Selling to Multiple No-Regret Buyers

Linda Cai[1](✉), S. Matthew Weinberg[1], Evan Wildenhain[2], and Shirley Zhang[3]

[1] Princeton University, Princeton, NJ 08544, USA
tcai@princeton.edu
[2] Google, Mountain View, CA 94043, USA
[3] Harvard University, Cambridge, MA 02138, USA

Abstract. We consider the problem of repeatedly auctioning a single item to multiple i.i.d buyers who each use a no-regret learning algorithm to bid over time. In particular, we study the seller's optimal revenue, if they know that the buyers are no-regret learners (but only that their behavior satisfies some no-regret property—they do not know the precise algorithm/heuristic used).

Our main result designs an auction that extracts revenue equal to the *full expected welfare* whenever the buyers are "mean-based" (a property satisfied by standard no-regret learning algorithms such as Multiplicative Weights, Follow-the-Perturbed-Leader, etc.). This extends a main result of [4] which held only for a single buyer.

Our other results consider the case when buyers are mean-based but never overbid. On this front, [4] provides a simple LP formulation for the revenue-maximizing auction for a single-buyer. We identify several formal barriers to extending this approach to multiple buyers.

1 Introduction

Classical Bayesian auction design considers a static auction where buyers participate once. Here, the study of truthful auctions is ubiquitous following Myerson's seminal work [31]. But many modern auction applications (such as ad auctions) are *repeated*: the same buyers participate in many auctions over time. Moreover, the vast majority of auction formats used in such settings are *not* truthful (e.g. first-price auctions, generalized first-price auctions, generalized second-price auctions). Even those that are based on a truthful format (such as the Vickrey-Clarke-Groves mechanism [10,21,34]) are no longer truthful when the repeated aspect is taken into account (because the seller may increase or decrease reserves in later rounds based on buyers' behavior in earlier rounds). As such, it is imperative to have a study of non-truthful repeated auctions.

Over the past several years, this direction has seen significant progress on numerous fronts (we overview related work in Sect. 1.1). Our paper follows a recent direction initiated by [4] and motivated by empirical work of [32]. Specifically, [32] find that bidding behavior on Bing largely satisfies the no-regret guarantee (that is, there exist values for the buyers such that their bidding behavior

J. Garg et al. (Eds.): WINE 2023, LNCS 14413, pp. 113–129, 2024.
https://doi.org/10.1007/978-3-031-48974-7_7

guarantees low regret—the paper makes no claims about any particular algorithm the buyers might be using). This motivates the following question: *if buyer behavior guarantees no-regret, what auction format for the designer maximizes her expected revenue*?

[4] initiated this study for a single buyer. The main focus of our paper is to initiate the study for multiple buyers. We formally pose the model in Sect. 2, and overview our main results here.

The concept of a "mean-based" no-regret learning algorithm appears in [4], and captures algorithms which pull an arm with very high probability when it is historically better than all other arms (formal definition in Sect. 2). While it is common to design non-mean-based algorithms for dynamic environments, standard no-regret algorithms such as Multiplicative Weights, EXP3, etc. are all mean-based.[1]

Main Result (Informal—See Theorem 3:) When any number of i.i.d. buyers use bidding strategies which satisfy the mean-based no-regret guarantee, there exists a repeated single-item auction for the seller which guarantees them expected revenue arbitrarily close to the optimal expected *welfare*.

One main result of [4] proves the special case with just a single buyer. While we defer technical details of our construction to Sect. 3, we briefly overview the main challenges here. The one-buyer [4] auction is already surprising, as it requires the seller to both (a) give the buyer the item every round, yet (b) charge them their full value (without knowing their value). The key additional challenge for the multi-buyer setting is that the seller must now give the item not just to *a* buyer in every round, but to *the* buyer with the highest value. This means, in particular, that we must set up the auction so that buyers will pull a distinct arm for each possible value, and yet we must still charge each buyer their full expected value by the end of the auction.

Our auction, like that of [4], is fairly impractical (for example, it alternates between running a second-price auction every round, and charging huge surcharges to the winner) and is not meant to guide practice. Still, Theorem 3 establishes that full surplus is possible for multiple mean-based buyers, and therefore sets a high benchmark for this setting without further modeling assumptions.

No Overbidding. Indeed, the main impracticality in our Full-Surplus-Extraction auction is that lures buyers into *overbidding* significantly, and eventually paying more than their value. In practice, it may be reasonable for buyers to be *clever*, and just remove from consideration all bids exceeding their value (but guarantee no-regret on the remaining ones). To motivate this, observe that overbidding is a dominated strategy in all the aforementioned non-truthful auctions. So we next turn to analyze auctions for clever, mean-based buyers.

[1] Note also that these canonical algorithms are mean-based even if the learning rate changes over time, as long as the learning rate is $\omega(1/T)$.

Here, the second main result of [4] characterizes the revenue-optimal repeated auction via a linear program, and shows that it takes a particularly simple form. On this front, we identify *several* formal barriers to extending this result to multiple buyers. Specifically:

- For a single buyer, [4] write a concise polytope (we call it the 'BMSW polytope') characterizing auctions which can be implemented for a single clever mean-based buyer (i.e. being in this polytope is necessary and sufficient to be implementable). We show that two natural extensions of this polytope to multiple buyers contain auctions which *cannot* be implemented for multiple clever mean-based buyers (we show that being in either natural polytope is necessary, but not sufficient). This is in Sect. 4.1.
- For a single buyer, [4] shows that the optimal auction is "pay-your-bid with declining reserve.[2]" We show that a natural generalization of this "pay-your-bid uniform auctions with declining reserve" to multiple buyers captures many extreme points of the multi-buyer BMSW polytope. But, we also show that such auctions are not necessarily optimal (meaning that this aspect of [4] does not generalize to multiple buyers either). This is in Sect. 4.2.
- Finally, we establish that not only does the particular multi-buyer BMSW polytope not capture all implementable auctions for clever mean-based buyers, but the space of implementable auctions is *not even convex*! This is in Sect. 4.3.

While our results are not a death sentence for the future work in the [4] model for clever mean-based buyers, the barriers do shut down most natural multi-buyer extensions of their approach. Still, these barriers also help focus future work towards other potentially fruitful approaches.

