ε-factor of the revenue of Myerson's auction, for any constant $\varepsilon > 0$.

In Section 5, we extend our main result to the multiple-identical-item case. When we have k identical items to sell and the buyers are unit-demand, we extend our mechanism and achieve a very similar trade-off between the two sides of the revenue objective, losing a factor of $O(\ln^2 k)$ in the trade-off.[6] We prove:

Theorem (Informal)**.** For n buyers and k items, there exists a prior-independent randomized mechanism achieving a revenue of at least

$$\min\left(\Omega\Big(\frac{1}{\ln(k\tau)}\cdot\frac{1}{\ln k}\Big)\cdot\mathsf{Myerson},\ \tau\cdot\mathsf{VCG}\right).$$

[5]This is in stark contrast to the case of point distributions in the i.i.d. setting, where getting Myerson's revenue is easy by using e.g. a second-price auction.

[6]This loss is $O(\ln k)$ when τ is large.

Finally, in Appendix B we provide a characterization of optimal prior-independent mechanisms in our heterogeneous setting for 2 buyers (similar to the characterization for the i.i.d. setting in [AB20]).

Techniques

Our main result uses a simple class of randomized Threshold Mechanisms (see Section 2.2 for definition). Informally, our threshold mechanism chooses a threshold from a carefully constructed distribution over thresholds and only allocate the item if the winner's bid is higher than the next highest bid by a factor of this threshold. It is not hard to see intuitively why one turns to this class of mechanisms in the prior-independent setting: Since a mechanism should be scale-free (see Appendix B), we exclude mechanisms such as posted-price or using reserve prices. Moreover, the mechanism has to be randomized, again by considering the simple two-bidder point distributions example from before. Our idea of randomizing from a set of geometrically separated thresholds can be found in many works in the field of approximation algorithms.

Our novel benchmark poses interesting new mathematical challenges: We emphasize that our benchmark is the minimum (which is a non-linear operator) over the expected revenues of the two canonical auctions over the distributions – it is a much stronger requirement than beating the minimum of the two auctions in each realization of values. There have been related work on prior-free auctions (discussed in Section 1.2) – they can been seen as prior-independent auctions used on *point* value distributions (i.e. the much weaker requirement mentioned in the preceding sentence) and hence are very special cases of our setting of *regular* value distributions. (Recall that we made the regularity assumption since a similar result is impossible without any distributional assumption, according to Theorem 3.1.) To achieve our results, we first use the median of the highest order statistic as an upper bound on the potential revenue achieved by Myerson's auction. We next show that the second-price auction's revenue is related to lower order statistics. We then argue that, depending on the gap between the order statistics, our threshold mechanism can either approximate the highest order statistic or beat the lower order statistics by a suitable margin. This relies on the randomness in our mechanism as it is oblivious to the order-statistics information. The main challenge and novelty in the proof is the *distributional* analysis for arbitrary regular distributions – existing works on prior-free auctions, as the name suggests, do not share with our work this challenge. We believe our techniques have potential applications in other settings related to regular distributions.

As we extend the result to the multi-item setting, these techniques do not suffice. We need to be competitive over a larger set of possibilities depending on the particular decomposition of the distributions. A simple extension would lose a factor of k (the number of items) in the trade-off. To hedge against all these instances, we introduce the technique of randomly limiting the number of items sold, in a manner that does not lose too much in expected revenue. On the one hand, we could sell a large number of items at a small price or sell a few items for a large price. By randomly choosing the number, we hope to set the correct price more often. One particular challenge that arises is that to bound expected revenue of the optimal Myerson auction is non-trivial here. We use a subtle definition of the appropriate order statistics of a carefully chosen subset of the bidders to argue that we can still achieve a constant fraction of Myerson or beat the Vickrey Auction appropriately. This helps us reduce the loss in the trade-off to only $O(\ln^2 k)$. Whether this can be reduced to a constant is an intriguing open question.

1.2 Related Work

There is a lot of related work in the area related to revenue maximization starting with the work of [Mye81]. Due to the practical and technical difficulty of dealing with priors as mentioned earlier, there has been a great deal of work in coming up with prior-independent mechanisms. These works can be classified into a few major related strands.

The first major strand is when the auction is standard and simple, e.g. Vickrey, but the mechanism tries to recruit additional bidders to ensure that the new instance can compete with respect to the optimal revenue. This line of work was initiated by [BK96]. This work gained a great deal of interest in the algorithmic game theory community through the result of [HR09] who showed revenue guarantees for the Vickrey Auction with additional bidders. In particular, they showed that in quite general settings, a VCG Auction with n additional bidders can compete with the optimal revenue when the bidders are drawn from heterogeneous regular distributions. This was extended to a more general class of distributions in the work of [SS13]. The best results along these lines were obtained by [FLR19] who showed that by carefully choosing the extra bidders, one needs only one extra bidder for the single-item case. A number of recent works (see e.g. [FFR18, EFF$^+$17, BW19, CS21]) showed that recruiting additional bidders provides a simple set of mechanisms that achieve near optimal mechanisms even with multiple items and multi-dimensional valuations. Note that in the latter case, the optimal mechanisms are known to be extremely complex. All of the above results argue that a certain level of additional bidders allows a simple mechanism such as the Vickery Auction to compete with the revenue-optimal mechanism. However, it may not be possible to recruit additional bidders because they may not exist or may come at a great cost. Furthermore, these results offer no guarantees with respect to the instance with those additional bidders.

Another line of work attempts to produce bounds on the approximation factor of the revenue of certain simple mechanisms, potentially with knowledge of the prior, compared to the optimal revenue on the same instance. Notably, the bidder-augmentation result of [BK96] can be reinterpreted ([DRY15]) to show that the expected revenue of the Vickrey Auction is at least $^{(n-1)}/_{n} \geq {}^1/_2$ of the expected revenue of Myerson's auction for i.i.d. distributions. A series of works, starting with [HR09], study *simple* auctions which are competitive to the optimal auction even in heterogeneous settings. These results, including [AHN$^+$19, JLQ$^+$19], consider the competitiveness of simple auctions such as the second-price auction with an anonymous reserve. Note that these latter auctions, although simple, still depend on the knowledge of the prior.

Recently, there has been a third line of research building on and improving the result that Vickrey is a $^1/_2$-approximation in the i.i.d. setting, via new prior-independent auctions in two papers highly relevant to our work. Firstly, [FILS15] showed that in the case of a single-item and i.i.d. bidders with a regular distribution, one can beat the revenue approximation guarantees of the Vickrey Auction through the use of randomization. [FILS15] introduced a randomized prior-independent auction called (ε, δ)-inflated second-price auction: with probability ε it runs a second-price auction, and with the rest of the probability the highest bidder wins only if its bid is greater than the next highest bid by a factor of $\alpha \geq 1$; otherwise the item is unallocated. They proved that this auction achieves a fraction of the optimal revenue strictly larger than $\frac{n-1}{n}$ fraction (for $n \geq 2$ bidders), and an improved factor of 0.512 for two bidders, thus beating the approximation guarantee of the second-price auction. This result was further greatly generalized in [AB20] which introduced a family of prior-independent auctions called *threshold-auctions*, and proved stronger results for revenue in the prior-independent two bidder setting. They show an approximation factor of 0.715 for i.i.d. MHR and 0.519 for i.i.d. regular distributions. The factor for MHR distributions is achieved by the second-price auction and is shown to be optimal, while the factor for regular distributions is achieved by a new auction in the class of threshold mechanisms; an upper bound of 0.556 is also shown (under a technical assumption of finite Arzelà variation). This gap between 0.519 and 0.556 was finally closed (under the same technical assumption) by the work of [HJL20], who showed that the optimal prior-independent auction gives an approximation factor of 0.524 for two bidders with i.i.d. regular distributions. Our paper lies in this thread of research and studies the non-identical setting which is arguably more relevant practically as discussed earlier. We note that while some prior work, e.g., [DRY15], does consider prior-independent approximation in a heterogeneous setting, such results still need to assume the existence of multiple bidders with the same attribute, i.e., distribution, which can essentially serve as i.i.d. replacements for each other. We note that our auction GTM also lies in the family of threshold-auctions introduced in [AB20]. This family has been further studied (see e.g. [Meh22, LMP23]) to show stronger welfare guarantees beyond VCG in the auto-bidding setting, which is an increasingly important area in the online advertising industry. This suggests our results on revenue guarantees may be of significant practical interest.

Another way to deal with the distributional assumption is to understand the cost of learning the prior distribution from repeated auctions. [KL03] considered the case where one must learn the value of the buyers' distributions using posted-price mechanisms. [CR14, GHZ19] considered the question of determining how many samples one needs from a distribution to compute the optimal mechanism for revenue maximization. There is a line of work on approximately revenue-optimal auctions with access of 1 sample (see e.g. [DRY15, AKW14, CDFS19]).

Early works on the closely related direction of prior-free auctions include [GHW01, GHK+06, CGL14], where the guarantees are worst-case and valuations are not even drawn from prior distributions. From another point of view, prior-free settings are prior-independent settings limited to (heterogeneous) point distributions. There are other algorithms that study the prior-free setting with additional assumptions such as the buyers having a specific form as such following a low-regret algorithm or participating in a dynamic auction where the state changes as the auctioneer must have limited liability (see e.g. [DSS19, BSW21]). Another approach towards robust auction design is that of distributionally-robust auctions which assumes that the auctioneer has knowledge of some summary statistics of the distribution such as the mean and the upper limit of the support, and characterizes the max-min performance, i.e., under the worst case distribution (see [BTC22, Che22]. A recent work of [ABB22] also tackles the question of designing optimal mechanisms for prior-independent distributions but considers the benchmark of regret. Their results focus on categorizing additive loss between the best mechanism and the optimal welfare that can be achieved and do not translate to giving multiplicative approximations as in our work. Recently, there is another line of work on "revelation gaps" that studies non-truthful auctions in the prior-independent framework [FH18, FHL21]. [HJ21] give lower bounds on prior-independent auctions in the i.i.d. setting.

2 Preliminaries

We consider the setting of selling one indivisible item to n buyers. Buyer i has valuation v_i for the item. v_i is drawn from a distribution V_i, and they are mutually independent. Different from a classical setting where V_i's are public information, we assume the seller and the buyers do not have access to these distributions. The seller, therefore, must use a *prior-independent* auction to sell the item. Her goal is to maximize her expected revenue, using a direct, dominant-strategy incentive-compatible (DSIC) mechanism (meaning that the mechanism asks each buyer for their valuation, and for each buyer, reporting their true valuation is a dominant strategy).

We use $\mathsf{Myerson}$ to denote the revenue-optimal auction characterized by the seminal work of [Mye81], and use SPA to denote the second-price auction. Notice that $\mathsf{Myerson}$ is not prior-independent while SPA is. If the context is clear, we also use $\mathsf{Myerson}$ and SPA to denote their respective expected revenue on some given instance.

We use $v^{(k)}$ to denote the k-th maximum value in $\{v_1, v_2, \ldots, v_n\}$. In particular, $v^{(1)}$ is the maximum value and $\mathsf{SPA} = \mathrm{E}_{(v_1,\ldots,v_n)\sim(V_1,\ldots,V_n)}\left[v^{(2)}\right]$. Moreover, we use s_k to denote the median of the distribution of $v^{(k)}$.

2.1 Regularity

Many of our results use the notion of *regularity* on the value distributions. The regularity assumption is frequently imposed in the literature of auction design, ever since the seminal work of [Mye81].

In the proof of our positive results, we restrict regular distributions to be continuous (which is commonly assumed). This allows us to define the median s of the distribution for some random value v to satisfy $\Pr[v \geq s] = 1/2$, and similarly for other quantiles, to simplify the exposition. The proofs directly generalize to point distributions (i.e. deterministic values; they can be considered as more broadly construed "regular distributions"). However, in our negative results, we consider point distributions (i.e. deterministic values) to be regular.[7]

Formally, regular distributions are the distributions where the virtual value function $\varphi(v) := v - \frac{1-F(v)}{f(v)}$ is nondecreasing in v, where $f(\cdot)$ and $F(\cdot)$ are the corresponding probability density function (PDF) and cumulative distribution function (CDF). Special cases of regular distributions include all monotone-hazard-rate (MHR) distributions – the distributions with hazard rate $\frac{f(v)}{1-F(v)}$ nondecreasing in v.

