

# Convergence Analysis of No-Regret Bidding Algorithms in Repeated Auctions

Zhe Feng,<sup>1</sup> Guru Guruganesh,<sup>2</sup> Christopher Liaw,<sup>3</sup> Aranyak Mehta,<sup>2</sup> Abhishek Sethi<sup>2</sup>

<sup>1</sup> Harvard University

<sup>2</sup> Google Research

<sup>3</sup> University of British Columbia

zhe\_feng@g.harvard.edu, gurug@google.com, cvliaw@cs.ubc.ca, aranyak@google.com, arsethi@google.com

## Abstract

The connection between games and no-regret algorithms has been widely studied in the literature. A fundamental result is that when all players play no-regret strategies, this produces a sequence of actions whose time-average is a coarse-correlated equilibrium of the game. However, much less is known about equilibrium selection in the case that multiple equilibria exist. In this work, we study the convergence of no-regret bidding algorithms in auctions. Besides being of theoretical interest, bidding dynamics in auctions is an important question from a practical viewpoint as well. We study repeated game between bidders in which a single item is sold at each time step and the bidder's value is drawn from an unknown distribution. We show that if the bidders use any mean-based learning rule then the bidders converge with high probability to the truthful pure Nash Equilibrium in a second price auction, in VCG auction in the multi-slot setting and to the Bayesian Nash equilibrium in a first price auction. We note mean-based algorithms cover a wide variety of known no-regret algorithms such as Exp3, UCB,  $\epsilon$ -Greedy etc. Also, we analyze the convergence of the individual iterates produced by such learning algorithms, as opposed to the time-average of the sequence. Our experiments corroborate our theoretical findings and also find a similar convergence when we use other strategies such as Deep Q-Learning.

## 1 Introduction

The connection between Learning and Games has proven to be a very innovative, practical, and elegant area of research (see, e.g., (Fudenberg and Levine 1998; Cesa-Bianchi and Lugosi 2006; Nisan et al. 2007)). Several fundamental connections have been made between Learning and Game Theory. A folklore result is that if all players play low-regret strategies in a repeated general-sum game, then the time-averaged history of the plays converges to a coarse correlated equilibrium (see, e.g., (Blum and Mansour 2007)). Similarly, low-swap-regret play leads to correlated equilibria.

In this work, we are interested in the setting of multi-agent learning in auction environments. Given the importance of auctions and bidding in the online advertising ecosystem, this is an important question from a practical point of view

as well. In this setting, an auction takes input bids and determines the allocation of ad-slots and the prices for each advertiser. This is a highly repeated auction setting over a huge number of queries which arrive over time. Advertisers get feedback on the allocation and cost achieved by their chosen bidding strategy, and can respond by changing their bids, and often do so in an automated manner (we will assume throughout that advertisers are *profit-maximizing*, although there are potentially other goals as well). As each advertiser responds to the feedback, it changes the cost landscape for every other competing advertiser via the auction rules. As advertisers respond to the auction and to each other, the dynamics lead them to an equilibrium. This results in the following fundamental question: Given a fixed auction rule, do such bidding dynamics settle in an equilibrium and if so, what equilibrium do they choose.

Surprisingly, neither the Auction Theory literature from Economics nor the literature on Learning in Games provides a definitive answer. Consider the simplest setting: a repeated auction for a single item, where the allocation is either first-price or second-price. Auction Theory suggests that bidders converge to *canonical equilibria* (see, e.g., (Krishna 2002)):

- For a second-price auction (or more generally a VCG auction), bidders will choose to be truthful (bid their true value every time) as this strategy weakly dominates every other strategy, i.e., no other strategy can yield more profit. This is a weakly dominating strategy Nash equilibrium (NE).
- For a first price auction in which each advertiser's value is picked from some commonly known distribution in an i.i.d. manner, each advertiser will underbid in a specific way to achieve a *Bayesian Nash Equilibrium* (BNE). For example, when there are two bidders and the value distribution is the uniform distribution on  $[0, 1]$ , then each advertiser will set its bid to be half of its true value.