1.1 Related Work

There is a vast body of work at the intersection of learning and auction design. Much of this considers learning from the perspective of the *seller* (e.g. sample complexity of revenue-optimal auctions), and is not particularly related [5,6, 18–20,22–25]. More related is the recent and growing literature on dynamic auctions [1,14,17,26,28–30,33] Like our model, the auction is repeated. The distinction between these works and ours is that they assume the buyer is *fully strategic* and processes fully how their actions today affect the seller's decisions tomorrow, perhaps the buyer needs to learn their (whereas we instead model buyers as no-regret learners).

[2] That is, each round there is a reserve. Any bid above the reserve wins the item, but pays their bid. The reserve declines over time.

The most related work to ours is in the [4] model itself, which studies the one seller one buyer scenario, where the buyer employs a mean-based no-regret algorithm. Follow-ups to [4] have extended the setting in [4] in a few different directions. First, [11,15,16] studies convergence of no-regret learning agents in a fixed mechanism such as first price auction, which diverges from the mechanism design perspective of [4]. More relevantly, [13,27] considers the problem of playing a two-player game against a no-regret learner. While technically not an auctions problem, there is thematic overlap with our main result. [12] extends the single-buyer results in [4] to be *prior-free*. Specifically, they show how to design auctions achieving the same guarantees as those in [4] but where the buyer's values are chosen adversarially. In comparison to these works, ours is the first to extend the model to consider multiple buyers.

Finally, [8] considers interaction between a learning buyer and a *learning* seller. Their seller does not have a prior against which to optimize, and instead itself targets a no-regret guarantee. In comparison, our seller (like the seller in all previously cited works) optimizes expected revenue with respect to a prior.

1.2 Organization

The rest of the paper will be organized as follows. In Sect. 2, we discuss our setting where buyer behavior falls under a broad class of no regret learning algorithms and introduce notations that will be used throughout the paper. In Sect. 3, we show that full surplus extraction is possible when buyers are using naive no-regret policies. In Sect. 4, we establish formal barriers in understanding optimal auction design when no regret buyers do not overbid their value. All missing proofs in Sect. 3 and Sect. 4 can be found in the full version of our paper[3].

2 Preliminaries

We consider the same setting as [4], extended to multiple buyers. Specifically, there are n buyers and T rounds. In each round, there is a single item for sale. Each buyer i has value $v_{i,t}$ for the item during round t, and each $v_{i,t}$ is drawn from $\mathcal{D}$ independently (that is, the buyers are i.i.d., and the rounds are i.i.d. as well). For simplicity of exposition (and to match prior work), we assume $\mathcal{D}$ has finite support $0 \leq w_1 < w_2 < \ldots < w_m \leq 1$ and we define q_j to be the probability w_j is drawn from $\mathcal{D}$.

Each round, the seller presents K arms for the buyers. Each arm is labeled with a bid, and we assume that one of the arms is labeled with 0 (to represent a bid of "don't participate"). Note that the same set of arms is presented to all buyers, and the same set of arms is presented in each round.

[3] https://arxiv.org/abs/2307.04175.

In each round t, the seller defines an anonymous auction. Specifically, for all i, t, the seller defines $a_{i,t}(\boldsymbol{b})$ to be the probability that buyer i gets the item in round t, and $p_{i,t}(\boldsymbol{b}) \in [0, b_i \cdot a_{i,t}(\boldsymbol{b})]$ to be the price buyer i pays, when each buyer j pulls the arm labeled b_j. To be anonymous, it must further be that for all permutations σ of the buyers that $(a_{\sigma(i),t}(\sigma(\boldsymbol{b})), p_{\sigma(i),t}(\sigma(\boldsymbol{b})) = (a_{i,t}(\boldsymbol{b}), p_{i,t}(\boldsymbol{b}))$ (the auction is invariant under relabeling buyers). The only additional constraints on a are that $\sum_i a_{i,t}(\boldsymbol{b}) \leq 1$, for all $t, \boldsymbol{b}$ (item can be awarded at most once), and that $b'_i > b_i \Rightarrow a_{i,t}(b_i; \boldsymbol{b}_{-i}) \leq a_{i,t}(b'_i; \boldsymbol{b}_{-i})$ for all $i, \boldsymbol{b}_{-i}, b_i, b'_i$ (allocation is monotone). p must also be monotone ($b'_i > b_i \Rightarrow p_{i,t}(b_i; \boldsymbol{b}_{-i}) \leq p_{i,t}(b'_i; \boldsymbol{b}_{-i})$). When we state prior work in the single-buyer setting, we may drop the buyer subscript of i (for instance, we will write $a_{1,t}(b_1)$ as $a_t(b)$).

2.1 Contextual Bandits

Like [4], we model the buyers as online learners. Also like [4], our results apply equally well to the experts and bandits model, where $v_{i,t}$ serves as buyer i's context for round t. Specifically:

- For all subsequent definitions below, fix a buyer i, fix a bid vector $\boldsymbol{b}_{-i,t}$ for all rounds t, and fix $a_{i,t}(\cdot)$.
- For any bid b, buyer i, and round t, define $r_{ibt}(v) := v \cdot a_{i,t}(b; \boldsymbol{b}_{-i}) - p_{i,t}(b; \boldsymbol{b}_{-i})$. That is, define $r_{ibt}(v)$ to be the utility during round t that buyer i would enjoy by bidding b with value v.
- For an algorithm S (decides a bid for round t based only on what it observes through rounds $t-1$, and its value $v_{i,t}$)[4] that submits bids b_{it} in round t, its total payoff is $P(S) := \mathbb{E}[\sum_t r_{ib_{it}t}(v_{i,t})]$. The expectation is over any randomness in the bids b_{it}, as S may be a randomized algorithm, and the randomness in $v_{i,t}$.
- An algorithm is fixed-bid if $v_{it} = v_{it'} \Rightarrow b_{it} = b_{it'}$. That is, the algorithm may make different bids in different rounds, but only due to changes in the buyer's value. Let $\mathcal{F}$ denote the set of all fixed-bid strategies.
- The *regret* of an online learning algorithm S is $\max_{F \in \mathcal{F}}\{P(F) - P(S)\}$.
- An algorithm is δ-low regret if it guarantees regret at most δ *on every fixed sequence of auctions, and fixed bids of the other players.* We say that an algorithm is *no-regret* if it is δ-low regret for some $\delta = o(T)$.