Lemma 2.1 states a property of a regular distribution, which we will use later.

Lemma 2.1. *For v drawn from a regular distribution V, let r be its Myerson's reserve (i.e., $r \in \arg\max_p p \cdot \Pr[v \geq p]$), and s be its median (i.e. $\Pr[v \geq s] = \frac{1}{2}$). We have*

1. *$r \cdot \Pr[v \geq r] \leq s$. In other words, Myerson's revenue is at most s, and thus at most twice the revenue of selling at s.*
2. *If $s \leq \ell \leq r$ for some ℓ, then $r \cdot \Pr[v \geq r] \leq 2 \cdot \ell \cdot \Pr[v \geq \ell]$.*

Proof. If $r \leq s$, then clearly $r \cdot \Pr[v \geq r] \leq r \leq s$. Otherwise, let $q(p) = \Pr[v \geq p]$ and consider any $\ell \in [s, r]$. For a regular distribution, the revenue is concave in the quantile space, i.e., $v \cdot q(v)$ is concave in $q(v)$. Therefore,

$$\ell \cdot q(\ell) \geq 0 \cdot q(0) \cdot \frac{q(\ell) - q(r)}{q(0) - q(r)} + r \cdot q(r) \cdot \frac{q(0) - q(\ell)}{q(0) - q(r)} = r \cdot q(r) \cdot \frac{1 - q(\ell)}{1 - q(r)} \geq \frac{1}{2} \cdot r \cdot q(r)$$

and therefore proves Statement 2.

Noticing that $q(s) = \frac{1}{2}$, we have $s \geq r \cdot q(r)$ if we set $\ell = s$, which proves Statement 1. $\square$

[7] In the negative results of Theorem 3.1, Theorem 4.5, Theorem 4.6 and Theorem 4.7, we can slightly perturb the distributions to make them continuous regular, without changing the message of the proof. Theorem 3.2 does utilize the fact that there can be a probability mass.

2.2 The Threshold Mechanisms

We will use a class of prior-independent mechanisms: the *threshold mechanisms*. A threshold mechanism uses a finite number of thresholds $\{\lambda_1, \lambda_2, \ldots, \lambda_m\}$, where λ_i happens with probability w_i, with $\sum_{i=1}^{m} w_i = 1$. For a value profile $(v_1, v_2, \ldots, v_n)$, the mechanism generates a random threshold λ_i according to the probabilities $(w_1, w_2, \ldots, w_m)$. It then looks at the top two values, $v^{(1)}$ and $v^{(2)}$. If $v^{(1)} \geq \lambda_i \cdot v^{(2)}$, the item is allocated to the buyer with the highest value for a price of $\lambda_i \cdot v^{(2)}$. Otherwise, the item is not allocated. (Since we are considering continuous distributions, ties happen with probability 0, and we omit the mechanism's behavior to resolve ties. Ties can be broken in any consistent way without affecting our results.) Note that the threshold mechanisms are dominant-strategy incentive-compatible (DSIC).

We also generalize this definition to the multi-item case, which can be found in Section 5.

3 Impossibility Results

As we noted in the introduction, our goal is to obtain a mechanism that is competitive with $\mathsf{Myerson}$ or significantly beats SPA in a prior-independent setting. However, if distributions can be different across all the bidders, we will need to make some assumptions on the class of distributions of the buyers. As we show in Theorem 3.1, no such mechanism can achieve the desired form of guarantees for all general distributions. (Following the convention of this line of work on prior-independent auctions, we limit our attention to DSIC mechanisms.)

Theorem 3.1. *For any constant $\varepsilon > 0$, no DSIC mechanism can guarantee a revenue of $\min(\varepsilon \cdot \mathsf{Myerson}, (1+\varepsilon) \cdot \mathsf{SPA})$ for general valuation distributions, even when there are only two buyers.*

Proof. Fix an integer $k > 1$, and define for $j = 1, \ldots, k$, the instance $\mathcal{I}_j$ to be the following:

- $v_1 = \sqrt{k} \cdot 2^j$ with probability 2^{-j}; and $v_1 = 1$ with probability $1 - 2^{-j}$.
- $v_2 = 1$ with probability 1.

We have $\mathsf{SPA}(\mathcal{I}_j) = 1$, since $v^{(2)} = 1$; and $\mathsf{Myerson}(\mathcal{I}_j) \geq \sqrt{k}$, since selling to Buyer 1 at price $\sqrt{k} \cdot 2^j$ already gives revenue of $\sqrt{k}$. Imagine an adversary who picks instance $\mathcal{I}_j$ with probability $\frac{1}{k}$ for each $j = 1, 2, \ldots, k$. A prior-independent mechanism cannot distinguish which instance the adversary is picking, and it is essentially run on the mixed instance $\mathcal{I}^*$ where

- $v_1 = \sqrt{k} \cdot 2^j$ with probability $2^{-j} \cdot \frac{1}{k}$, for each $j = 1, 2, \ldots, k$; and $v_1 = 1$ with the rest probability.
- $v_2 = 1$ with probability 1.

However, even the optimal auction has $\mathsf{Myerson}(\mathcal{I}^*) = \max_p \left(p \cdot \Pr[v_1 \geq p] + v_2 \cdot (1 - \Pr[v_1 \geq p])\right) \leq 1 + \frac{2}{\sqrt{k}}$. Therefore, the revenue of any prior-independent mechanism on this mixed instance $\mathcal{I}^*$ is at most $1 + \frac{2}{\sqrt{k}}$ too. This means for any prior-independent mechanism, its revenue is at most $1 + \frac{2}{\sqrt{k}}$ for some $\mathcal{I}_j$ where $j \in \{1, 2, \ldots, n\}$. This is only $O\left(\frac{1}{\sqrt{k}}\right) \cdot \mathsf{Myerson}$ or $\left(1 + O\left(\frac{1}{\sqrt{k}}\right)\right) \cdot \mathsf{SPA}$ when $k \to +\infty$. $\square$

This motivates us to limit the class of value distributions. We find that *regularity*, which is a widely imposed assumption in auction theory, suffices for us to show the positive results. Some common examples of regular distributions are uniform, exponential, equal-revenue, and MHR ones.

Now we move on to show that having randomness in our mechanism is also mandatory for a non-trivial guarantee. We state this result in Theorem 3.2.

Theorem 3.2. *For any constant $\varepsilon > 0$, no deterministic DSIC mechanism can guarantee a revenue of $\min(\varepsilon \cdot \mathsf{Myerson}, (1+\varepsilon) \cdot \mathsf{SPA})$ for single-point valuation distributions, even when there are only two buyers.*

Proof. Suppose for the purpose of contradiction that such a deterministic truthful auction exists. It has to always allocate the item to someone, otherwise it gets 0 revenue on that value pair and thus does not meet the theorem condition there.

For $M > 1 + \frac{1}{\varepsilon}$, it has to allocate to Buyer 2 on value pair $(v_1, v_2) = (1, M)$, and the payment has to be at least $(1+\varepsilon)$.

Therefore, on value pair $(1, 1+0.5\varepsilon)$, it has to allocate to Buyer 1. Otherwise Buyer 2 will deviate to report $\hat{v}_2 = (1+0.5\varepsilon)$ on value pair $(1, M)$, to still get the item with less payment.

By the same logic, it has to allocate to Buyer 2 on value pair $\left(1, \frac{1}{1+0.5\varepsilon}\right)$. Therefore the allocation rule is not monotone – Buyer 2 loses the allocation when his value increases – and thus cannot be truthful. □

4 Main Result: the Single-Item Case

We discuss our main result in this section. We give a family of parameterized prior-independent mechanisms that, informally speaking, always achieve a revenue (in expectation) either at least a certain fraction of the optimal prior-dependent expected revenue of Myerson, or much better than the expected revenue of SPA. Our result holds for any set of buyers whose values are drawn from independent, but not necessarily identical, regular distributions.

We first define our geometric-threshold mechanisms, and then present our main result in Theorem 4.2.

Definition 4.1 (Geometric-Threshold Mechanisms)**.** Given parameters $\alpha \geq 1$ and $k \in \mathbb{Z}_+$, we denote $\mathsf{GTM}(\alpha, k)$ as the threshold mechanism with the following $k+1$ thresholds: $\lambda_1 = 1$ and $w_1 = \frac{1}{2}$; $\lambda_i = \alpha^{\frac{i-1}{k}}$ and $w_i = \frac{1}{2k}$ for $i = 2, \ldots, k+1$. Moreover, given a set of distributions on the private values of buyers, we abuse the notation to denote $\mathsf{GTM}(\alpha, k)$ also as the expected revenue of the seller using the corresponding mechanism.

Theorem 4.2. *For any parameter $\tau > e$, there exists some $\alpha = O(\tau \ln^2 \tau)$ and $k = \ln \alpha \in \mathbb{Z}_+$ such that for n buyers with private values drawn from independent (but not necessarily identical) regular distributions,*

$$\mathsf{GTM}(\alpha, k) \geq \min\left(\Omega\left(\frac{1}{\ln \tau}\right) \cdot \mathsf{Myerson},\ \tau \cdot \mathsf{SPA}\right).$$

Recall $\mathsf{GTM}(\alpha, k)$, Myerson and SPA denote the expected revenue achieved by the respective mechanisms from selling one item to the n buyers.

Corollary 4.3 is immediately implied by Theorem 4.2, and captures our main message in a simpler form – for any constant τ, there is a prior-independent mechanism that either beats SPA by a factor of τ, or constant-approximates Myerson.

Corollary 4.3. *For any parameter $\tau = O(1)$, there exists some α and k such that for n buyers with private values drawn from independent (but not necessarily identical) regular distributions,*

$$\mathsf{GTM}(\alpha, k) \geq \min\left(\Omega(1) \cdot \mathsf{Myerson},\ \tau \cdot \mathsf{SPA}\right).$$

Before proving Theorem 4.2, we first present a lemma which bounds the revenue of Myerson in our setting.

Lemma 4.4. *For n buyers with values $v_1, \dots, v_n$ drawn from independent regular distributions $V_1, \dots, V_n$, recall that s_1 is the median of the distribution of $v^{(1)} = \max_{i=1}^n v_i$. The expected revenue of* Myerson *for selling one item to the n buyers satisfies*

$$\frac{1}{2} s_1 \leq \mathsf{Myerson} \leq (1 + 2\ln 2) s_1.$$

Proof. For the first inequality, notice that sequentially posting a price of s_1 for every buyer sells the item with probability $\frac{1}{2}$. Therefore, $\mathsf{Myerson} \geq \frac{1}{2} s_1$, since the revenue-optimal mechanism Myerson gets revenue at least that of the sequential posted-price mechanism.

For the second inequality, consider the virtual welfare (which equals the revenue) achieved by the optimal mechanism. Let p_i be the probability that buyer i wins in the optimal mechanism and z_i be the $(1 - p_i)$-th quantile of V_i, i.e., $\Pr[v_i \geq z_i] = p_i$. Then using the ex-ante relaxation, the virtual welfare from buyer i is at most $p_i \cdot z_i$. If $z_i \leq s_1$, then $p_i \cdot z_i \leq p_i \cdot s_1$. Otherwise (i.e. if $z_i > s_1$), then $p_i \cdot z_i \leq 2 \cdot \Pr[v_i \geq s_1] \cdot s_1$ by regularity of V_i and Lemma 2.1. Therefore,

$$\mathsf{Myerson} \leq s_1 \cdot \sum_{i=1}^n \left(p_i + 2\Pr[v_i > s_1]\right) \leq s_1 \cdot \left(1 + 2\sum_{i=1}^n \Pr[v_i > s_1]\right) \leq (1 + 2\ln 2) s_1.$$

The last step uses the fact that $\sum_{i=1}^n \Pr[v_i > s_1] > \ln 2$ would imply

$$\Pr\left[\max_{i=1}^n v_i > s_1\right] = 1 - \prod_{i=1}^n (1 - \Pr[v_i > s_1]) \geq 1 - \exp\left(-\sum_{i=1}^n \Pr[v_i > s_1]\right) > \frac{1}{2},$$

which contradicts with the definition of s_1 as the median of the distribution of $\max_{i=1}^n v_i$. $\square$

We note that there is a line of work on approximating revenue using simple mechanisms such as anonymous pricing (i.e., sequential posted pricing with the same price); see e.g. [AHN+19, JLQ+19]. Lemma 4.4 gives a simple bound using s_1, where the constant factor should be improvable with techniques from aforementioned work. We now proceed to the proof of our main theorem of this section.