While these canonical equilibria make intuitive sense, they are not the only equilibria in the respective games. For example, consider a single-item setting in which bidder 1 has a value of 1.0 and bidder 2 has a value of 0.5 for the item. While truthful bidding is an NE, any two values  $b_1, b_2 \in [0.5, 1.0]^2$ , with  $b_1 > b_2$  also form an NE (in fact, an Envy-Free NE). Thus, there are an infinite number of NEs, with very different revenues. Similarly, in the Bayesian set-

ting where the values are drawn from say, a uniform distribution, there are many NEs as well. For example, one player could bid 1 and the other player always bids 0 regardless of their valuations. This issue of multiple equilibria is treated in the Economics literature via various notions of equilibrium selection but it is not clear if such selection occurs naturally in settings such as ours, especially via bidding dynamics.

To take the Learning in Games approach to answering the question, we have to fix the bidding dynamics. We assume bidders use no-regret (mean-based) online learning algorithms; these are natural and powerful strategies that we may expect advertisers to use. A lot of commonly used no-regret learning algorithms, e.g. multiplicative weights update (MWU), follow the perturbed leader (FTPL), EXP3, UCB, and  $\epsilon$ -Greedy, are all special cases of mean-based no-regret learning algorithms.

Indeed, there has been considerable work recently which studies various questions in the online advertising setting under this assumption (see, e.g., (Nekipelov, Syrgkanis, and Tardos 2015)). Folklore results in Learning imply that under low-regret dynamics the time-average of the advertisers' bids will converge to a *coarse correlated equilibrium (CCE)* of the underlying game. Hartline, Syrgkanis, and Tardos (2015) shows that no-regret learning in Bayesian games converges to Bayesian CCE in the time-average manner. However, there could be many CCEs in a game as well. Since every NE is a CCE as well, the above examples hold for the second-price auction. For the first price setting (even when the values are drawn uniformly at random) another CCE is for the two bidders to bid  $(v_1 + v_2)/2 + \epsilon$  and  $(v_1 + v_2)/2$  where  $v_1$  and  $v_2$  are the drawn values, and  $v_1 > v_2$ . Further, since any convex combination of these CCEs is a valid CCE (the set of CCE forms a polytope), there is an infinite number of CCEs in both first and second price auctions. So again we are left with the prediction dilemma: it is not clear which of these CCEs a low-regret algorithm will converge to. Some of the CCEs are clearly not good for the ecosystem, for revenue or efficiency. Further, as we elaborate below, even if there is a unique CCE in the games of interest, the convergence guarantee is only for the time-average rather than point-wise.

## Questions

These examples directly motivate the questions we ask in this paper:

- If the values are drawn in a bayesian setting, and the bidders follow a low-regret learning algorithm in a repeated second-price auction, do they converge to the truthful equilibrium?
- Similarly, in a first price auction with i.i.d. values with bidders values drawn from a uniform distribution  $[0, 1]$ , do they converge to the Bayesian Nash equilibrium (with two bidders) of bidding half of true value?
- Do the bidding dynamics converge to such an equilibrium point-wise, i.e., in the last-iterate sense, or only in the time-average sense?
- When there are multiple slots, do the bidders converge to truthful equilibrium under VCG settings?

Given the current state of the literature, we see these as fundamental questions to ask. The only guarantees we have are those known for general games: Low-regret dynamics converge to some CCE and there is no guarantee for the last-iterate convergence. If it is the case that only the time average converges, then that means that bidders may be required to keep changing their bidding strategy at every time step (see the discussion on the non-point-wise-convergence results in (Bailey and Piliouras 2018) below), and would achieve very different rewards over time. This would not be a satisfactory situation in the practical setting.

## Our Results

Our main result is that when each of the bidders use a mean-based learning rule (see Definition 2.2) then all bidders converge to truthful bidding in a second-price auction and a multi-position VCG auction and to the Bayes Nash Equilibrium in a first-price auction.