Like [4], we are particularly interested in algorithms "like Multiplicative Weights Update:"

Definition 1 (Mean-Based Online Learning Algorithm, [4]). *Let* $\sigma_{i,b,s}(v) := \sum_{t<s} r_{ibt}(v)$. *An algorithm is* γ-mean-based *if whenever* $\sigma_{i,b,s}(v_{i,s}) < \sigma_{i,b',s}(v_{i,s}) - \gamma T$ *(for any* b, b'*), then the probability that the algorithm bids* b *during round* s *is at most* γ. *An algorithm is* mean-based *if it is* γ*-mean-based for some* $\gamma = o(1)$.

[4] In the bandits model, buyer i learns only $r_{ibt}(v)$ for the bid $b := b_{it}$ after each round t (and all v). In the experts model, buyer i learns $r_{ibt}(v)$ for all b (and all v).

As noted in [4], natural extensions of Multiplicative Weights, EXP3, Follow the Perturbed Leader, etc. to the contextual setting are all mean-based online learning algorithms.

2.2 Learners and Benchmarks

Before formally stating our main results, we first provide relevant benchmarks. We use $\mathsf{Val}_n(\mathcal{D}) := \mathbb{E}_{\boldsymbol{v} \leftarrow \mathcal{D}^n}[\max_i v_i]$ to denote the expected maximum value among the n buyers. We use $\mathsf{Mye}_n(\mathcal{D})$ to denote the expected revenue of the optimal truthful auction when n buyers have values drawn from $\mathcal{D}$. We make the following quick observation, which holds for *any* low regret learning algorithm (and extends an observation made in [4] for a single buyer).

Observation 1. *The seller cannot achieve expected revenue beyond* $T \cdot \mathsf{Val}_n(\mathcal{D}) + o(T)$ *when buyers guarantee no-regret, even if the seller knows precisely what algorithms the buyers will use.*

Finally, we will consider two types of no-regret learners. One type we will consider is simply no-regret learners who use a mean-based learning algorithm. Second, we will consider no-regret learners who use a no-regret learning algorithm but *never overbid*. Specifically, such learners immediately remove from consideration all bids $b_{it} > v_{i,t}$, but otherwise satisfy the no-regret guarantee. We refer to such learners are *clever*. [4] motivate such learners by observing that in most (perhaps all) standard non-truthful auction formats, overbidding is a *dominated strategy*. For example, it is always better to bid truthfully than to overbid in a first-price auction, generalized first-price auction, generalized second-price auction, and all-pay auction.

2.3 Border's Theorem

Some of our work will use Border's theorem [2], which considers the following. Consider a monotone, anonymous (not necessarily truthful) single-item auction, and a fixed strategy $s(\cdot)$ which maps values to actions. Let $x(w_j)$ denote the probability that a buyer using action $s(w_j)$ wins the item, assuming that all other buyers' values are drawn i.i.d. from $\mathcal{D}$ and use strategy s as well. Border's theorem asks the following: when given some vector $\langle x_1, \ldots, x_m \rangle$, does there exists a monotone anonymous (not necessarily truthful) single-item auction such that $x(w_j) = x_j$ for all j? If so, we say that $\boldsymbol{x}$ is *Border-feasible*. Below is Border's theorem. We will not actually use the precise Border conditions in any of our proofs, just the fact that they exist and are linear in $\boldsymbol{x}$.

Theorem 2 (Border's Theorem [3,7,9]). *When the buyers are drawn i.i.d from $\mathcal{D}$ (meaning each buyer's probability of valuing the item at w_j is q_j), $\boldsymbol{x}$ is Border-feasible if and only if it satisfies the* Border conditions*:*

$$n \sum_{\ell \geq j} q_j \cdot x_j \leq 1 - (1 - \sum_{\ell \geq j} q_j)^n \;\; \forall j \in [m].$$

3 Full Surplus Extraction from Mean-Based Buyers

Here, we show a repeated auction which achieves expected revenue arbitrarily close to $T \cdot \mathsf{Val}_n(\mathcal{D})$ when buyers are mean-based (but consider overbidding). We also note that our auction does not depend on the particular mean-based algorithms used. The auction does *barely* depend on $\mathcal{D}$, but only in initial "setup rounds" (the auction during almost all rounds does not depend on $\mathcal{D}$). Recall this guarantee is the best possible, due to Observation 1.

Theorem 3. *Whenever n buyers use strategies satisfying the mean-based guarantee, there exists a repeated auction which obtains revenue $T \cdot (1-\delta)\mathsf{Val}_n(\mathcal{D}) - o(T)$ for any constant $\delta < 1$.*

In this language, one main result of [4] proves Theorem 3 when $n = 1$. Before diving into our proof, we remind the reader of the main challenge. In order to possibly extract this much revenue, the auction must somehow both (a) charge each winning buyer their full value, leaving them with zero utility, yet also (b) figure out which buyer has the highest value in each round, and give them the item. The distinction between the $n = 1$ and $n > 1$ case is in (b). When $n = 1$, it is still challenging to give the buyer the item every round while charging their full value, but at least the buyer does not need to convey any information to the seller (so, for example, it is not necessary to incentivize the buyer to pull distinct arms for each possible value—the buyers just need to pay their full value on average by the end). When $n > 1$, we need the buyer to pull a distinct arm for each of their possible values, because we need to make sure that the highest buyer wins the item (and the only information we learn about each buyer's value is the arm they pull).

Additional Notation. We now provide our auction and analysis, beginning with some additional notation for this section. We will divide the T rounds of the auction into *phases* of $2R$ consecutive rounds, where $R = \Omega(T)$. There are $P := T/(2R)$ total phases (so P is a constant, but it will be a large constant depending on δ). In our contruction, the first $m-1$ phases will be the setup phases and the last $P - m + 1$ phases will be the main phases. The goal of the setup phases is to align buyer's incentives so that they will behave in a particular manner in later phases. The main phases are where we will extract most of our revenue.