Proof of Theorem 4.2. Again we let s_1 be the median of the distribution of $\max_{i=1}^n V_i$. Without loss of generality, we assume V_1 is the distribution that maximizes $\Pr_{v_i \sim V_i}[v_i \geq s_1]$ over $i \in [n]$. Furthermore, we define u_2 to be the median of the distribution of $\max_{i=2}^n V_i$. It is straightforward from the definitions that $s_1 \geq u_2$. Most of our proof works with a generic $\alpha \geq e$ that gives $k = \ln \alpha \in \mathbb{Z}_+$, and we pick the appropriate α to get the guarantees in terms of τ in the theorem statement at the end of our proof.

We consider the following three cases: (1) When no single value distribution frequently exceeds s_1; (2) When some value distribution frequently exceeds s_1, and s_1 and u_2 are relatively close; (3) When some value distribution frequently exceeds s_1, and $s_1 \gg u_2$.

Case (1): If $\Pr[v_1 \geq s_1] \leq \frac{1}{4}$, we will show $\mathsf{GTM}(\alpha, k)$ is a constant approximation to $\mathsf{Myerson}$. The intuition is that the second highest value is at least s_1 with constant probability; and if there are at least two values exceeding s_1, we will gain a revenue of at least $\frac{s_1}{2}$, since our mechanism uses a threshold $\lambda_1 = 1$ with probability $w_1 = \frac{1}{2}$. Formally, we know

$$\mathsf{GTM}(\alpha, k) \geq \frac{s_1}{2} \cdot \Pr\left[v^{(2)} \geq s_1\right].$$

We will show that $\Pr\left[v^{(2)} \geq s_1\right]$ is at least a constant in this case. Enumerating which two values are at least s_1, we have

$$\begin{aligned}
\Pr\left[v^{(2)} \geq s_1\right] &\geq \sum_{1 \leq i < j \leq n} \Pr[v_i \geq s_1] \cdot \Pr[v_j \geq s_1] \cdot \Pr\left[\max_{t \neq i,j} v_t < s_1\right] \\
&\geq \sum_{1 \leq i < j \leq n} \Pr[v_i \geq s_1] \cdot \Pr[v_j \geq s_1] \cdot \Pr\left[\max_{1 \leq t \leq n} v_t < s_1\right] \\
&= \frac{1}{2} \cdot \sum_{1 \leq i < j \leq n} \Pr[v_i \geq s_1] \cdot \Pr[v_j \geq s_1] \\
&= \frac{1}{4} \cdot \sum_{1 \leq i \leq n} \Pr[v_i \geq s_1] \cdot \sum_{j \neq i} \Pr[v_j \geq s_1],
\end{aligned}$$

where the second last step uses the definition of s_1 being the median of the (continuous) distribution of $v^{(1)}$. Further, since $\sum_{1 \leq j \leq n} \Pr[v_j \geq s_1] \geq \Pr[v^{(1)} \geq s_1] \geq \frac{1}{2}$ and $\Pr[v_i \geq s_1] \leq \Pr[v_1 \geq s_1] \leq \frac{1}{4}$ by our assumption in this case, we have $\sum_{j \neq i} \Pr[v_j \geq s_1] \geq \frac{1}{2} - \frac{1}{4} = \frac{1}{4}$. Therefore,

$$\Pr\left[v^{(2)} \geq s_1\right] \geq \frac{1}{4} \cdot \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{32},$$

and thus

$$\mathsf{GTM}(\alpha, k) \geq \frac{s_1}{64} \geq \frac{1}{64(1 + 2\ln 2)} \cdot \mathsf{Myerson}.$$

Case (2): If $\Pr[v_1 \geq s_1] > \frac{1}{4}$ and $s_1 \leq 12\alpha u_2$, we will show $\mathsf{GTM}(\alpha, k)$ is an $\Omega\left(\frac{1}{k}\right)$-approximation to $\mathsf{Myerson}$. Note that by our choice of $k = \ln \alpha$, the consecutive thresholds in our mechanism are separated by a constant factor of $\alpha^{1/k} = e$.

Observe that if $v_1 \geq s_1$ and $\max_{i=2}^{n} v_i \geq u_2$, then $v^{(2)}$ (i.e. the second highest value) is at least u_2 and $v^{(1)} \geq s_1$. In this case, the thresholds in our mechanisms are $v^{(2)}, e \cdot v^{(2)}, e^2 \cdot v^{(2)}, \ldots, \alpha v^{(2)}$. When $\frac{v^{(1)}}{v^{(2)}} \leq \alpha$, there exists a threshold setting the price to be at least $\frac{v^{(1)}}{e} \geq \frac{s_1}{e}$, and when $\frac{v^{(1)}}{v^{(2)}} > \alpha$, the largest threshold will set the price to be $\alpha v^{(2)} \geq \alpha u_2 \geq s_1/12$. The mechanism will pick each threshold with probability (at least) $\frac{1}{2k}$, and thus $\mathsf{GTM}(\alpha, k)$ gets at least a revenue of $\frac{1}{2k} \cdot \frac{s_1}{12}$ by just looking at when $v_1 \geq s_1$ and $\max_{i=2}^n v_i \geq u_2$. This allows us to show

$$\begin{aligned}
\mathsf{GTM}(\alpha, k) &\geq \frac{1}{2k} \cdot \frac{s_1}{12} \cdot \Pr\left[v_1 \geq s_1 \wedge \max_{i=2}^n v_i \geq u_2\right] \\
&= \frac{1}{2k} \cdot \frac{s_1}{12} \cdot \Pr\left[v_1 \geq s_1\right] \cdot \Pr\left[\max_{i=2}^n v_i \geq u_2\right] \\
&\geq \frac{1}{2k} \cdot \frac{s_1}{12} \cdot \frac{1}{4} \cdot \frac{1}{2} = \frac{s_1}{192k} \geq \frac{1}{192(1+2\ln 2)k} \cdot \mathsf{Myerson}.
\end{aligned}$$

Case (3): Otherwise (i.e., $\Pr[v_1 \geq s_1] > \frac{1}{4}$ and $s_1 > 12\alpha u_2$), we will show $\mathsf{GTM}(\alpha, k) = \Omega(\frac{\alpha}{k \ln \alpha}) \cdot \mathsf{SPA}$. Notice that if $v_1 \geq s_1$ and $u_2 \leq \max_{2 \leq i \leq n} v_i \leq \frac{s_1}{\alpha}$, then $v^{(1)} = v_1$ and $v^{(2)} = \max_{2 \leq i \leq n} v_i \leq \frac{v^{(1)}}{\alpha}$, which means our mechanism will have revenue at least $\frac{1}{2k} \cdot \alpha \cdot \max_{2 \leq i \leq n} v_i$ by using the threshold $\alpha \cdot v^{(2)}$ with probability $\frac{1}{2k}$. Therefore,

$$\begin{aligned}
\mathsf{GTM}(\alpha, k) &\geq \mathrm{E}\left[\frac{1}{2k} \cdot \alpha \cdot \max_{i=2}^n v_i \,\middle|\, v_1 \geq s_1 \wedge u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right] \cdot \Pr\left[v_1 \geq s_1 \wedge u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right] \\
&= \frac{\alpha}{2k} \cdot \mathrm{E}\left[\max_{i=2}^n v_i \,\middle|\, u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right] \cdot \Pr\left[v_1 \geq s_1\right] \cdot \Pr\left[u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right] \\
&\geq \frac{\alpha}{2k} \cdot \mathrm{E}\left[\max_{i=2}^n v_i \,\middle|\, u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right] \cdot \frac{1}{4} \cdot \left(\Pr\left[\max_{i=2}^n v_i \geq u_2\right] - \Pr\left[\max_{i=2}^n v_i > \frac{s_1}{\alpha}\right]\right) \\
&\geq \frac{\alpha}{2k} \cdot \mathrm{E}\left[\max_{i=2}^n v_i \,\middle|\, u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right] \cdot \frac{1}{4} \cdot \left(\frac{1}{2} - \frac{1+2\ln 2}{12}\right) \\
&= \frac{(5 - 2\ln 2)\alpha}{96k} \cdot \mathrm{E}\left[\max_{i=2}^n v_i \,\middle|\, u_2 \leq \max_{i=2}^n v_i \leq \frac{s_1}{\alpha}\right].
\end{aligned}$$

The second-last step uses the fact that $\Pr\left[\max_{i=2}^n v_i \geq 12u_2\right] \leq \frac{1+2\ln 2}{12}$; If that doesn't hold, sequential posted pricing at $12u_2$ on buyers $2, 3, \ldots, n$ would give revenue more than $(1+2\ln 2)u_2$, which contradicts Lemma 4.4.

Next, we give an upper bound of similar form for SPA. Let $\sec_{i=j}^n v_i$ denote the second largest value from the set $\{v_j, v_{j+1}, \ldots, v_n\}$. We have

$$\mathsf{SPA} = \mathrm{E}\left[\sec_{i=1}^n v_i\right] = \int_0^{+\infty} \Pr\left[\sec_{i=1}^n v_i \geq t\right] \, \mathrm{d}t.$$

Evaluating the integral separately at $t \in [0, u_2)$, $t \in [u_2, s_1/\alpha)$, $t \in [s_1/\alpha, s_1)$, and $t \in [s_1, +\infty)$, we get

$$
\begin{aligned}
\mathsf{SPA} \le &\int_0^{u_2} \mathrm{d}t + \int_{u_2}^{s_1/\alpha} \Pr\left[\max_{i=2}^n v_i \ge t\right] \mathrm{d}t + \int_{s_1/\alpha}^{s_1} \Pr\left[\max_{i=2}^n v_i \ge t\right] \mathrm{d}t + \\
&\int_{s_1}^{+\infty} \left(\Pr\left[v_1 \ge t \wedge \max_{i=2}^n v_i \ge t\right] + \Pr\left[\sec_{i=2}^n v_i \ge t\right]\right) \mathrm{d}t \\
= &\, u_2 + \int_{u_2}^{s_1/\alpha} \left(\Pr\left[t \le \max_{i=2}^n v_i \le s_1/\alpha\right] + \Pr\left[\max_{i=2}^n v_i > s_1/\alpha\right]\right) \mathrm{d}t + \int_{s_1/\alpha}^{s_1} \Pr\left[\max_{i=2}^n v_i \ge t\right] \mathrm{d}t + \\
&\int_{s_1}^{+\infty} \left(\Pr\left[v_1 \ge t \wedge \max_{i=2}^n v_i \ge t\right] + \Pr\left[\sec_{i=2}^n v_i \ge t\right]\right) \mathrm{d}t \\
\le &\, u_2 + \int_{u_2}^{s_1/\alpha} \left(\Pr\left[t \le \max_{i=2}^n v_i \le s_1/\alpha\right] + \frac{(1+2\ln 2)u_2\alpha}{s_1}\right) \mathrm{d}t + \int_{s_1/\alpha}^{s_1} \frac{(1+2\ln 2)u_2}{t} \mathrm{d}t + \\
&\int_{s_1}^{+\infty} \left(\frac{s_1}{t} \cdot \frac{(1+2\ln 2)u_2}{t} + \Pr\left[\sec_{i=2}^n v_i \ge t\right]\right) \mathrm{d}t.
\end{aligned}
$$