**Informal Main Theorem.** *Suppose  $n$  players whose value are drawn from a distribution over space  $\{\frac{1}{H}, \frac{2}{H}, \dots, 1\}$  bid according to a mean-based learning algorithm in either (i) a second-price auction; (ii) a first-price auction (for  $n = 2$  and uniform value distribution); or (iii) a multi-position VCG auction. Then, after an initial exploration phase, each bidder bids the canonical Bayes Nash equilibrium with high probability.*

The formal statement of this theorem appears in Theorem 3.1, Theorem 4.1, and Theorem 4.3 for second-price auctions, first-price auctions, and multi-position VCG auctions, respectively. Moreover, we show each bidder converges to bid canonical Nash Equilibrium point-wise for any time, which is in sharp contrast with previous time-average convergence analysis.

Throughout this paper, we assume the learning algorithms that the bidders use may be oblivious to the auction format that is used by the seller.

We complement these results by simulating the above model with experiments. In particular, we show that these algorithms converge and produce truthful (or the canonical) equilibria. Furthermore, we show that the algorithm converges much quicker than the theory would predict. These results indicate that these low-regret algorithms are quite robust in the context of Auctions.

## Our Techniques

Our proof techniques involve a few key observations. For second-price and VCG auctions, we want to show that the bidding converges to truthful bidding. Firstly, note that if the other bidders bid completely randomly then the truthful arms have a slight advantage in profit. However, the other bidders themselves bid according to their own instance of the low-regret algorithms, thus the environment that a given bidder sees is not necessarily random; hence we need more insight. Fix a particular bidder, say bidder 1. In the beginning of the learning algorithm, all bidders do bid randomly. The next observation is that if the other bidders *happened* to converge to truthful bidding, then again bidder 1 will see completely random bids, because the other bidders' values

are picked randomly at each stage. Hence we can say that both in the beginning and also if other bidders happen to converge to (or for some reason were restricted to) truthful bidding, then bidder 2 will see an advantage in truthful bidding and converge to that. It remains to show that in the interim period, when all bidders are learning, exploring, and exploiting, the truthful strategy builds and retains an advantage.

In a first price auction, the proofs follow the same structure. However there are some technical difficulties that one must overcome. Initially, both bidders bid uniformly at random and it is not difficult to show that bidding according to the canonical NE gives an advantage. If the bidders happen to converge to the BNE then, of course, bidding according to the BNE is the optimal strategy for either player. It is not clear, however, that when the opposing bid is *not* uniform or bidding according to the BNE that an advantage is maintained. Our technical contribution here is to show that an advantage for bidding the BNE is maintained which allows both bidders to converge to the BNE.

Our results show that we can achieve high probability results we can show that the model will bid truthfully *for all time* (assuming a modest exploratory period in the beginning). This requires a new partitioning argument which enables us to apply concentration results for all times. This technique may be of independent interest.

## Related Works

Our work lies in a wide area of the inter-disciplinary research between mechanism design and online learning algorithms, e.g., (Auer et al. 1995; Cesa-Bianchi and Lugosi 2006; Blum and Mansour 2007), and we only point out a few lines of research which are more closely related to our work.

In online advertising, the setting where bidders may be running no-regret algorithms rather than best response, has recently been investigated and has garnered a significant amount of interest. For example, Nekipelov, Syrgkanis, and Tardos (2015) study how to infer advertisers’ values under this assumption. However, Bailey and Piliouras (2018) show, somewhat surprisingly, that even in very simple games (e.g., a zero-sum matching pennies game), the no-regret dynamics of MWU do *not* converge, and in fact the individual iterates tend to *diverge* from the CCE. On the other hand, recent results in (Daskalakis and Panageas 2018; Mertikopoulos et al. 2019) show that certain Optimistic variants of Gradient Descent and Mirror Descent converge in the last-iterate to the NE in certain zero-sum games. Our result can be seen as a contribution in this stream of work as well, in that we show that for the (non-zero sum) games arising from auctions that we study, mean-based learning algorithms converge in the last iterate, and, in fact, to the natural (Bayes) NE. On a related note, Papadimitriou and Piliouras (2019) shows there is a conflict between the economic solution concepts and those predicted by Learning Dynamics. In that framework, one can consider this work as suggesting that perhaps there is no such conflict between economic theory and learning, in the context of games arising from auctions, as learning converges to the solutions predicted by auction theory.