Recall that there are K non-zero arms labeled $b_1 < \ldots < b_K$. Our construction will use $K := P$ arms. Because the buyers consider overbidding, the precise bid labels are not important, so long as they are sufficiently large (concretely, we set $b_i := 2w_m + i$). We will sometimes index arms using $b_j^\tau := b_{P+j-\tau}$. This notation will be helpful to remind the reader that b_j^τ is the arm that we intend to be pulled by a buyer with value w_j during main phase τ.[5]

3.1 Defining the Auction

Intuitively, our auction tries to do the following. In each phase τ, there is a targeted arm b_j^τ for each possible value w_j, so there are m arms that are (intended to be) pulled during each phase. Ideally, since w_j needs to transition from pulling $b_j^{\tau-1}$ in phase $\tau - 1$ to pulling b_j^τ in phase τ, at the beginning of phase τ, w_j should be indifferent between pulling $b_j^{\tau-1}$, w_j's favourite arm in phase $\tau - 1$; and b_j^τ,w_j's intended arm for phase τ. Let us for now assume this is true and see how we design the auction during phase τ (which contains $2R$ rounds).

The base auction each round is just a second-price auction, where pulling arm b_j^τ submits a bid of w_j. For the first R rounds of each phase, this is exactly the auction executed. Because the second-price auction is dominant strategy truthful, this lures a mean-based buyer with value w_j into having high cumulative reward for arm b_j^τ (and in particular, strictly higher than any other arm). For the second R rounds of each phase, the base auction will still be the same second-price auction, except we will now overcharge each buyer so that their average utility during all $2R$ rounds of auction in phase τ is close to zero. In principle, this is possible because the buyers have high cumulative utility for this arm from the first R rounds, and are purely mean-based (and so they will pay more than their value to pull an arm which is historically much better than all others).

Now, by design our auction in phase τ gives the item to the highest buyer most of the time, therefore the expected welfare is almost optimal. Meanwhile, the expected utility is close to 0, which means we have managed to extract revenue that is almost the full welfare in phase τ. Lastly, notice that cumulative utility for arm $b_j^{\tau+1}$ increases during phase τ, so our phase cannot last forever. If we set the phase length to be too long, then $b_j^{\tau+1}$ will become w_j's favourite arm before phase τ ends. This is exactly why we need multiple phases instead of one phase. Let us set our phase length in such a way that at the end of τ, the increase in cumulative utility for arm $b_j^{\tau+1}$ is just enough for w_j to be indifferent between b_j^τ and $b_j^{\tau+1}$, then the exact condition we assume at the start of phase τ is satisfied, but for phase $\tau + 1$. Thus we can safely start a new phase $\tau + 1$ and extract almost full welfare by the same auction design.

Of course, this is just intuition for why an auction like this could possibly work—significant details remain to prove that it does in fact work (including precisely the choice of overcharges, analyzing incentives between phases, etc.). The formal description of our auction can be found in the full version of our paper.

Definition 2 (Full Surplus Extraction Auction). *The FSE Auction uses the following allocation and payment rule in each round. There are two steps in*

[5] So for example, arm b_{P-m} will first be (intended to be) pulled by buyers with value w_1 in phase $m + 1$, then by buyers with value w_2 in phase $m + 2$, etc.

each round. First, based on the arm pulled, a bid is submitted on behalf of the buyer into a secondary auction. Then, the secondary auction is resolved. There are three types of arms:

- *Some arms are* dormant. *These arms don't enter the secondary auction (i.e. no item and* 0 *payment).*
- *Some arms are* active. *Pulling arm* $b_{P-\tau+j} = b_j^\tau$ *enters a bid of* w_j *into a secondary auction.*
- *Some arms are* retired. *Pulling a retired arm enters a bid of* $w_m + 1$ *into a secondary auction.*

Which arms are dormant/active/retired change each phase. In addition, the secondary auction resolves differently for the first $m-1$ *phases (we call these the setup phases) versus the last* $P-m+1$ *phases (we call these the main phases). Think of* $P \gg m$*, so the main phases are what matter most, the setup phases are just a technical setup to get incentives to work out. In any main phase* ($\tau \geq m$)*:*

- *Active arms:* $b_{P-\tau+1} = b_1^\tau$ *through* $b_{P-\tau+m} = b_m^\tau$*. Dormant: below* $b_{P-\tau+1}$*. Retired: above* $b_{P-\tau+m}$*. Note that by our definition, the index of active arms decreases as* τ *increases. For instance, if in* n^{th} *phase the active arms are* $b_h, b_{h-1}, \cdots, b_l$*, then in the* $n+1^{th}$ *phase the active arms are* $b_{h-1}, b_{h-2}, \cdots, b_{l-1}$*.*
- *The secondary auction awards the item to a uniformly random buyer who submits the highest bid.*
- *If the winning arm was retired (i.e., submitted a bid of* $w_m + 1$*), they pay* $2w_m$*.*
- *If the winning arm was active, the winner pays the second-highest bid.*
- *Additionally, in the second* R *rounds of a phase, if the highest bid is* w_j *and the second-highest bid is* w_ℓ*, the winner pays an additional surcharge of* $2(w_j - w_\ell)$*.*

Due to space constraint, the description of the setup phase can be found in the full version of the paper.

We first quickly confirm that the FSE Auction is monotone.

Observation 4. *The allocation and payment rule for the FSE auction are both monotone.*

3.2 Mean-Based Behavior

Before analyzing the expected revenue of the seller, we first analyze the behavior of mean-based buyers. The challenge, of course, is that the payoff from each arm depends on the behavior of the other buyers, who are themselves mean-based. So our goal is to establish that mean-based learning in the FSE auction forms some sort of "equilibrium", in the sense that one mean-based buyer pulls the desired arm almost-always provided that all other buyers pull the desired arm almost-always. Our first step is characterizing a buyer's payoff for each arm at each round, assuming that all other buyers pull the intended arm almost always.

The main steps in our proof are as follows. First, we analyze the cumulative payoff for a buyer with each possible value for each possible arm, assuming that each other buyer pulls their intended arm. We then conclude that a buyer with value w_j has highest cumulative utility for their intended arm for the entirety of each phase. However, we also establish that the utility they enjoy during each phase for their intended arm is 0. This means that every buyer has 0 utility at the end (up to $o(T)$), meaning that the seller's revenue is equal to the expected welfare. Because we give the item to the highest value buyer whenever they pull the intended arm, the welfare is $T \cdot \mathsf{Val}(\mathcal{D})$. We now proceed with each step.