In the last step we used the fact that $\Pr[v_1 \ge t] \le \frac{s_1}{t}$, since pricing at t for Buyer 1 should not give revenue more than s_1 by Lemma 2.1; and $\Pr\left[\max_{i=2}^n v_i \ge t\right] \le \frac{(1+2\ln 2)u_2}{t}$, since sequential posted pricing at t for Buyer $2, 3, \ldots, n$ should not give revenue more than $(1+2\ln 2)u_2$, which is an upper bound for the optimal revenue given by Lemma 4.4 applied to (only) buyers $2, \ldots, n$. To continue with our derivation, we have

$$
\begin{aligned}
\mathsf{SPA} \le &\, u_2 + \int_{u_2}^{s_1/\alpha} \left(\Pr\left[t \le \max_{i=2}^n v_i \le s_1/\alpha\right] + \frac{(1+2\ln 2)u_2\alpha}{s_1}\right) \mathrm{d}t + \int_{s_1/\alpha}^{s_1} \frac{(1+2\ln 2)u_2}{t} \mathrm{d}t + \\
&\int_{s_1}^{+\infty} \left(\frac{(1+2\ln 2)s_1u_2}{t^2} + \Pr\left[\sec_{i=2}^n v_i \ge t\right]\right) \mathrm{d}t \\
\le &\, u_2 + \int_{u_2}^{s_1/\alpha} \left(\Pr\left[t \le \max_{i=2}^n v_i \le s_1/\alpha \;\middle|\; u_2 \le \max_{i=2}^n v_i \le s_1/\alpha\right]\right) \mathrm{d}t + (1+2\ln 2)u_2 + (1+2\ln 2)u_2 \ln\alpha + \\
&\frac{(1+2\ln 2)s_1u_2}{s_1} + \int_{s_1}^{+\infty} \Pr\left[\sec_{i=2}^n v_i \ge t\right] \mathrm{d}t \\
\le &\, u_2 + \int_{u_2}^{s_1/\alpha} \left(\Pr\left[t \le \max_{i=2}^n v_i \le s_1/\alpha \;\middle|\; u_2 \le \max_{i=2}^n v_i \le s_1/\alpha\right]\right) \mathrm{d}t + \\
&\int_{s_1}^{+\infty} \Pr\left[\sec_{i=2}^n v_i \ge t\right] \mathrm{d}t + (1+2\ln 2)(2+\ln\alpha)u_2 \\
= &\, \mathrm{E}\left[\max_{i=2}^n v_i \;\middle|\; u_2 \le \max_{i=2}^n v_i \le s_1/\alpha\right] + \int_{s_1}^{+\infty} \Pr\left[\sec_{i=2}^n v_i \ge t\right] \mathrm{d}t + ((3+4\ln 2) + (1+2\ln 2)\ln\alpha)u_2.
\end{aligned}
$$

Finally, notice that

$$\begin{aligned}\int_{s_1}^{+\infty} \Pr\left[\sec_{i=2}^{n} v_i \geq t\right] \mathrm{d}t &\leq \int_{s_1}^{+\infty} \Pr\left[\max_{i=2}^{n} v_i \geq t\right]^2 \mathrm{d}t \\ &\leq \int_{s_1}^{+\infty} \frac{((1+2\ln 2)u_2)^2}{t^2} \mathrm{d}t \\ &= \frac{((1+2\ln 2)u_2)^2}{s_1} < \frac{(1+2\ln 2)^2}{12} \cdot u_2,\end{aligned}$$

where we once again used Lemma 4.4. Therefore,

$$\begin{aligned}\mathsf{SPA} &\leq \mathrm{E}\left[\max_{i=2}^{n} v_i \;\middle|\; u_2 \leq \max_{i=2}^{n} v_i \leq s_1/\alpha\right] + \frac{(1+2\ln 2)^2}{12} \cdot u_2 + ((3+4\ln 2) + (1+2\ln 2)\ln\alpha)u_2 \\ &\leq \mathrm{E}\left[\max_{i=2}^{n} v_i \;\middle|\; u_2 \leq \max_{i=2}^{n} v_i \leq s_1/\alpha\right] \cdot \left(1 + \frac{(1+2\ln 2)^2}{12} + (3+4\ln 2) + (1+2\ln 2)\ln\alpha\right).\end{aligned}$$

Thus,

$$\frac{\mathsf{GTM}(\alpha,k)}{\mathsf{SPA}} \geq \frac{\frac{(5-2\ln 2)\alpha}{96k}}{\left(1 + \frac{(1+2\ln 2)^2}{12} + (3+4\ln 2) + (1+2\ln 2)\ln\alpha\right)} \geq \frac{1}{256} \cdot \frac{\alpha}{k\ln\alpha},$$

when $\alpha \geq e$.

Taking $k = \ln\alpha$ and thus $\alpha^{1/k} = e$, we get that in all cases, GTM is either at least $\frac{\alpha}{256\ln^2\alpha} \cdot \mathsf{SPA}$ (i.e. Case (3)) or $\Omega(1/\ln\alpha) \cdot \mathsf{Myerson}$ (i.e. Cases (1),(2)).

To get the guarantees in the theorem and corollary statements, when $\tau \in (1, e]$, it suffices to take $\alpha = e^{12}$ to get $\frac{\alpha}{256\ln^2\alpha} > e \geq \tau$ on the SPA side, and $\Omega(1/\ln\alpha)$ is $\Omega(1)$ on the Myerson side. When $\tau > e$, it suffices to take $\alpha \in [e^{11}, e^{12}] \cdot \tau\ln^2\tau$ (with $k \in \mathbb{Z}_+$). It is easy to check $\frac{\alpha}{256\ln^2\alpha} \geq \tau$ on the SPA side, and $\Omega(1/\ln\alpha)$ is $\Omega(1/\ln\tau)$ on the Myerson side. ∎

4.1 Lower Bounds

We complement our positive result with Theorem 4.5, which states that the guarantee in Theorem 4.2 is tight up to constants, even for 2 buyers with deterministic value distributions.

Theorem 4.5. *For 2 buyers with deterministic values and any $\tau \geq 3$, there is no prior-independent DSIC mechanism* Mec *with revenue satisfying*

$$\mathsf{Mec} \geq \min\left(\frac{2.5}{\ln\tau} \cdot \mathsf{Myerson},\ \tau \cdot \mathsf{SPA}\right).$$

Proof. Suppose for the purpose of contradiction that such a mechanism Mec exists. Fix a parameter $m \in \mathbb{Z}_+$ to be decided later based on τ, and consider the family of m examples each with two buyers whose values are $v_1 = 1$ and $v_2(k) = 2^k$ for $k \in \{1, 2, \ldots, m\}$. Let Mec_k, $\mathsf{Myerson}_k$ and SPA_k be the respective revenues from the point-distribution instance $(v_1, v_2(k))$. Treating each example as its own point-distribution, we need to satisfy the guarantee $\mathsf{Mec}_k \geq \min\left(\frac{2.5}{\ln\tau} \cdot \mathsf{Myerson}_k,\ \tau \cdot \mathsf{SPA}_k\right)$ for $(v_1, v_2(k))$ for all $k \in \{1, 2, \ldots, m\}$ simultaneously. Note that $\mathsf{Myerson}_k = 2^k$ and $\mathsf{SPA}_k = 1$ for all k.

We argue that Mec cannot achieve this guarantee, by first proving $\sum_{k=1}^{m} p_k \cdot \mathsf{Mec}_k \le 3$ where $p_k := \frac{1}{2^k} \cdot \frac{2^m}{2^m - 1}$ gives a probability distribution over k. To show this we consider a single instance with distribution over a support on the aforementioned m examples. Suppose the instance $(v_1, v_2(k))$ appears with probability p_k for each k, then Mec can get a revenue of at most 3 (at most 1 from Buyer 1 and at most 2 from Buyer 2 at any price) on this randomized instance. Choose m so that $2^m + 1 \le \tau < 2^{m+1} + 1$, and thus $\mathsf{Mec}_k \le 2^m < \tau \cdot \mathsf{SPA}_k$. This means for Mec to exist we must have $\mathsf{Mec}_k \ge \frac{2.5}{\ln \tau} \cdot \mathsf{Myerson}_k = \frac{2.5}{\ln \tau} \cdot 2^k$ for all k, then

$$\begin{aligned}\sum_{k=1}^{m} p_k \cdot \mathsf{Mec}_k &\ge \frac{2.5}{\ln \tau} \cdot \sum_{k=1}^{m} p_k \cdot \mathsf{Myerson}_k \\ &= \frac{2.5}{\ln \tau} \cdot \sum_{k=1}^{m} \frac{1}{2^k} \cdot \frac{2^m}{2^m - 1} \cdot 2^k \\ &= \frac{2.5m}{\ln \tau} \cdot \frac{2^m}{2^m - 1} \quad > 3.\end{aligned}$$

The contradiction implies the theorem statement. □

We note that constructions in the work of [AB20] can give lower bounds when τ is close to 1. We present Theorem 4.6 and Theorem 4.7 in addition to our lower bound of Theorem 4.5.

Theorem 4.6 follows from the same instance as in [AB20], where SPA is the optimal prior-independent auction there. Therefore, their lower bound that no prior-independent mechanism (with a technical assumption) can beat $0.715 \cdot \mathsf{Myerson}$ in the instance implies Theorem 4.6.

Theorem 4.6 ([AB20])**.** *Even for two i.i.d. MHR distributions, for any $\varepsilon > 0$, no prior-independent DSIC mechanism* Mec *with finite Arzelà variation (see [AB20, Section 8]) can always satisfy* $\mathsf{Mec} \ge \min(0.715 \cdot \mathsf{Myerson}, (1+\varepsilon) \cdot \mathsf{SPA})$ *.*

For Theorem 4.7, we again look at the same family of instances as in [AB20]. However, here we need to take into account the performance of SPA and balance the parameters. They show no prior-independent mechanism can beat $0.556 \cdot \mathsf{Myerson}$, and later [HJL20] give a stronger, tight impossibility result that no prior-independent mechanism can beat $0.524 \cdot \mathsf{Myerson}$. We show no prior-independent mechanism can beat $\min(0.572 \cdot \mathsf{Myerson}, (1+\varepsilon) \cdot \mathsf{SPA})$.

Theorem 4.7. *Even for two i.i.d. regular distributions, for any $\varepsilon > 0$, no prior-independent DSIC mechanism* Mec *with finite Arzelà variation can always satisfy* $\mathsf{Mec} \ge \min((\frac{4}{7} + \varepsilon) \cdot \mathsf{Myerson}, (1 + \varepsilon) \cdot \mathsf{SPA})$*.*

Proof. Following the work of [AB20], we look at the regular distribution:

$$F_a(v) = \begin{cases} 1 - \frac{1}{v+1} & \text{if } v < a \\ 1 & \text{if } v \ge a \end{cases}.$$

On a pair of these distributions, $\mathsf{SPA} = \int_0^a \left(\Pr_{v \sim F_a}[v \ge t]\right)^2 \mathrm{d}t = \int_0^a \left(\frac{1}{t+1}\right)^2 \mathrm{d}t = \frac{a}{a+1}$, and $\mathsf{Myerson} = a \cdot (1 - (\frac{a}{a+1})^2) = \frac{a(2a+1)}{(a+1)^2}$. [AB20] show that for any $\varepsilon > 0$, no prior-independent mechanism Mec

can guarantee

$$\frac{\mathsf{Mec}}{\mathsf{Myerson}} \geq (1+\varepsilon)\cdot \max\Bigg(\frac{1}{2-q}, \max_{\gamma>1}\left(\frac{q}{2-q} + 2\cdot\frac{\gamma}{\gamma-1}\cdot\frac{1}{1-q}\cdot\frac{1}{2-q}\cdot\left(\frac{1-q}{1-q+\gamma q} - \frac{1}{\gamma-1}\ln\frac{\gamma}{1-q+\gamma q}\right)\right)\Bigg)$$

where $q = \frac{1}{1+a}$.

Let $a = 3$. We have $q = \frac{1}{4}$, and it is impossible to beat $\frac{4}{7}\mathsf{Myerson} = \mathsf{SPA}$. $\square$

5 Selling Multiple Identical Items

In this section, we consider a generalization where the seller is selling k identical items. The buyers are unit-demand, meaning each of them can only receive at most one copy. Without loss of generality, we assume $\log_2 k$ is an integer to simply our exposition. (In general, we can reduce k to the nearest power of 2, and the loss of constant factors are absorbed in the theorem statements in this section.)