Our work is also related with *Learning to bid* literature, e.g., (Weed, Perchet, and Rigollet 2016; Feng, Podimata, and Syrgkanis 2018; Balseiro et al. 2019), where these papers focus on designing a good learning algorithm for the bidders in repeated auctions. In addition, Braverman et al. (2018) considers how to design a mechanism to maximize revenue against bidders who adopt mean-based learning algorithms. In contrast, the auctions are fixed in our setting and we are interested in understanding the bidder dynamics.

Last but not least, Feldman, Lucier, and Nisan (2016) characterize multiple equilibria (NE, CE, and CCE) in first price auctions, under the prior-free (non-Bayesian) setting, and study the revenue and efficiency properties of these equilibria. They show there are auctions in which a CCE can have as low as  $1 - 2/e \simeq 0.26$  factor of the second highest value (although not lower), and there are auctions in which a CCE can have as low as 0.81 of the optimal efficiency (but not lower). However, our results show that even though there may be “Bad” CCEs, the natural dynamics do not reach them, and instead, converge to the classic canonical Nash equilibrium.

## 2 Model and Notations

We consider the setting that there is a single seller repeatedly selling one good to  $n$  bidders per round. At each time  $t$ , each bidder  $i$ ’s valuation  $v_{i,t}$  is i.i.d. drawn from an unknown (CDF) distribution  $F_i$  and bidder  $i$  will submit a bid  $b_{i,t}$  based on  $v_{i,t}$  and historical information. In this paper, we assume the value and the bid of each bidder at any time are always in a  $\frac{1}{H}$ -evenly-discretized space  $V = \{\frac{1}{H}, \frac{2}{H}, \dots, 1\}$ , i.e.,  $v_{i,t}, b_{i,t} \in V, \forall i, t$ . Let  $v_t = (v_{1,t}, \dots, v_{n,t})$  be the valuation profile of  $n$  bidders at time  $t$ ,  $v_{-i,t}$  be the valuation profile of bidders other than  $i$ , and similarly for  $b_t$  and  $b_{-i,t}$ . Let  $F$  be the (CDF) distribution of  $v_t$  and  $f_i$  be the probability density function (PDF) of bidder  $i$ ’s value. Denote  $m_{i,t} = \max_{j \neq i} b_{j,t}$  as the maximum bid of the bidders other than  $i$  and  $z_{i,t} = \max_{j \neq i} v_{j,t}$  be the maximum value of the bidders other than  $i$ . We denote  $G_i$  as the (CDF) distribution of  $z_{i,t}$  and  $g_i$  as the associated PDF. For theoretical purpose, we propose an assumption about  $G_i$  in the following,

**Assumption 2.1** (Thickness Assumption of  $G_i$ ). *There exists a constant  $\tau > 0$  (may depend on  $n$ ), s.t.,  $g_i(v) \geq \tau, \forall i \in [n], v \in V$ . Without loss of generality<sup>1</sup>, we assume  $\tau \leq \frac{1}{H^{n-1}}$ .*

We assume the each bidder runs a no-regret learning algorithm to decide her bid at each time. Specifically, in this paper, we are interested in a broad class of no-regret learning algorithm known as mean-based (contextual) learning algorithm (Braverman et al. 2018); these include the multiplicative weights update (MWU), Exp3, and  $\varepsilon$ -Greedy algorithms as special cases.

In this paper, we focus on the contextual version of mean-based learning algorithms, which can be used to model learning algorithms of the bidders in repeated auctions, defined in the following.

<sup>1</sup>It is without loss of generality, since if  $\tau > \frac{1}{H^{n-1}}$ , we redefine  $\tau := \min \{ \tau, \frac{1}{H^{n-1}} \}$ .