In each of the technical lemmas below, we let $H_s(v, b)$ denote the cumulative payoff during rounds 0 to s that a buyer with value v would have enjoyed in hindsight by pulling arm b in the FSE Auction, assuming that all other buyers pull their intended arm for at least a $1 - o(1)$ fraction of the rounds during every main phase τ, and that they pull either their intended arm (if it exists) or b_P (otherwise) during every setup phase τ. We let $X_{VCG}(v)$ denote the probability that a bidder with value v wins a second-price auction when bidding truthfully against $n-1$ values drawn independently from $\mathcal{D}$ (ties broken randomly). And we let $P_{VCG}(v)$ denote the interim payment made by a bidder with value v to a second-price auction, in expectation over $n-1$ other values drawn independently from $\mathcal{D}$.[6]

Lemma 1. *At the end of phase τ, the change in cumulative payoff of a buyer with value v for each arm satisfies:*

- *For dormant arms b, $H_{2R\tau}(v, b) - H_{2R(\tau-1)}(v, b) = 0$.*
- *For active arms: $H_{2R\tau}(v, b_j^\tau) - H_{2R(\tau-1)}(v, b_j^\tau) = 2R \cdot (v - w_j) \cdot X_{VCG}(w_j) \pm o(T)$.*
- *For retired arms: $H_{2R\tau}(v, b_j) - H_{2R(\tau-1)}(v, b_j) = 2R \cdot (v - 2w_m) \pm o(T)$.*

Corollary 1. *At the end of phase τ, the cumulative payoffs for a buyer with value v satisfy:*

- *If b_j is dormant during phase τ ($j \leq P - \tau$): $H_{2R\tau}(v, b) = 0$.*
- *If b_j is active during phase τ ($j \in [P - \tau + 1, P - \tau + m]$):*

$$H_{2R\tau}(v, b_j) = 2R \cdot \left(\sum_{k=1}^{j+\tau-P} (v - w_k) \cdot X_{VCG}(w_k) \right) \pm o(T).$$

- *If b_j is retired during τ ($j \geq P - \tau + m + 1$):*

$$H_{2R\tau}(v, b) = 2R \cdot \left((\tau - m) \cdot (v - 2w_m) + \sum_{k=1}^{m} (v - w_k) \cdot X_{VCG}(w_k) \right) \pm o(T).$$

[6] Formally, let X_i be independent draws from $\mathcal{D}$ for $i = 1$ to $n - 1$. Define $X_0 := v$. Let $X := \max_{i \geq 1}\{X_i\}$, and Y be an indicator random variable for the event that a uniformly random element in $\arg\max_{i \geq 0}\{X_i\}$ is 0. Then $P_{VCG}(v) := \mathbb{E}[X \cdot Y]$.

Lemma 2. *For all τ, at the start of each phase τ, when $j \leq \tau$, a buyer with value w_j has highest cumulative utility for arm b_j^τ, and also $b_j^{\tau-1}$. Specifically, for all other arms b_ℓ:*

$$H_{2R(\tau-1)}(w_j, b_j^{\tau-1}) \pm o(T) = H_{2R(\tau-1)}(w_j, b_j^\tau) > H_{2R(\tau-1)}(w_j, b_\ell) + \Omega(T).$$

When $j > \tau$, for all $b_\ell \neq \tau$, $H_{2R(\tau-1)}(w_j, b_\tau^\tau) > H_{2R(\tau-1)}(w_j, b_\ell) + \Omega(T)$.

Lemma 3. *For all τ, assuming that all other buyers pull their intended arm except for $o(T)$ rounds, a mean-based buyer with value w_j pulls arm b_j^τ (if it exists) for the first R rounds, except for at most $o(T)$ rounds. Otherwise, they pull arm $b_P = b_\tau^\tau$ for the first R rounds, except for at most $o(T)$ rounds.*

Lemma 4. *For all τ, assuming that all other buyers pull their intended arm except for $o(T)$ rounds, a mean-based buyer with value w_j pulls arm b_j^τ (if it exists) for the last R rounds, except for at most $o(T)$ rounds. Otherwise, they pull arm $b_P = b_\tau^\tau$ for the last R rounds, except for at most $o(T)$ rounds.*

Finally, we combine everything together to conclude the following:

Proposition 1. *When all buyers are mean-based, they all pull their intended arm in the FSE Auction, except for at most $o(T)$ rounds.*

3.3 Analyzing the Revenue

Finally, we show that when all buyers pull their intended arm, the FSE auction extracts full surplus.

Proof (Proof of Theorem 3). Except for the setup phases, and for rounds where buyers do not pull their intended arm, the auction gives the item to the highest buyer. Therefore, the expected welfare of the auction is at least $(1 - m/P)T \cdot \mathsf{Val}(\mathcal{D}) - o(T)$. Moreover, Lemma 1 establishes that through an entire phase, the cumulative utility of a buyer for pulling their intended arm is $0 \pm o(T)$. Therefore, the total utility of the mean-based buyer is at most $o(T)$. Therefore, the revenue is at least $(1 - m/P)T \cdot \mathsf{Val}(\mathcal{D}) - o(T)$. Setting $P \geq m/\delta$ completes the proof.

4 Clever Mean-Based Buyers

In this section we consider clever mean-based buyers, and identify three formal barriers to developing optimal auctions for multiple clever mean-based buyers. We develop each barrier in the subsections below. Section 4.1 reminds the reader of the [4] Linear Program, which exactly captures the optimal auction for a single clever mean-based buyer, and provides a natural extension to multiple buyers.

4.1 A Linear Programming Upper Bound

We first remind the reader of the [4] Linear Program, and give a natural extension to multiple buyers. We first explicitly define variables for the results of a repeated auction.