Multi-Item Threshold Mechanisms We generalize the threshold mechanisms in Section 2.2 to the multi-item case. Such a mechanism with capacity parameter t still uses a finite number of thresholds $\{\lambda_1, \lambda_2, \ldots, \lambda_m\}$, where λ_i happens with probability w_i with $\sum_{i=1}^m w_i = 1$. For a value profile $(v_1, v_2, \ldots, v_n)$, the mechanism generates a random threshold λ_i according to the probabilities $(w_1, w_2, \ldots, w_m)$. It then looks at values $v^{(j)}$ and $v^{(t+1)}$, separately for each $j \in [t]$. If $v^{(j)} \geq \lambda_i \cdot v^{(t+1)}$, one item is allocated to buyer j for a price of $\lambda_i \cdot v^{(t+1)}$. Otherwise, no item is allocated to buyer j.

Definition 5.1 (Multi-Item Geometric-Threshold Mechanisms)**.** The mechanism $\mathsf{MGTM}(\tau)$ runs $\mathsf{GTM}_{2^j}(\alpha, c)$ for $j \in \{0, 1, \ldots, \log_2 k\}$ each with probability $\frac{1}{1+\log_2 k}$, where $\alpha = O\left(k^2\tau\ln^2(k\tau)\right)$ and $c = \ln\alpha \in \mathbb{Z}_+$. $\mathsf{GTM}_{2^j}(\alpha, c)$ is the multi-item threshold mechanism with capacity parameter $t = 2^j$ and $m+1$ thresholds: $\lambda_1 = 1$ and $w_1 = \frac{1}{2}$; $\lambda_i = \alpha^{\frac{i-1}{k}}$ and $w_i = \frac{1}{2m}$ for $i = 2, \ldots, m+1$.

Theorem 5.2. *For n buyers and k items, the mechanism $\mathsf{MGTM}(\tau)$ satisfies*

$$\mathsf{MGTM}(\tau) \geq \min\left(\Omega\Big(\frac{1}{\ln(k\tau)}\cdot\frac{1}{\ln k}\Big)\cdot\mathsf{Myerson},\ \tau\cdot\mathsf{VCG}\right).$$

5.1 An Upper Bound for Myerson

To prove Theorem 5.2, we first give an upper bound for $\mathsf{Myerson}$. Recall we denote s_i as the median of the distribution of $v^{(i)}$.

Lemma 5.3. $\mathsf{Myerson} \leq 12\cdot\sum_{j=0}^{\log_2 k} 2^j\cdot s_{2^j}$.

Proof. The statement clearly follows from Lemma 4.4 when $k = 1$, and thus we assume $k \geq 2$ in the rest of the proof.

Suppose we have the following $(\log_2 k)+1$ buckets $[0, s_{2^{\log_2 k}}], (s_{2^{\log_2 k}}, s_{2^{(\log_2 k)-1}}], \dots, (s_2, s_1]$. Bucket j is the one whose largest value is s_{2^j}, where $j \in \{0, 1, \dots, \log_2 k\}$. Put each value distribution V_i into one of the buckets depending on which range the median of V_i is in. (Note that the median of V_i cannot exceed s_1, since s_1 is the median of the distribution of $v^{(1)}$.) Let B_j be the set of indices of distributions in Bucket j, and let $n_j := |B_j|$.

Now we describe a relaxation of the allocation constraint. Instead of allocating to at most k buyers, we allow allocation to any buyer in $B_0 \cup B_1 \cup \dots \cup B_{(\log_2 k)-1}$ arbitrarily, and require that we allocate to at most k buyers in $B_{\log_2 k}$. The maximum revenue under these new relaxed constraints is an upper bound for $\textsf{Myerson}$.

The revenue we can get from buyers in $B_0 \cup B_1 \cup \dots \cup B_{(\log_2 k)-1}$ is at most

$$\sum_{j=0}^{(\log_2 k)-1} n_j \cdot s_{2^j}$$

by Lemma 2.1.

The virtual welfare, which is equal to the revenue, that we can get from buyers in $B_{\log_2 k}$ can be upper-bounded similar to the proof of Lemma 4.4. Let p_i be the allocation probability to buyer i, and let w_i be the $(1-p_i)$-th quantile of V_i, i.e., $\Pr[v_i \geq w_i] = p_i$. The virtual welfare from buyer i is at most $p_i \cdot w_i$. If $w_i \leq s_k$, then $p_i \cdot w_i \leq p_i \cdot s_k$. Otherwise, $p_i \cdot w_i \leq 2 \cdot \Pr[v_i \geq s_k] \cdot s_k$ by Lemma 2.1. To summarize, we have

$$\begin{aligned}
\textsf{Myerson} &\leq \left(\sum_{j=0}^{(\log_2 k)-1} n_j \cdot s_{2^j}\right) + \left(\sum_{i \in B_{\log_2 k}} p_i \cdot s_k\right) + \left(\sum_{i \in B_{\log_2 k}} 2 \cdot \Pr[v_i \geq s_k] \cdot s_k\right) \\
&\leq \left(\sum_{j=0}^{(\log_2 k)-1} n_j \cdot s_{2^j}\right) + k \cdot s_k + \left(\sum_{i \in B_{\log_2 k}} 2 \cdot \Pr[v_i \geq s_k] \cdot s_k\right).
\end{aligned}$$

Further, using Lemma 5.4 and Lemma 5.5 that we are about to prove, we get

$$\begin{aligned}
\textsf{Myerson} &\leq \left(\sum_{j=0}^{(\log_2 k)-1} 12 \cdot 2^j \cdot s_{2^j}\right) + k \cdot s_k + 6 \cdot k \cdot s_k \\
&\leq 12 \cdot \sum_{j=0}^{\log_2 k} 2^j \cdot s_{2^j},
\end{aligned}$$

and thus finishes the proof. □

Lemma 5.4. *For $j < \log_2 k$, we have $n_j \leq 12 \cdot 2^j$.*

Proof. Let $X_i = 1$ denote the event of $v_i \geq s_{2^{j+1}}$ and $X_i = 0$ otherwise. X_i's are mutually independent for $i \in B_j$. Let $\mu := \mathrm{E}\left[\sum_{i \in B_j} X_i\right]$. Using the Chernoff bound, we have

$$\Pr\left[\sum_{i \in B_j} X_i \geq \frac{\mu}{2}\right] \geq 1 - e^{-\mu/8}.$$

Note that for each $i \in B_j$, $\Pr[v_i \geq s_{2^{j+1}}] \geq \frac{1}{2}$. Therefore, $\mu \geq \frac{n_j}{2}$ and thus

$$\Pr\left[\sum_{i \in B_j} X_i \geq \frac{n_j}{4}\right] \geq 1 - e^{-n_j/16}.$$

If $n_j > 12 \cdot 2^j$, then

$$\Pr\left[\sum_{i \in B_j} X_i \geq 2^{j+1}\right] \geq 1 - e^{-12/16} > \frac{1}{2},$$

contradicting with the definition of $s_{2^{j+1}}$, as too frequently 2^{j+1} values exceed $s_{2^{j+1}}$. $\square$

Lemma 5.5. *If $k \geq 2$, then $\sum_{i \in B_{\log_2 k}} \Pr[v_i \geq s_k] \leq 3k$.*

Proof. Let $X_i = 1$ denote the event of $v_i \geq s_k$ and $X_i = 0$ otherwise. X_i's are mutually independent for $i \in B_{\log_2 k}$. Let $\mu := \mathrm{E}\left[\sum_{i \in B_{\log_2 k}} X_i\right]$. Using the Chernoff bound, we have

$$\Pr\left[\sum_{i \in B_{\log_2 k}} X_i \geq \frac{\mu}{2}\right] \geq 1 - e^{-\mu/8}.$$

If $\mu > 3k$, then

$$\Pr\left[\sum_{i \in B_{\log_2 k}} X_i \geq k\right] \geq 1 - e^{-3k/8} > \frac{1}{2},$$

contradicting with the definition of s_k, which is the median of the distribution of the k-th maximum value. $\square$

5.2 Extending the Single-Item Case

Lemma 5.6. *For n buyers, k items and parameter τ', there exists some $\alpha = O(\tau' \ln^2 \tau')$ and $c = \ln \alpha \in \mathbb{Z}_+$ such that for each $t \in [k]$,*

$$\mathsf{GTM}_t(\alpha, c) \geq \min\left(\Omega\Big(\frac{1}{\ln \tau'}\Big) \cdot t \cdot s_t,\ \tau' \cdot t \cdot \mathrm{E}[v^{(t+1)}]\right).$$

The proof of Lemma 5.6 is conceptually similar to that of Theorem 4.2, and therefore we postpone it to the appendix. That said, the generalization does require new technical insights. In particular, we prove new technical steps in the form of Lemma A.1 and Lemma A.2.

Lemma A.1 is the core lemma to generalize Case (1) of Theorem 4.2. By definition of s_t, the number of values above s_t is at least t with probability $\frac{1}{2}$. What Lemma A.1 states is under the case condition, the number of values above s_t is at least $t+1$ with constant probability as well, in the style of an "anti-concentration" bound. Lemma A.2 is a major step towards generalizing Case (3) of Theorem 4.2. It is a tail upper bound for the minimum of multiple regular distributions, a counterpart of Lemma 2.1 in the single-item case.

5.3 Completing the Proof

Proof of Theorem 5.2. Notice that $\mathsf{MGTM}(\tau)$ uses the mechanism $\mathsf{GTM}_t(\alpha, c)$ with probability $\Omega\left(\frac{1}{\ln k}\right)$ for each $t = 2^j$ where $j \in \{0, 1, \dots, \log_2 k\}$. For any t, we can use Lemma 5.6 with $\tau' = k^2\tau$ to get

$$\begin{aligned}
\mathsf{GTM}_t(\alpha, c) &\geq \min\left(\Omega\Big(\frac{1}{\ln \tau'}\Big) \cdot t \cdot s_t,\ \tau' \cdot t \cdot \mathrm{E}[v^{(t+1)}]\right) \\
&\geq \min\left(\Omega\Big(\frac{1}{\ln(k^2\tau)}\Big) \cdot t \cdot s_t,\ k^2\tau \cdot t \cdot \mathrm{E}[v^{(k+1)}]\right) \\
&\geq \min\left(\Omega\Big(\frac{1}{\ln(k\tau)}\Big) \cdot t \cdot s_t,\ k\tau \cdot \mathsf{VCG}\right).
\end{aligned}$$

The last step above uses $\mathsf{VCG} = k \cdot \mathrm{E}[v^{(k+1)}]$.

Now we break into two cases:

Case (1): Suppose for some $j \in \{0, 1, \dots, \log_2 k\}$, $\mathsf{GTM}_{2^j}(\alpha, c) \geq k\tau \cdot \mathsf{VCG}$. Since $\mathsf{MGTM}(\tau)$ uses $\mathsf{GTM}_{2^j}(\alpha, c)$ with probability $\frac{1}{1+\log_2 k}$, we know $\mathsf{MGTM}(\tau)$ gets revenue at least $\tau \cdot \mathsf{VCG}$.

Case (2): Otherwise, we know for every $j \in \{0, 1, \dots, \log_2 k\}$, $\mathsf{GTM}_{2^j}(\alpha, c) \geq \Omega\Big(\frac{1}{\ln(k\tau)}\Big) \cdot 2^j \cdot s_{2^j}$. Additionally, Lemma 5.3 states that

$$\sum_{j=0}^{\log_2 k} 2^j \cdot s_{2^j} \geq \frac{1}{12} \cdot \mathsf{Myerson}.$$

Therefore,

$$\begin{aligned}
\mathsf{MGTM}(\tau) &\geq \Omega\Big(\frac{1}{\ln(k\tau)}\Big) \cdot \frac{1}{1 + \log_2 k} \sum_{j=0}^{\log_2 k} 2^j \cdot s_{2^j} \\
&\geq \Omega\Big(\frac{1}{\ln(k\tau)} \cdot \frac{1}{\ln k}\Big) \cdot \mathsf{Myerson}.
\end{aligned}$$

Combining the two cases gives the theorem statement. □

6 Conclusions

In this work, we studied the design of prior-independent auctions for bidders with heterogeneous value distributions. We showed a mechanism that can either achieve a constant fraction of the optimal revenue of any mechanism that knows the value distributions, or beat the revenue of the second-price auction by an arbitrarily large constant factor. Our mechanism has asymptotically optimal trade-off between the constants. We generalized our result to selling multiple identical items and gave a similar message. A possible future direction is to give better bounds and to consider further generalizations.