**Definition 2.2** (Mean-based Contextual Learning Algorithm). Let  $r_{a,t}(c)$  be the reward of action  $a$  at time  $t$  when the context is  $c$  and  $\sigma_{a,t}(c) = \sum_{s=1}^t r_{a,s}(c)$ . An algorithm for the contextual bandits problem is  $\gamma_t$ -mean-based if it is the case that whenever  $\sigma_{a,t}(c) < \sigma_{b,t}(c) - \gamma_t t$ , then the probability  $p_{a,t}(c)$  that the algorithm plays action  $a$  on round  $t + 1$ , given context  $c$ , is at most  $\gamma_t$ . We say an algorithm is mean-based if it is  $\gamma_t$ -mean-based for some  $\gamma_t$  such that  $\gamma_t t$  is increasing and  $\gamma_t \rightarrow 0$  as  $t \rightarrow \infty$ .<sup>2</sup>

In the repeated auctions setting, the context information received by each bidder  $i$  at time  $t$  is the realization of the valuation  $v_{i,t}$ . The reward function  $r_{b,t}^i$  for bidder  $i$  can be defined as

$$\forall v \in [0, 1], r_{b,t}^i(v) := u_{i,t}((b, b_{-i,t}); v), \quad (1)$$

where  $u_{i,t}((b, b_{-i,t}); v)$  is the utility of bidder  $i$  at time  $t$  when the bidder  $i$  bids  $b$  and the others bid  $b_{-i,t}$ , if bidder  $i$  values the good  $v$ .

### Learning Algorithms of Mean-Based Bidders

In this paper, we focus on the setting where each bidder  $i$  runs a  $\gamma_t$ -means-based contextual learning algorithm to submit the bid<sup>3</sup>. In addition, we assume each bidder runs several pure exploration steps in the beginning to estimate the reward of each action (bid) for each context (value). We assume each bidder runs  $T_0$  pure exploration steps: at each pure exploration step, each bidder  $i$  uniformly generates a bid from  $B$  at random, regardless of the realization of value  $v_{i,t}$ . To summarize, we describe the learning algorithm of mean-based bidders in Algorithm 1.

---

**Algorithm 1** Mean-based (Contextual) Learning Algorithm of Bidder  $i$

---

```

1: procedure MBL( $\gamma_t, T_0$ )
2:   for  $t = 1, 2, \dots, T_0$  do
3:     Choose bid  $b_{i,t}$  uniformly from  $V$  at random.
4:   end for
5:   for  $t = T_0 + 1, T_0 + 2, \dots$  do
6:     Observes value  $v_t$ .
7:     Choose bid  $b_{i,t}$  following a  $\gamma_t$ -mean-based
      learning algorithm.
8:   end for
9: end procedure

```

---

In the *learning to bid* literature, there are different feedback models: full information feedback (Papadimitriou and Piliouras 2019), bandit feedback (Weed, Perchet, and Rigollet 2016; Feng, Podimata, and Syrgkanis 2018), or cross-learning feedback (Balseiro et al. 2019). However, our results hold for any feedback model, as long as each bid-

<sup>2</sup>The mean-based learning algorithms proposed by (Braverman et al. 2018) set  $\gamma_t$  be a constant, which only depends on total number of time steps  $T$ . Here we extend it to be a time-dependent variable, which is used to show our *anytime* convergence results.

<sup>3</sup>Indeed, our analysis can be extended to the setting where each mean-based bidder has different  $\gamma_t$  parameters. We assume they share parameters  $\gamma_t$ , for notation simplicity.

der uses the general mean-based learning algorithm to bid, shown in Algorithm 1.

## 3 Second Price Auctions with Mean-based Bidders

In this section, we analyze the learning dynamics of mean-based bidders in (repeated) second price auctions. In second price auctions, the utility function of each bidder  $i$  at time  $t$  can be represented as,

$$u_{i,t}((b, b_{-i,t}); v) = (v - m_{i,t}) \cdot \mathbb{I}\{b \geq m_{i,t}\} \quad (2)$$

Since the bids of each bidder are in a discrete space, we break ties randomly throughout this paper. We first show the following main theorem in this section, which proves that the mean-based learners converge to truthful reporting pointwisely, in the repeated second price auctions.