Let A be a repeated auction with n i.i.d buyers of value distribution $\mathcal{D}$. For each buyer i, let S_i denote a strategy which takes as input a value v_{it} for round t (and all other information available from previous rounds) and outputs an arm $b_{it}^{S_i}(v_{it})$ to pull in round t. Let $\boldsymbol{v} := \langle v_{it} \rangle_{i \in [n], t \in T}$, which is drawn from the product distribution $\times_{nT}\mathcal{D}$. We use the following notation:

$$Rev_A(\mathcal{D}, S_1, \ldots, S_n) := \mathbb{E}_{\boldsymbol{v}}\left[\sum_{i=1}^{n}\sum_{t} p_{i,t}\Big(b_{it}^{S_i}(v_{it}); b_{-it}^{S_{-i}}(v_{-it})\Big)\right]$$

$$Rev_n(\mathcal{D}, S_1, \ldots, S_n) := \max_{A}\{Rev_A(\mathcal{D}, S_1, \ldots, S_n)\}$$

$$X_{ij}^{A}(\mathcal{D}, S_1, \ldots, S_n) = \frac{1}{T}\mathbb{E}_{\boldsymbol{v}}\left[\sum_{t} a_{it}\Big(b_{it}^{S_i}(w_j); b_{-it}^{S_{-i}}(v_{-it})\Big)\right]$$

$$Y_{ij}^{A}(\mathcal{D}, S_1, \ldots, S_n) = \frac{1}{T}\mathbb{E}_{\boldsymbol{v}}\left[\sum_{t} a_{it}\Big(w_j; b_{-it}^{S_{-i}}(v_{-it})\Big)\right]$$

$$U_{ij}^{A}(\mathcal{D}, S_1, \ldots, S_n) = \frac{1}{T}\mathbb{E}_{\boldsymbol{v}}\left[\sum_{t} w_j \cdot a_{it}\Big(b_{it}^{S_i}(w_j); b_{-it}^{S_{-i}}(v_{-it})\Big) - p_{it}\Big(b_{it}^{S_i}(w_j); b_{-it}^{S_{-i}}(v_{-it})\Big)\right]$$

Definition 3 (Auction Feasible). *A tuple of m-vectors (x^*, y^*, u^*) is n-buyer auction feasible if there exists a repeated auction A, such that for all $\gamma = o(T)$, whenever n buyers with values drawn i.i.d. from $\mathcal{D}$ run clever γ-mean-based strategies $S_1, \ldots, S_n$, then $\forall i, X_{ij}^{A}(\mathcal{D}, S_1, \ldots, S_n) = x_j^* \pm O(\gamma); Y_{ij}^{A}(\mathcal{D}, S_1, \ldots, S_n) = y_j^* \pm O(\gamma); U_{ij}^{A}(\mathcal{D}, S_1, \ldots, S_n) = u_j^* \pm O(\gamma)$. We call (x^*, u^*) n-buyer auction feasible if there exists y^* such that (x^*, y^*, u^*) is n-buyer auction feasible.*

One key insight in [4] is that the space of auction feasible tuples is convex and can be characterized by simple linear equations. Below, note that the "only if" direction is slightly non-trivial, and we rederive it later for arbitrary n. The "if" direction requires designing an auction (for which we refer the interested reader to [4, Theorem 3.4]). We will not rederive the "if" direction, although we define the relevant auction later as well.

Theorem 5 ([4]). *(x, u) is 1-buyer auction feasible if and only if it satisfies the* BMSW constraints:[7]

$$\begin{aligned} &u_i \geq (w_i - w_j) \cdot x_j, \quad \forall i, j \in [m] : i > j, \\ &x_i \geq x_j, \quad \forall i \in [m], i > j, \\ &u_i \geq 0, x_i \in [0, 1], \quad \forall i \in [m]. \end{aligned}$$

[7] In fact, the 'only if' portion of this theorem holds when replacing the clever mean-based buyer with just a clever buyer. But the 'if' portion requires the stronger assumption of mean-based learning.

Intuitively, the first BMSW constraint is the interesting one, which is necessary for the buyer to not regret pulling arm b_j when their value is w_i (again recall this is non-trivial, but we argue this shortly as a special case for general n). The second constraint is necessary because the auction must be monotone. The final constraint is necessary because the auction must have a null arm, and because all allocation probabilities must be in $[0, 1]$ every round.

[4] also observe that the expected revenue of an auction A can be computed as a linear function of $X_{ij}^A(\mathcal{D}, S_1, \ldots, S_n)$ and $U_{ij}^A(\mathcal{D}, S_1, \ldots, S_n)$ (because revenue = welfare − utility). Therefore, Theorem 5 enables a simple LP formulation to find the optimal auction for clever buyers.

We consider two natural attempts to generalize Theorem 5, and show that both hold only in the 'only if' direction. The reason the BMSW constraints don't work verbatim for multiple buyers is that the feasibility constraints are wrong: it is not feasible to (for instance) have each buyer win the item with probability 1 every round. Indeed, there is only one copy of the item, implying (for instance) that $n \sum_i q_i \cdot x(w_i) \leq 1$, but also stronger conditions. These conditions are known as *Border's constraints* from Theorem 2 [2].

Proposition 2. *A tuple* (x, y, u) *is* n*-buyer auction feasible only if it satisfies the* n-buyer BMSW constraints *below. A tuple* (x, u) *is* n*-buyer auction feasible only if it satisfies the* reduced n-buyer BMSW constraints.[8]

***n*-buyer BMSW Constraints**

$u_i \geq (w_i - w_j) \cdot y_j, \ \ \forall i > j \in [m],$

$y_i \geq x_i, \ \ \forall i \in [m],$

$u_i \geq 0, \ \ \forall i \in [m],$

$\boldsymbol{x}$ *satisfies Border's constraints,*

$\boldsymbol{x}, \boldsymbol{y}$ *monotone.*

Reduced n-buyer BMSW Constraints

$u_i \geq (w_i - w_j) \cdot x_j, \ \ \forall i > j \in [m],$

$u_i \geq 0, \ \ \forall i \in [m],$

$\boldsymbol{x}$ *satisfies Border's constraints,*

$\boldsymbol{x}$ *monotone.*

We next turn to see whether the other direction holds, as in Theorem 5 for the single-buyer case. If it did, then we could again write a linear program to find the optimal n-buyer feasible auction, because the expected revenue can be written as a function of (x, u). However, we provide an example showing that this extension is *false*.