As another intriguing future direction, one can consider other ways to measure the effectiveness of prior-independent auctions for heterogeneous bidders. What does “approximately optimal” mean and how can we “rank” different mechanisms? We leave alternative answers to these questions for future work.

References

[AB20] Amine Allouah and Omar Besbes. Prior-independent optimal auctions. *Management Science*, 66(10):4417–4432, 2020.

[ABB22] Jerry Anunrojwong, Santiago R. Balseiro, and Omar Besbes. On the robustness of second-price auctions in prior-independent mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 151–152. ACM, 2022.

[AHN+19] Saeed Alaei, Jason D. Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*, 118:494–510, 2019.

[AKW14] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1358–1377, 2014.

[BK96] Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *American Economic Review*, 86(1):180–194, 1996.

[BSW21] Mark Braverman, Jon Schneider, and S. Matthew Weinberg. Prior-free dynamic mechanism design with limited liability. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 204–223, 2021.

[BTC22] Nir Bachrach and Inbal Talgam-Cohen. Distributional robustness: From pricing to auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, page 150. Association for Computing Machinery, 2022.

[BW19] Hedyeh Beyhaghi and S. Matthew Weinberg. Optimal (and benchmark-optimal) competition complexity for additive buyers over independent items. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, page 686–696, 2019.

[CDFS19] José R. Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior. Prophet inequalities for I.I.D. random variables from an unknown distribution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 3–17, 2019.

[CGL14] Ning Chen, Nick Gravin, and Pinyan Lu. Optimal competitive auctions. In David B. Shmoys, editor, *Proceedings of the 46th Annual Symposium on Theory of Computing*, pages 253–262. ACM, 2014.

[Che22] Ethan Che. Robustly optimal auction design under mean constraints. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, page 153–181. Association for Computing Machinery, 2022.

[CR14] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th annual ACM symposium on Theory of computing*, pages 243–252, 2014.

[CS21] Linda Cai and Raghuvansh R. Saxena. 99% revenue with constant enhanced competition. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 224–241. ACM, 2021.

[DRY15] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[DSS19] Yuan Deng, Jon Schneider, and Balasubramanian Sivan. Prior-free dynamic auctions with low regret buyers. In *Proceedings of the 2019 Conference on Neural Information Processing Systems*, pages 4804–4814, 2019.

[EFF+17] Alon Eden, Michal Feldman, Ophir Friedler, Inbal Talgam-Cohen, and S. Matthew Weinberg. The competition complexity of auctions: A bulow-klemperer result for multi-dimensional bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, page 343. ACM, 2017.

[FFR18] Michal Feldman, Ophir Friedler, and Aviad Rubinstein. 99% revenue via enhanced competition. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 443–460, 2018.

[FH18] Yiding Feng and Jason D. Hartline. An end-to-end argument in mechanism design (prior-independent auctions for budgeted agents). In *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science*, pages 404–415, 2018.

[FHL21] Yiding Feng, Jason D. Hartline, and Yingkai Li. Revelation gap for pricing from samples. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1438–1451, 2021.

[FILS15] Hu Fu, Nicole Immorlica, Brendan Lucier, and Philipp Strack. Randomization beats second price as a prior-independent auction. In *Proceedings of the 16th ACM Conference on Economics and Computation*, pages 323–323, 2015.

[FLR19] Hu Fu, Christopher Liaw, and Sikander Randhawa. The vickrey auction with a single duplicate bidder approximates the optimal revenue. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 419–420, 2019.

[GHK+06] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael E. Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

[GHW01] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Proceedings of the 12th Annual Symposium on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.

[GHZ19] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673. ACM, 2019.

[GLMN21] Negin Golrezaei, Max Lin, Vahab S. Mirrokni, and Hamid Nazerzadeh. Boosted second price auctions: Revenue optimization for heterogeneous bidders. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 447–457. ACM, 2021.

[HJ21] Jason D. Hartline and Aleck Johnsen. Lower bounds for prior independent algorithms. *arXiv preprint arXiv:2107.04754*, 2021.

[HJL20] Jason D. Hartline, Aleck C. Johnsen, and Yingkai Li. Benchmark design and prior-independent optimization. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science*, pages 294–305, 2020.

[HR09] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, pages 225–234, 2009.

[JLQ+19] Yaonan Jin, Pinyan Lu, Qi Qi, Zhihao Gavin Tang, and Tao Xiao. Tight approximation ratio of anonymous pricing. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, page 674–685, 2019.

[KL03] Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 594–605. IEEE, 2003.

[LMP23] Christopher Liaw, Aranyak Mehta, and Andrés Perlroth. Efficiency of non-truthful auctions in auto-bidding: The power of randomization. In *Proceedings of the 2023 ACM Web Conference*, pages 3561–3571. ACM, 2023.

[Meh22] Aranyak Mehta. Auction design in an auto-bidding setting: Randomization improves efficiency beyond VCG. In *Proceedings of the 2022 ACM Web Conference*, pages 173–181. ACM, 2022.

[Mye81] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[Nee03] Zvika Neeman. The effectiveness of english auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

[SS13] Balasubramanian Sivan and Vasilis Syrgkanis. Vickrey auctions for irregular distributions. In *Proceedings of the 9th International Conference on Web and Internet Economics*, pages 422–435. Springer, 2013.

[Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[Wil89] Robert Wilson. Game theoretic analysis of trading. In *Advances in Economic Theory: Fifth World Congress*, pages 33–70. CUP Archive, 1989.

A Missing Proofs in Section 5

Lemma A.1 (Anti-Concentration)**.** *Let $X_1, \ldots, X_n$ be mutually independent Bernoulli random variables. $k < n$ is a positive integer. Suppose we have all three conditions below:*

1. $\mathrm{E}[X_1] \geq \mathrm{E}[X_2] \geq \cdots \geq \mathrm{E}[X_n]$.
2. $\Pr\left[\sum_{i=1}^{n} X_i \geq k\right] = \frac{1}{2}$.
3. $\Pr\left[X_1 = X_2 = \cdots = X_k = 1\right] \leq \frac{1}{4}$.

Then,

$$\Pr\left[\sum_{i=1}^{n} X_i \geq k+1\right] > 0.01.$$

Proof. We divide our proof into two cases:
Case (1): Suppose $\mathrm{E}[X_k] \leq \frac{1}{4}$. Define $Q_j = 1$ if $\sum_{i=1}^{j} X_i \geq k$ and $Q_j = 0$ otherwise. We notice that

1. $\mathrm{E}[Q_k] \leq \frac{1}{4}$.
2. $\mathrm{E}[Q_n] = \frac{1}{2}$.
3. $0 \leq \mathrm{E}[Q_{j+1}] - \mathrm{E}[Q_j] \leq \mathrm{E}[X_{j+1}] \leq \frac{1}{4}$, for each $j \in \{k, k+1, \ldots, n-1\}$.

In other words, the sequence $\mathrm{E}[Q_j]$ is increasing in j but cannot have a jump over $\frac{1}{4}$ in a single step. Therefore, there is some $j^* \in \{k, k+1, \ldots, n-1\}$, so that $\mathrm{E}[Q_{j^*}] \in [\frac{1}{8}, \frac{3}{8}]$.

Since $\Pr\left[\sum_{i=1}^{n} X_i \geq k\right] = \frac{1}{2}$ and $\Pr\left[\sum_{i=1}^{j^*} X_i \geq k\right] \leq \frac{3}{8}$, we know

$$\Pr\left[\sum_{i=j^*+1}^{n} X_i \geq 1\right] \geq \frac{1}{2} - \frac{3}{8} = \frac{1}{8}.$$

Further,

$$\Pr\left[\sum_{i=1}^{n} X_i \geq k+1\right] \geq \Pr\left[\sum_{i=1}^{j^*} X_i \geq k\right] \cdot \Pr\left[\sum_{i=j^*+1}^{n} X_i \geq 1\right] \geq \frac{1}{64}.$$

Case (2): Suppose $\mathrm{E}[X_k] > \frac{1}{4}$. Define $R_j = 1$ if $\sum_{i=1}^{j} X_i = j$ and $R_j = 0$ otherwise. Notice that

1. $\mathrm{E}[R_0] = 1$.
2. $\mathrm{E}[R_k] \leq \frac{1}{4}$.
3. $\frac{1}{4} < \mathrm{E}[X_{j+1}] = \frac{\mathrm{E}[R_{j+1}]}{\mathrm{E}[R_j]} \leq 1$, for each $j \in \{0, 1, \ldots, k-1\}$.

In other words, the sequence $\mathrm{E}[R_j]$ is decreasing in j but cannot decrease by a factor more than $\frac{1}{4}$ in a single step. Therefore, there is some $j^* \in \{1, 2, \ldots, k\}$, so that $\mathrm{E}[R_{j^*}] \in [0.1, 0.4]$.

Since $\Pr\left[\sum_{i=1}^{n} X_i \geq k\right] = \frac{1}{2}$ and $\Pr\left[\sum_{i=1}^{j^*} X_i \geq j^*\right] \leq 0.4$, we know

$$\Pr\left[\sum_{i=j^*+1}^{n} X_i \geq k - j^* + 1\right] \geq \frac{1}{2} - 0.4 = 0.1.$$

Therefore,

$$\Pr\left[\sum_{i=1}^{n} X_i \geq k+1\right] \geq \Pr\left[\sum_{i=1}^{j^*} X_i \geq j^*\right] \cdot \Pr\left[\sum_{i=j^*+1}^{n} X_i \geq k - j^* + 1\right] \geq 0.01. \qquad \square$$

Proof of Lemma 5.6. Similar to the proof of Theorem 4.2, we denote $v^{(t)}$ as the t-th largest value among $v_1, \dots, v_n$, and let s_t be the median of the distribution of $v^{(t)}$. Without loss of generality, we index the bidders in non-increasing order according to $\Pr_{v_i \sim V_i}[v_i \ge s_t]$ for $i \in [n]$. Furthermore, we define u_{t+1} to be the median of the distribution of $\max_{i=t+1}^n v_i$. Again we work with a generic $\alpha \ge e$ that gives $c = \ln \alpha \in \mathbb{Z}_+$, and we pick the appropriate α to get the guarantees in terms of τ'' in the theorem statement at the end of our proof.

We can naturally extend the three cases considered in the proof of Theorem 4.2 to multiple items as follows.

Case (1): When $\Pr[v_i \ge s_t \forall i \in [1,t]] \le \frac{1}{4}$, we will show $\mathsf{GTM}(\alpha, k)$ is at least a constant factor of s_t. Define X_i as the indicator of event $v_i \ge s_t$ for $i = 1, \dots, n$. Note our ordering of bidders, together with the definition of s_t and the assumption of this specific case, allows us to use Lemma A.1, which gives

$$\Pr\left[v^{(t+1)} \ge s_t\right] \ge 0.01,$$

so with at least a constant probability, we will have at least $t+1$ values that are at least s_t. When this happens, we will gain a revenue of at least $\frac{t \cdot s_t}{2}$, since our mechanism uses a threshold $\lambda_1 = 1$ with probability $w_1 = \frac{1}{2}$. Consequently, we know

$$\mathsf{GTM}_t(\alpha, c) \ge \frac{t \cdot s_t}{2} \cdot \Pr\left[v^{(t+1)} \ge s_t\right] \ge 0.005 \cdot t \cdot s_t.$$

Case (2): If $\Pr[v_i \ge s_t \forall i \in [1,t]] > \frac{1}{4}$ and $s_t \le 12\alpha u_{t+1}$, we will show $\mathsf{GTM}_t(\alpha, c)$ is at least an $\Omega\left(\frac{1}{k}\right)$ factor of s_t. Note that by our choice of $c = \ln \alpha$, the consecutive thresholds in our mechanism are separated by a constant factor of $\alpha^{1/k} = e$.