**Theorem 3.1.** Suppose assumption 2.1 holds and  $T_0$  is large enough, such that  $\exp\left(-\frac{\tau^2 T_0}{32n^2 H^2}\right) \leq \frac{1}{2}$  and  $\gamma_t \leq \frac{\tau}{8nH}, \forall t \geq T_0$ . Then at time  $t > T_0$ , each  $\gamma_t$ -mean-based learner  $i$  will submit  $b_t = v_{i,t}$  in repeated second price auctions with probability at least  $p(t) = 1 - H\gamma_t - 4 \exp\left(-\frac{\tau^2 T_0}{32n^2 H^2}\right)$ , for any fixed  $v_{i,t}$ . Note  $p(t) \rightarrow 1 - 4 \exp\left(-\frac{\tau^2 T_0}{32n^2 H^2}\right)$  when  $t \rightarrow \infty$ .

Our main results for second price auctions show an *anytime* convergence for each bidder: as long as  $T_0$  is large enough, each bidder will bid truthfully at any time  $t$ , with high probability. The main technical contribution in this paper is the proof for Theorem 3.1.

### Proof of Theorem 3.1

In this section, we summarize the proof of Theorem 3.1. To show that, we propose the following lemmas, in which the complete proofs are deferred to Appendix B.

Firstly, we characterize that in the pure exploration phase, each bidder gains significantly greater utility when bidding truthfully.

**Lemma 3.2.** For any fixed value  $v$ , any bid  $b \neq v$  and any time  $t \leq T_0$ , we have for each bidder  $i$ ,

$$\mathbb{P}\left(u_{i,t}((v, b_{-i,t}); v) - u_{i,t}((b, b_{-i,t}); v) \geq \frac{1}{H}\right) \geq \frac{\tau}{n}$$

Then by a standard Chernoff bound, we can argue the accumulative utility advantage obtained by bidding truthfully will be large enough, for any time  $t$  in exploration phase.

**Lemma 3.3.** For any fixed  $v$ , any bid  $b \neq v$  and any time  $t \leq T_0$ , we have for each bidder  $i$ ,

$$\sum_{s \leq t} u_{i,s}((v, b_{-i,s}); v) - u_{i,s}((b, b_{-i,s}); v) \geq \frac{\tau t}{2nH}$$

holds with probability at least  $1 - \exp\left(-\frac{\tau^2 t}{2n^2 H^2}\right)$ .

Finally, we show that if the cumulative utility advantage of truthful bidding is large enough to satisfy requirements of the mean-based learning algorithms, then truthful bidding still gains significant greater utility for each bidder at time  $t > T_0$ .

**Lemma 3.4.** For any  $t > T_0$ , suppose  $\sum_{s \leq t} u_{i,s}((v, b_{-i,s}); v) - u_{i,s}((b, b_{-i,s}); v) \geq \gamma t$  holds for any fixed  $v, b \neq v$  and each bidder  $i$ , then

$$u_{i,t+1}((v, b_{-i,t+1}); v) - u_{i,t+1}((b, b_{-i,t+1}); v) \geq \frac{1}{H}$$

holds with probability at least  $\frac{\tau}{2n}$ , for any fixed value  $v$ , bid  $b \neq v$  and each bidder  $i$ .

Given the above three auxiliary lemmas, we prove Theorem 3.1 in the following.

*Proof of Theorem 3.1.* One of the key techniques used in this paper is the partitioning of the time steps into buckets with a geometric partitioning scheme. In particular, we divide time steps  $t > T_0$  to several episodes as follows,  $\Gamma_1 = [T_0 + 1, T_1]$ ,  $\Gamma_2 = [T_1 + 1, T_2]$ , ..., such that  $\forall k \geq 1$ ,  $T_k = \lfloor \frac{\tau T_{k-1}}{4\gamma T_k n H} \rfloor$ . We always choose the smallest  $T_k$  to satisfy this condition.<sup>4</sup> The total time steps of each episode  $|\Gamma_k| = T_k - T_{k-1}$ ,  $\forall k \geq 1$ . Then we show the following claim, which states that in each time bucket the expected utility doesn't deviate too much.