Theorem 6. *There exist* (x, y, u) *that satisfy the* n*-buyer BMSW Constraints but are not* n*-buyer auction feasible, and* (x, u) *that satisfy the Reduced* n*-buyer BMSW Constraints but are not* n*-buyer auction feasible.*

4.2 Uniform Auctions with Declining Reserves

In this section, we consider the following possibility: although n-buyer BMSW constraints don't characterize the n-buyer feasible auctions, it is conceivable

[8] In fact, this claim holds when replacing mean-based clever buyers with just clever buyers, just like the 'only if' part of Theorem 5.

(although perhaps unlikely) that the Linear Programming solution (optimizing expected revenue subject to n-buyer BMSW constraints) happens to always yield an n-buyer feasible auction. The reason this is a priori possible is because the objective function and BMSW constraints are related: the objective function depends on $\mathcal{D}$, and so do the n-buyer Border constraints (this is another way in which n-buyer and 1-buyer auctions differ: 1-buyer Border constraints don't depend on $\mathcal{D}$).

For the single-buyer case, [4] shows that not only is every tuple satisfying the BMSW constraints 1-buyer auction feasible, but the auction witnessing this is particularly simple. First, whenever the buyer gets the item, they pay their bid (and in each round, each arm gives the item with probability 0 or 1). Second, the minimum winning bid is declining over time. We generalize both definitions below to multiple buyers, and show a connection between these auctions and certain types of tuples which satisfy the n-buyer BMSW conditions.

Definition 4 (Pay-your-bid). *A repeated auction is pay-your-bid if* $p_{i,t}(\boldsymbol{b}) = b_i \cdot a_{i,t}(\boldsymbol{b})$ *for all* i, t.

Definition 5 (Uniform Auction with Declining Reserve). *A repeated auction is a* uniform auction with declining reserve *when: (a) there exists a reserve* r_t *for every round* t *which is monotonically decreasing in* t*, and (b) in each round the item is awarded to a uniformly random buyer among those with* $b_{it} \geq r_t$.

Definition 6 (Correspondence). *We call* (x, y, u) *the* corresponding tuple *of repeated auction* A *if for* 0*-mean-based strategies* $S_1, \ldots, S_n$ *and all* i, j*:*
$X_{ij}^A(\mathcal{D}, S_1, \ldots, S_n) = x_j$; $Y_{ij}^A(\mathcal{D}, S_1, \ldots, S_n) = y_j$; $U_{ij}^A(\mathcal{D}, S_1, \ldots, S_n) = u_j$.[9]

In this language, [4] shows that when $n = 1$, every tuple which satisfies the BMSW conditions can be implemented as a pay-your-bid uniform auction with declining reserve (and this establishes the 'if' direction of Theorem 5). Due to Theorem 6, this claim clearly cannot extend to $n > 1$. However, we show that certain kinds of natural tuples can all be implemented as pay-your-bid uniform auctions with declining reserves.

Theorem 7. *Consider any repeated auction* A *and its corresponding tuple* (x, y, u)*. If* (x, y, u) *satisfies the* n*-buyer BMSW constraints, and* $x = y$*, and* A *is pay-your-bid, then* A *is a uniform auction with declining reserve.*

With Theorem 7 in mind, another possible avenue towards characterizing optimal n-buyer feasible auctions would be through pay-your-bid uniform auctions with declining reserves. To this end, we first show that the optimal pay-your-bid uniform auction with declining reserve can be found by a linear program. However, we also show that examples exist where the optimal n-buyer feasible auction strictly outperforms the best pay-your-bid uniform auction with declining reserve.

[9] Observe that this is always well-defined, as the unique 0-mean-based strategy is Follow-the-leader.

Theorem 8. *The optimal*[10] *pay-your-bid uniform auction with declining reserve can be found by a linear program of size* $\mathsf{Poly}(m)$. *However, there exist* 2*-buyer instances where the optimal* 2*-buyer feasible auction strictly outperforms the best pay-your-bid uniform auction with declining reserve.*

4.3 Non-convexity of N-Buyer Feasible Auctions

Finally, we consider the possibility that while the n-buyer BMSW constraints don't capture the space of n-buyer feasible auctions, perhaps some other compact, convex space does. This too is not the case, as we show that the space of n-buyer feasible triples is non-convex (subject to one technical restriction).

Theorem 9. *Let* P *denote the set of all* (x, y, u) *that are* n*-buyer feasible auctions* where the bid space is equal to the support of $\mathcal{D}$. *Then* P *is not necessarily convex, even when* $n = 2$.

References

1. Ashlagi, I., Daskalakis, C., Haghpanah, N.: Sequential mechanisms with ex-post participation guarantees. In: Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016, Maastricht, The Netherlands, 24–28 July 2016, pp. 213–214 (2016). https://doi.org/10.1145/2940716.2940775
2. Border, K.C.: Implementation of reduced form auctions: a geometric approach. Econometrica **59**(4), 1175–1187 (1991)
3. Border, K.C.: Reduced form auctions revisited. Econ. Theor. **31**, 167–181 (2007)
4. Braverman, M., Mao, J., Schneider, J., Weinberg, M.: Selling to a no-regret buyer. In: Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, 18–22 June 2018, pp. 523–538 (2018). https://doi.org/10.1145/3219166.3219233
5. Brustle, J., Cai, Y., Daskalakis, C.: Multi-item mechanisms without item-independence: learnability via robustness. In: Biró, P., Hartline, J., Ostrovsky, M., Procaccia, A.D. (eds.) EC 2020: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, 13–17 July 2020, pp. 715–761. ACM (2020). https://doi.org/10.1145/3391403.3399541
6. Cai, Y., Daskalakis, C.: Learning multi-item auctions with (or without) samples. In: 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, 15–17 October 2017, pp. 516–527 (2017). https://doi.org/10.1109/FOCS.2017.54
7. Cai, Y., Daskalakis, C., Weinberg, S.M.: A constructive approach to reduced-form auctions with applications to multi-item mechanism design. arXiv preprint arXiv:1112.4572 (2011)
8. Camara, M.K., Hartline, J.D., Johnsen, A.C.: Mechanisms for a no-regret agent: beyond the common prior. In: 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, 16–19 November 2020, pp. 259–270. IEEE (2020). https://doi.org/10.1109/FOCS46700.2020.00033

[10] By optimal we mean the auction achieves the best revenue if all buyers run 0-regret algorithms. It is easy to see that when buyers have γ-regret, the revenue is within $O(n\gamma)$ of the revenue when buyers have 0-regret.