Observe that when $v_i \ge s_t$ for all $i \in [1,t]$ and $\max_{j=t+1}^n v_j \ge u_{t+1}$, then $v^{(t+1)}$ is at least $\min(s_t, u_{t+1})$, and thus at least $\frac{s_t}{12\alpha}$ from the assumption in this case. When this happens, the thresholds in our mechanisms are $v^{(t+1)}, e \cdot v^{(t+1)}, e^2 \cdot v^{(t+1)}, \dots, \alpha v^{(t+1)}$. When $\frac{v^{(t)}}{v^{(t+1)}} \le \alpha$, there exists a threshold setting the price to be at least $\frac{v^{(t)}}{e} \ge \frac{s_t}{e}$, and when $\frac{v^{(t)}}{v^{(t+1)}} > \alpha$, the largest threshold will set the price to be $\alpha v^{(t+1)} \ge s_t/12$. The mechanism will pick each threshold with probability (at least) $\frac{1}{2k}$, and thus $\mathsf{GTM}_t(\alpha, c)$ gets at least a revenue of $\frac{t}{2k} \cdot \frac{s_t}{12}$ (since we sell t items) by just looking at when $v_i \ge s_t$ for all $i \in [1,t]$ and $\max_{j=t+1}^n v_j \ge u_{t+1}$. This allows us to show

$$\begin{aligned} \mathsf{GTM}_t(\alpha, c) &\ge \frac{t}{2k} \cdot \frac{s_t}{12} \cdot \Pr\left[v_i \ge s_t \forall i \in [1,t] \wedge \max_{j=t+1}^n v_j \ge u_{t+1}\right] \\ &= \frac{t}{2k} \cdot \frac{s_t}{12} \cdot \Pr[v_i \ge s_t \forall i \in [1,t]] \cdot \Pr\left[\max_{j=t+1}^n v_j \ge u_{t+1}\right] \\ &\ge \frac{t}{2k} \cdot \frac{s_t}{12} \cdot \frac{1}{4} \cdot \frac{1}{2} = \frac{s_t}{192k}, \end{aligned}$$

where in the last step we used the assumption of this case and the definition of u_{t+1} as the median of $\max_{j=t+1}^n v_j$.

Case (3): Otherwise (i.e., $\Pr[v_i \ge s_t \forall i \in [1,t]] > \frac{1}{4}$ and $s_t > 12\alpha u_{t+1}$) . Notice that if $v_i \ge s_t \forall i \in [1,t]$ and $u_{t+1} \le \max_{t+1 \le j \le n} v_j \le \frac{s_t}{\alpha}$, then $v^{(t)} \ge s_t$ and $v^{(t+1)} = \max_{t+1 \le j \le n} v_j \le \frac{v^{(t)}}{\alpha}$, which means our mechanism will have revenue at least $\frac{t}{2k} \cdot \alpha \cdot \max_{t+1 \le j \le n} v_j$ by using the threshold

$\alpha \cdot v^{(t+1)}$ with probability $\frac{1}{2k}$. Therefore,

$$\begin{aligned}
\mathsf{GTM}_t(\alpha, c) \geq & \mathrm{E}\left[\frac{t}{2k} \cdot \alpha \cdot \max_{j=t+1}^{n} v_j \;\middle|\; v_i \geq s_t \forall i \in [1,t] \wedge u_{t+1} \leq \max_{j=t+1}^{n} v_j \leq \frac{s_t}{\alpha}\right] \\
& \cdot \Pr\left[v_i \geq s_t \forall i \in [1,t] \wedge u_{t+1} \leq \max_{j=t+1}^{n} v_j \leq \frac{s_t}{\alpha}\right] \\
= & \frac{\alpha t}{2k} \cdot \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; u_{t+1} \leq \max_{j=t+1}^{n} v_j \leq \frac{s_t}{\alpha}\right] \cdot \Pr\left[v_i \geq s_t \forall i \in [1,t]\right] \cdot \Pr\left[u_{t+1} \leq \max_{j=t+1}^{n} v_j \leq \frac{s_t}{\alpha}\right] \\
\geq & \frac{\alpha t}{2k} \cdot \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; u_{t+1} \leq \max_{j=t+1}^{n} v_j \leq \frac{s_t}{\alpha}\right] \cdot \frac{1}{4} \cdot \left(\Pr\left[\max_{j=t+1}^{n} v_j \geq u_{t+1}\right] - \Pr\left[\max_{j=t+1}^{n} v_j > \frac{s_t}{\alpha}\right]\right) \\
\geq & \frac{\alpha t}{2k} \cdot \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; u_{t+1} \leq \max_{j=t+1}^{n} v_j \leq \frac{s_t}{\alpha}\right] \cdot \frac{1}{4} \cdot \left(\frac{1}{2} - \frac{1+2\ln 2}{12}\right) \\
= & \frac{(5 - 2\ln 2)\alpha t}{96k} \cdot \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; \max_{j=t+1}^{n} v_j \in [u_{t+1}, \frac{s_t}{\alpha}]\right].
\end{aligned}$$

The third-last step uses the assumption of this case, and the second-last step uses the definition of u_{t+1} as the median, and also fact that $\Pr\left[\max_{j=t+1}^{n} v_j \geq 12u_{t+1}\right] \leq \frac{1+2\ln 2}{12}$; If that doesn't hold, in the case of selling a single item to buyers $t+1, \ldots, n$, the sequential posted pricing at $12u_{t+1}$ would give revenue more than $(1+2\ln 2)u_{t+1}$, which contradicts Lemma 4.4.

Next, we give an upper bound of similar form for $\mathrm{E}\left[v^{(t+1)}\right]$. Again, we use $\sec_{i=j}^{n} v_i$ to denote the second largest value from the set $\{v_j, v_{j+1}, \ldots, v_n\}$. We have

$$\mathrm{E}\left[v^{(t+1)}\right] = \int_0^{+\infty} \Pr\left[v^{(t+1)} \geq x\right] \mathrm{d}x.$$

Evaluating the integral separately at $x \in [0, u_{t+1})$, $x \in [u_{t+1}, s_t/\alpha)$, $x \in [s_t/\alpha, s_t)$, and $x \in [s_t, +\infty)$, we get

$$\begin{aligned}
\mathrm{E}\left[v^{(t+1)}\right] \leq & \int_0^{u_{t+1}} \mathrm{d}x + \int_{u_{t+1}}^{s_t/\alpha} \Pr\left[\max_{j=t+1}^{n} v_j \geq x\right] \mathrm{d}x + \int_{s_t/\alpha}^{s_t} \Pr\left[\max_{j=t+1}^{n} v_j \geq x\right] \mathrm{d}x + \int_{s_t}^{+\infty} \Pr\left[v^{(t+1)} \geq x\right] \mathrm{d}x \\
\leq & u_{t+1} + \int_{u_{t+1}}^{s_t/\alpha} \left(\Pr\left[\max_{j=t+1}^{n} v_j \in [x, s_t/\alpha]\right] + \Pr\left[\max_{j=t+1}^{n} v_j > s_t/\alpha\right]\right) \mathrm{d}x + \int_{s_t/\alpha}^{s_t} \Pr\left[\max_{j=t+1}^{n} v_j \geq x\right] \mathrm{d}x + \\
& \int_{s_t}^{+\infty} \left(\Pr\left[v_i \geq x \forall i \in [1,t] \wedge \max_{j=t+1}^{n} v_j \geq x\right] + \Pr\left[\sec_{j=t+1}^{n} v_j \geq x\right]\right) \mathrm{d}x \\
\leq & u_{t+1} + \int_{u_{t+1}}^{s_t/\alpha} \left(\Pr\left[\max_{j=t+1}^{n} v_j \in [x, s_t/\alpha]\right] + \frac{(1+2\ln 2)u_{t+1}\alpha}{s_t}\right) \mathrm{d}x + \int_{s_t/\alpha}^{s_t} \frac{(1+2\ln 2)u_{t+1}}{x} \mathrm{d}x + \\
& \int_{s_t}^{+\infty} \left(\frac{3s_t}{x} \cdot \frac{(1+2\ln 2)u_{t+1}}{x} + \Pr\left[\sec_{j=t+1}^{n} v_j \geq x\right]\right) \mathrm{d}x.
\end{aligned}$$

The first step uses the fact that we always have $v^{(t+1)} \leq \max_{j=t+1}^{n} v_j$. In the last step we used Lemma A.2 and $\frac{8}{e} \leq 3$; and $\Pr\left[\max_{j=t+1}^{n} v_j \geq x\right] \leq \frac{(1+2\ln 2)u_{t+1}}{x}$, since sequential posted pricing at x for Buyer $t+1, \ldots, n$ should not give revenue more than $(1+2\ln 2)u_{t+1}$, which is an upper bound for the optimal revenue given by Lemma 4.4 applied to selling one item to buyers $t+1, \ldots, n$. To continue with our derivation, we have

$$\begin{aligned}
\mathrm{E}\left[v^{(t+1)}\right] \le\ & u_{t+1} + \int_{u_{t+1}}^{s_t/\alpha} \left(\Pr\left[\max_{j=t+1}^{n} v_j \in [x, {}^{s_t}\!/\!_{\alpha}]\right] + \frac{(1+2\ln 2)u_{t+1}\alpha}{s_t} \right) \mathrm{d}x + \int_{s_t/\alpha}^{s_t} \frac{(1+2\ln 2)u_{t+1}}{x} \mathrm{d}x \\
& + \int_{s_t}^{+\infty} \left(\frac{3(1+2\ln 2)s_t u_{t+1}}{x^2} + \Pr\left[\sec_{j=t+1}^{n} v_j \ge x\right] \right) \mathrm{d}x. \\
\le\ & u_{t+1} + \int_{u_{t+1}}^{s_t/\alpha} \left(\Pr\left[\max_{j=t+1}^{n} v_j \in [x, {}^{s_t}\!/\!_{\alpha}] \;\middle|\; \max_{j=t+1}^{n} v_j \in [u_{t+1}, {}^{s_t}\!/\!_{\alpha}]\right] \right) \mathrm{d}x + (1+2\ln 2)u_{t+1} \\
& + (1+2\ln 2)u_{t+1}\ln\alpha + \frac{3(1+2\ln 2)s_t u_{t+1}}{s_t} + \int_{s_t}^{+\infty} \Pr\left[\sec_{j=t+1}^{n} v_j \ge x\right] \mathrm{d}x \\
\le\ & u_{t+1} + \int_{u_{t+1}}^{s_t/\alpha} \left(\Pr\left[\max_{j=t+1}^{n} v_j \in [x, {}^{s_t}\!/\!_{\alpha}] \;\middle|\; \max_{j=t+1}^{n} v_j \in [u_{t+1}, {}^{s_t}\!/\!_{\alpha}]\right] \right) \mathrm{d}x \\
& + \int_{s_t}^{+\infty} \Pr\left[\sec_{j=t+1}^{n} v_j \ge x\right] \mathrm{d}x + (1+2\ln 2)(4+\ln\alpha)u_{t+1} \\
=\ & \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; \max_{j=t+1}^{n} v_j \in [u_{t+1}, {}^{s_t}\!/\!_{\alpha}]\right] + \int_{s_t}^{+\infty} \Pr\left[\sec_{j=t+1}^{n} v_j \ge x\right] \mathrm{d}x \\
& + ((5+8\ln 2) + (1+2\ln 2)\ln\alpha)u_{t+1}.
\end{aligned}$$

Finally, notice that

$$\begin{aligned}
\int_{s_t}^{+\infty} \Pr\left[\sec_{j=t+1}^{n} v_j \ge x\right] \mathrm{d}x &\le \int_{s_t}^{+\infty} \Pr\left[\max_{j=t+1}^{n} v_j \ge x\right]^2 \mathrm{d}x \\
&\le \int_{s_t}^{+\infty} \frac{((1+2\ln 2)u_{t+1})^2}{x^2} \mathrm{d}x \\
&= \frac{((1+2\ln 2)u_{t+1})^2}{s_t} < \frac{(1+2\ln 2)^2}{12} \cdot u_{t+1},
\end{aligned}$$

where we once again used Lemma 4.4 in the second step, and our assumption $s_t > 12\alpha u_{t+1}$ with $\alpha \ge e$ in the last step.