**Claim 3.5.** Let event  $\mathcal{E}_k$  be  $\sum_{s \leq T_k} u_{i,s}((v, b_{-i,s}); v) - u_{i,s}((b, b_{-i,s}); v) \geq \frac{\tau T_k}{4nH}$  holds for all  $i$ , given any fixed  $v, b \neq v$ . Then the event  $\mathcal{E}_k$  holds with probability at least  $1 - \sum_{\ell=0}^k \exp\left(-\frac{|\Gamma_\ell| \tau^2}{32n^2 H^2}\right)$ .

We prove the above claim by induction. If  $k = 0$ , the claim holds by Lemma 3.3. We assume the claim holds for  $k$ , then we argue the claim still holds for  $k + 1$ . We consider any time  $t \in \Gamma_{k+1}$ , given event  $\mathcal{E}_k$  holds, we have

$$\begin{aligned} & \sum_{s \leq t} u_{i,s}((v, b_{-i,s}); v) - u_{i,s}((b, b_{-i,s}); v) \\ & \geq \sum_{s \leq T_k} u_{i,s}((v, b_{-i,s}); v) - u_{i,s}((b, b_{-i,s}); v) \\ & \geq \frac{\tau T_k}{4nH} \geq \gamma t, \end{aligned} \quad (3)$$

where the first inequality is based on the fact that truth-telling is the dominant strategy in second price auctions, the second inequality holds because of the induction assumption and the last inequality hold because  $\forall t \in \Gamma_{k+1}$ ,

$$\gamma t \leq \gamma T_{k+1} T_{k+1} = \gamma T_{k+1} \left\lfloor \frac{\tau T_k}{4\gamma T_{k+1} n H} \right\rfloor \leq \frac{\tau T_k}{4nH}.$$

Then by Lemma 3.4, given  $\mathcal{E}_k$  holds, for any  $t \in \Gamma_{k+1}$  we have,

$$\mathbb{P}\left(u_{i,t}((v, b_{-i,t}); v) - u_{i,t}((b, b_{-i,t}); v) \geq \frac{1}{H} \middle| \mathcal{E}_k\right) \geq \frac{\tau}{2n}$$

Thus,  $\mathbb{E}\left[u_{i,t}((v, b_{-i,t}); v) - u_{i,t}((b, b_{-i,t}); v) \middle| \mathcal{E}_k\right] \geq \frac{\tau}{2nH}$  for any  $t \in \Gamma_{k+1}$ . Letting  $\Delta_s = u_{i,s}((v, b_{-i,s}); v) - u_{i,s}((b, b_{-i,s}); v)$ , by Azuma's inequality (for martingales),

<sup>4</sup> $T_k$  always exists since  $\gamma t \rightarrow \infty$  as  $t \rightarrow \infty$ .

we have

$$\begin{aligned} & \mathbb{P}\left(\sum_{s \in \Gamma_{k+1}} \Delta_s \leq \frac{\tau |\Gamma_{k+1}|}{4nH} \middle| \mathcal{E}_k\right) \\ & \leq \mathbb{P}\left(\sum_{s \in \Gamma_{k+1}} \Delta_s \leq \sum_{s \in \Gamma_{k+1}} \mathbb{E}\left[\Delta_s \middle| \mathcal{E}_k\right] - \frac{\tau |\Gamma_{k+1}|}{4nH} \middle| \mathcal{E}_k\right) \\ & \leq \exp\left(-\frac{|\Gamma_{k+1}| \tau^2}{32n^2 H^2}\right) \end{aligned}$$

Therefore, the event  $\mathcal{E}_{k+1}$  holds with probability at least

$$\left(1 - e^{-\frac{|\Gamma_{k+1}| \tau^2}{32n^2 H^2}}\right) \cdot \mathbb{P}(\mathcal{E}_k) \geq 1 - \sum_{\ell=0}^{k+1} \exp\left(-\frac{|\Gamma_\ell| \tau^2}{32n^2 H^2}\right),$$

which completes the induction and verifies the correctness of Claim 3.5. Given Claim 3.5, we have the following argument,