9. Che, Y.K., Kim, J., Mierendorff, K.: Generalized reduced-form auctions: a network-flow approach. Econometrica **81**, 2487–2520 (2013)
10. Clarke, E.H.: Multipart pricing of public goods. Public Choice **11**(1), 17–33 (1971)
11. Deng, X., Hu, X., Lin, T., Zheng, W.: Nash convergence of mean-based learning algorithms in first price auctions. In: Laforest, F., et al. (eds.) WWW 2022: The ACM Web Conference 2022, Virtual Event, Lyon, France, 25–29 April 2022, pp. 141–150. ACM (2022). https://doi.org/10.1145/3485447.3512059
12. Deng, Y., Schneider, J., Sivan, B.: Prior-free dynamic auctions with low regret buyers. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8–14 December 2019, Vancouver, BC, Canada, pp. 4804–4814 (2019). http://papers.nips.cc/paper/8727-prior-free-dynamic-auctions-with-low-regret-buyers
13. Deng, Y., Schneider, J., Sivan, B.: Strategizing against no-regret learners. In: Wallach, H.M., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E.B., Garnett, R. (eds.) Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8–14 December 2019, Vancouver, BC, Canada, pp. 1577–1585 (2019). http://papers.nips.cc/paper/8436-strategizing-against-no-regret-learners
14. Deng, Y., Zhang, H.: Prior-independent dynamic auctions for a value-maximizing buyer. In: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, 6–14 December 2021, Virtual, pp. 13847–13858 (2021)
15. Feng, S., Yu, F., Chen, Y.: Peer prediction for learning agents. In: NeurIPS (2022)
16. Feng, Z., Guruganesh, G., Liaw, C., Mehta, A., Sethi, A.: Convergence analysis of no-regret bidding algorithms in repeated auctions. In: Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, 2–9 February 2021, pp. 5399–5406. AAAI Press (2021). https://doi.org/10.1609/aaai.v35i6.16680
17. Golrezaei, N., Jaillet, P., Liang, J.C.N., Mirrokni, V.: Bidding and pricing in budget and ROI constrained markets. arXiv preprint arXiv:2107.07725 (2021)
18. Gonczarowski, Y.A., Nisan, N.: Efficient empirical revenue maximization in single-parameter auction environments. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, 19–23 June 2017, pp. 856–868 (2017). https://doi.org/10.1145/3055399.3055427
19. Gonczarowski, Y.A., Weinberg, S.M.: The sample complexity of up-to-ε multidimensional revenue maximization. In: 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS (2018)
20. Gonczarowski, Y.A., Weinberg, S.M.: The sample complexity of up-to-ϵ multidimensional revenue maximization. J. ACM **68**(3), 15:1–15:28 (2021). https://doi.org/10.1145/3439722
21. Groves, T.: Incentives in teams. Econometrica **41**(4), 617–631 (1973)
22. Guo, C., Huang, Z., Tang, Z.G., Zhang, X.: Generalizing complex hypotheses on product distributions: auctions, prophet inequalities, and Pandora's problem. arXiv preprint arXiv:1911.11936 (2019)

23. Guo, C., Huang, Z., Zhang, X.: Settling the sample complexity of single-parameter revenue maximization. In: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, 23–26 June 2019, pp. 662–673 (2019). https://doi.org/10.1145/3313276.3316325
24. Guo, W., Jordan, M.I., Zampetakis, E.: Robust learning of optimal auctions. In: Ranzato, M., Beygelzimer, A., Dauphin, Y.N., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, 6–14 December 2021, Virtual, pp. 21273–21284 (2021)
25. Hartline, J.D., Taggart, S.: Sample complexity for non-truthful mechanisms. In: Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, 24–28 June 2019, pp. 399–416 (2019). https://doi.org/10.1145/3328526.3329632
26. Liu, S., Psomas, C.: On the competition complexity of dynamic mechanism design. In: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, 7–10 January 2018, pp. 2008–2025 (2018). https://doi.org/10.1137/1.9781611975031.131
27. Mansour, Y., Mohri, M., Schneider, J., Sivan, B.: Strategizing against learners in bayesian games. In: Loh, P., Raginsky, M. (eds.) Conference on Learning Theory, 2–5 July 2022, London, UK. Proceedings of Machine Learning Research, vol. 178, pp. 5221–5252. PMLR (2022). https://proceedings.mlr.press/v178/mansour22a.html
28. Mirrokni, V.S., Leme, R.P., Tang, P., Zuo, S.: Dynamic auctions with bank accounts. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9–15 July 2016, pp. 387–393 (2016). http://www.ijcai.org/Abstract/16/062
29. Mirrokni, V.S., Leme, R.P., Tang, P., Zuo, S.: Non-clairvoyant dynamic mechanism design. In: Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, 18–22 June 2018, p. 169 (2018). https://doi.org/10.1145/3219166.3219224
30. Mirrokni, V.S., Leme, R.P., Tang, P., Zuo, S.: Optimal dynamic auctions are virtual welfare maximizers. In: The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, 27 January–1 February 2019, pp. 2125–2132. AAAI Press (2019). https://doi.org/10.1609/aaai.v33i01.33012125
31. Myerson, R.B.: Optimal auction design. Math. Oper. Res. **6**(1), 58–73 (1981)
32. Nekipelov, D., Syrgkanis, V., Tardos, E.: Econometrics for learning agents. In: Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC 2015, pp. 1–18. ACM, New York (2015). https://doi.org/10.1145/2764468.2764522
33. Papadimitriou, C.H., Pierrakos, G., Psomas, C., Rubinstein, A.: On the complexity of dynamic mechanism design. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, 10–12 January 2016, pp. 1458–1475 (2016). https://doi.org/10.1137/1.9781611974331.ch100
34. Vickrey, W.: Counterspeculations, auctions, and competitive sealed tenders. J. Financ. **16**(1), 8–37 (1961)