Therefore,

$$\begin{aligned}
& \mathrm{E}\left[v^{(t+1)}\right] \\
\le\ & \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; \max_{j=t+1}^{n} v_j \in [u_{t+1}, {}^{s_t}\!/\!_{\alpha}]\right] + \frac{(1+2\ln 2)^2}{12} \cdot u_{t+1} + ((5+8\ln 2) + (1+2\ln 2)\ln\alpha)u_{t+1} \\
\le\ & \mathrm{E}\left[\max_{j=t+1}^{n} v_j \;\middle|\; \max_{j=t+1}^{n} v_j \in [u_{t+1}, {}^{s_t}\!/\!_{\alpha}]\right] \cdot \left(1 + \frac{(1+2\ln 2)^2}{12} + (5+8\ln 2) + (1+2\ln 2)\ln\alpha\right),
\end{aligned}$$

where the last step is because the expectation is always at least u_{t+1}. Thus,

$$\frac{\mathsf{GTM}_t(\alpha, c)}{\mathrm{E}\left[v^{(t+1)}\right]} \ge \frac{\frac{(5-2\ln 2)\alpha t}{96k}}{\left(1 + \frac{(1+2\ln 2)^2}{12} + (5+8\ln 2) + (1+2\ln 2)\ln\alpha\right)} \ge \frac{1}{512} \cdot \frac{\alpha t}{k\ln\alpha},$$

when $\alpha \ge e$.

Taking $c = \ln\alpha$ and thus $\alpha^{1/c} = e$, we get that in all cases, $\mathsf{GTM}_t(\alpha, c)$ is either at least $\frac{\alpha}{512\ln^2 \alpha t} \cdot \mathrm{E}\left[v^{(t+1)}\right]$ (i.e. Case (3)) or $\Omega(t/\ln\alpha) \cdot s_t$ (i.e. Cases (1),(2)).

To get the guarantees in the theorem and corollary statements, when $\tau' \in (1, e]$, it suffices to take $\alpha = e^{20}$ to get $\frac{\alpha}{512\ln^2\alpha} > e >= \tau'$ on the $\mathrm{E}\left[v^{(t+1)}\right]$ side, and $\Omega(1/\ln\alpha)$ is $\Omega(1)$ on the s_t side. When $\tau' > e$, it suffices to take $\alpha \in [e^{19}, e^{20}] \cdot \tau' \ln^2 \tau'$ (with $c \in \mathbb{Z}_+$). It is easy to check $\frac{\alpha}{512\ln^2\alpha} \geq \tau'$ on the $\mathrm{E}\left[v^{(t+1)}\right]$ side, and $\Omega(1/\ln\alpha)$ is $\Omega(1/\ln\tau')$ on the s_t side. $\square$

Now we prove the lemma used in the above proof.

Lemma A.2. $\Pr[v_i \geq x \forall i \in [1, t]] \leq \frac{8}{e} \cdot \frac{s_t}{x}$ *for* $x > s_t$.

Proof. For a regular distribution V_i, the revenue is concave in probability of selling. Taking the probability of selling as 1, $\Pr[v_i \geq s_t]$, and $\Pr[v_i \geq x]$, we have

$$s_t \cdot \Pr[v_i \geq s_t] \geq x \cdot \Pr[v_i \geq x] \cdot \frac{1 - \Pr[v_i \geq s_t]}{1 - \Pr[v_i \geq x]} + 0.$$

Rearranging and using $\Pr[v_i \geq s_t] \leq 1$, we have

$$\Pr[v_i < x] \geq \frac{x}{s_t} \cdot \Pr[v_i \geq x] \cdot \Pr[v_i < s_t].$$

Define y so that $\min_{i=1}^t \Pr[v_i \geq y] = \frac{1}{2}$. Now we consider the following two cases:
Case (1): We have $x \leq y$. In this case,

$$\Pr[v_i < x] \geq \frac{x}{2s_t} \cdot \Pr[v_i < s_t].$$

Therefore,

$$\begin{aligned}
\Pr[v_i \geq x \forall i \in [1, t]] &= \prod_i \Pr[v_i \geq x] \\
&= \exp\left(\sum_i \ln(1 - \Pr[v_i < x])\right) \\
&\leq \exp\left(-\sum_i \Pr[v_i < x]\right) \\
&\leq \exp\left(-\sum_i \frac{x}{2s_t} \cdot \Pr[v_i < s_t]\right).
\end{aligned}$$

Note that $\sum_i \Pr[v_i < s_t] \geq \Pr[\min_{i=1}^t v_i < s_t] \geq \frac{1}{2}$. We get

$$\Pr[v_i \geq x \forall i \in [1, t]] \leq \exp\left(-\frac{x}{4s_t}\right) \leq \frac{4s_t}{ex}.$$

Case (2): Now consider the case where $x > y$. Let i^* be the buyer with $\Pr[v_{i^*} \geq y] = \frac{1}{2}$. We have

$$\Pr[v_{i^*} \geq x] \leq \frac{y}{x}$$

by Lemma 2.1. Therefore,

$$\begin{aligned}\Pr[v_i \geq x \forall i \in [1,t]] &\leq \Pr[v_i \geq y \forall i \in [1,t]] \cdot \Pr[v_{i^*} \geq x \mid v_{i^*} \geq y] \\ &\leq \frac{4s_t}{ey} \cdot \frac{2y}{x} < \frac{8s_t}{ex}.\end{aligned}$$

Combining the two cases gives the lemma statement. □

B Characterizations of Optimal Prior-Independent Mechanisms

In this section, we give characterizations for the optimal prior-independent mechanisms for 2 buyers. Our proofs are inspired by and generalize those of [AB20] for i.i.d. distributions.

Definition B.1 (Scale-Free Mechanisms [AB20])**.** A mechanism Mec is *scale-free* if

$$x_i(\theta v_i, \theta v_{-i}) = x_i(v_i, v_{-i})$$

for any $\theta > 0, v_i, v_{-i} \geq 0, i = 1, 2$.

Lemma B.2. *If the guarantee* $\mathsf{Mec} \geq \min(\alpha \cdot \mathsf{Myerson}, \beta \cdot \mathsf{SPA})$ *is satisfied by a prior-independent mechanism* Mec *with finite Arzelà variation, then the same guarantee can be satisfied by a scale-free prior-independent mechanism.*

Lemma B.2 intuitively makes sense: If the mechanism does not know the prior distributions, then after any scaling, the instance should be essentially the same. The proof of the i.i.d. case in [AB20] directly extends to this heterogeneous case.

Definition B.3 (Symmetric Mechanisms)**.** A mechanism Mec is *symmetric* if

$$x_1(v_1 = v, v_2 = v') = x_2(v_1 = v', v_2 = v)$$

for any $v, v' \geq 0$.

Lemma B.4. *If the guarantee* $\mathsf{Mec} \geq \min(\alpha \cdot \mathsf{Myerson}, \beta \cdot \mathsf{SPA})$ *is satisfied by a scale-free prior-independent mechanism* Mec*, then the same guarantee can be satisfied by a scale-free symmetric prior-independent mechanism.*

Proof. If a scale-free prior-independent mechanism Mec satisfies the lemma condition, then another mechanism Mec' constructed by switching the roles of Buyer 1 and Buyer 2 also satisfies the lemma condition. Randomizing between Mec and Mec' with equal probability is a scale-free symmetric prior-independent mechanism, and it also guarantees a revenue of $\min(\alpha \cdot \mathsf{Myerson}, \beta \cdot \mathsf{SPA})$. □

Lemma B.5 (Myerson's Lemma [Mye81])**.** $p_i(v_i, v_{-i}) = v_i \cdot x_i(v_i, v_{-i}) - \int_0^{v_i} x_i(t, v_{-i})\, \mathrm{d}t + p_i(0, v_{-i})$, *where* x_i *is the allocation function and* p_i *is the payment function.*

Without loss of generality, a revenue-maximizing auction should set $p_i(0, v_{-i})$ to be 0, as it cannot be positive by the individual rationality (IR) constraint. This allows us to derive the revenue for any mechanism using only the allocation functions. Let $\mathsf{Rev}_1(F, v_2)$ be the expected revenue from Buyer 1 when v_1 is drawn from F and v_2 is fixed. We have

$$\begin{aligned}\mathsf{Rev}_1(F, v_2) &= \int_0^M p_1(v_1, v_2)\,\mathrm{d}F(v_1) \\ &= \int_0^M \left[v_1 \cdot x_1(v_1, v_2) - \int_0^{v_1} x_1(t, v_2)\,\mathrm{d}t\right] \mathrm{d}F(v_1) \\ &= \int_0^M \left[v_1 \cdot x_1(v_1/v_2, 1) - \int_0^{v_1} x_1(t/v_2, 1)\,\mathrm{d}t\right] \mathrm{d}F(v_1)\end{aligned}$$

where the last step above uses Lemma B.2.

Similarly, let $\mathsf{Rev}_2(v_1, G)$ be the expected revenue from Buyer 2 when v_1 is fixed and v_2 is drawn from G. We have

$$\mathsf{Rev}_2(v_1, G) = \int_0^M \left[v_2 \cdot x_2(1, v_2/v_1) - \int_0^{v_2} x_2(1, t/v_1)\,\mathrm{d}t\right] \mathrm{d}G(v_2).$$

Using the symmetry condition of Lemma B.4, we get

$$\mathsf{Rev}_2(v_1, G) = \int_0^M \left[v_2 \cdot x_1(v_2/v_1, 1) - \int_0^{v_2} x_1(t/v_1, 1)\,\mathrm{d}t\right] \mathrm{d}G(v_2).$$

Abbreviating $x_1(v, 1)$ as $x_1(v)$, we have

$$\begin{aligned}\mathsf{Rev} &= \int_0^M \mathsf{Rev}_1(F, v_2)\,\mathrm{d}G(v_2) + \int_0^M \mathsf{Rev}_2(v_1, G)\,\mathrm{d}F(v_1) \\ &= \int_0^M \int_0^M \left[v_1 \cdot x_1(v_1/v_2) - \int_0^{v_1} x_1(t/v_2)\,\mathrm{d}t\right] \mathrm{d}F(v_1)\,\mathrm{d}G(v_2) + \\ &\quad \int_0^M \int_0^M \left[v_2 \cdot x_1(v_2/v_1) - \int_0^{v_2} x_1(t/v_1)\,\mathrm{d}t\right] \mathrm{d}G(v_2)\,\mathrm{d}F(v_1) \\ &= \int_0^M \int_0^M \left[v_1 \cdot x_1(v_1/v_2) + v_2 \cdot x_1(v_2/v_1) - \int_0^{v_1} x_1(t/v_2)\,\mathrm{d}t - \int_0^{v_2} x_1(t/v_1)\,\mathrm{d}t\right] \mathrm{d}F(v_1)\,\mathrm{d}G(v_2).\end{aligned}$$

Let $x_1(r) = \sum_{k=1}^{n} \frac{1}{n} \cdot \mathbf{1}[r \geq \gamma_k]$ with an even n, as an approximation to the possibly continuous increasing function of $x_1(r)$.[8] Assume the γ_k's are decreasing. The constraint on $x_1(r)$ is: $x_1(r) + x_1(1/r) \leq 1, \forall r$, which is equivalent to $\gamma_k \cdot \gamma_{n+1-k} > 1, \forall k$.

We rewrite Rev as:

$$\begin{aligned}\mathsf{Rev} = \sum_{k=1}^{n} \frac{1}{n} \cdot \int_0^M \int_0^M \Bigg[&v_1 \cdot \mathbf{1}[v_1/v_2 \geq \gamma_k] + v_2 \cdot \mathbf{1}[v_2/v_1 \geq \gamma_k] - \\ &\int_0^{v_1} \mathbf{1}[t/v_2 \geq \gamma_k]\,\mathrm{d}t - \int_0^{v_2} \mathbf{1}[t/v_1 \geq \gamma_k]\,\mathrm{d}t\Bigg] \mathrm{d}F(v_1)\,\mathrm{d}G(v_2).\end{aligned}$$

Note that one pair of $(\gamma_k, \gamma_{n+1-k})$ suffices to give optimal revenue (within the class of scale-free symmetric mechanisms) against any fixed instance.

[8] This approximation can be arbitrarily close similar to the proof in [AB20].