For any time  $t > T_0$ , there exists  $k(t)$ , s.t.,  $t \in \Gamma_{k(t)}$ , if the event  $\mathcal{E}_{k(t)}$  happens, the bidder  $i$  will report truthfully with probability at least  $1 - H\gamma_t$ , by the definition of  $\gamma_t$ -mean-based learning algorithms and the same argument as Eq. (3). Therefore, at any time  $t > T_0$ , each bidder  $i$  will report truthfully with probability at least

$$1 - H\gamma_t - \sum_{\ell=0}^{k(t)} \exp\left(-\frac{|\Gamma_\ell| \tau^2}{32n^2 H^2}\right)$$

Then we bound  $\sum_{\ell=0}^{k(t)} \exp\left(-\frac{|\Gamma_\ell| \tau^2}{32n^2 H^2}\right)$  through the following claim, where the proof is deferred to Appendix B.

**Claim 3.6.** Given  $\gamma_t \leq \frac{\tau}{8nH}$ ,  $\forall t > T_0$ ,  $\sum_{\ell=0}^{k(t)} \exp\left(-\frac{|\Gamma_\ell| \tau^2}{32n^2 H^2}\right) \leq 4 \exp\left(-\frac{\tau^2 T_0}{32n^2 H^2}\right)$ , when  $T_0$  is large enough s.t.  $\exp\left(-\frac{\tau^2 T_0}{32n^2 H^2}\right) \leq \frac{1}{2}$ .

Combining the above claim, we complete the proof for Theorem 3.1.  $\square$

## 4 Generalizations to Other Auctions

In this section, we generalize our results further to first price auctions when each bidder's valuation is drawn uniformly from  $V$ , and multi-position auctions when we run the Vickrey-Clarke-Groves (VCG) mechanism. For both cases, we break ties randomly.

### First Price Auctions

In first price auctions, the highest bidder wins and pays her bid. Therefore the utility function of each bidder  $i$  at time  $t$  can be defined as,

$$u_{i,t}((b, b_{-i,t}); v) = (v - b) \cdot \mathbb{I}\{b \geq m_{i,t}\} \quad (4)$$

It is well-known, first price auctions are not truthful and each bidder will underbid her value to manipulate the auctions. Bayesian Nash Equilibrium (BNE) bidding strategy is hard to characterize in general first price auctions. Only the

BNE bidding strategy for i.i.d bidders in first price auctions, is fully understood, e.g. (Krishna 2002).

In this paper, for simplicity, we focus on the setting that there are two i.i.d bidders with uniform value distribution over  $V$ . The BNE bidding strategy for each bidder is  $b = \frac{v}{2}$  when the value is  $v$  (Krishna 2002), if the value space  $V = [0, 1]$ . In this work, we assume  $V = \{\frac{1}{H}, \frac{2}{H}, \dots, 1\}$  and we break ties randomly. Under this case, we show each mean-based bidder will converge to play near-BNE bidding strategy point-wisely if the number of initial exploration steps  $T_0$  is large enough in the following.

**Theorem 4.1.** *Suppose there are two bidders and each bidder's value is i.i.d drawn uniformly from  $V$  at random,  $H$  is a even positive number, and  $T_0$  is large enough, such that  $\gamma_t \leq \frac{1}{4H^3}, \forall t \geq T_0$ . Then at time  $t > T_0$ , each  $\gamma_t$ -mean-based learner  $i$  in repeated first price auctions will bid  $b_t$ , s.t.  $\frac{v}{2} \leq b_t < \frac{v}{2} + \frac{1}{H}$  with probability at least  $1 - H\gamma_t - \exp\left(-\frac{(H-1)T_0}{32(4H^3+1)H^4}\right) \frac{\log t}{\log\left(\frac{4H^3+H}{4H^3+1}\right)}$ , for any fixed  $v_{i,t}$ .*

The proof is significantly different than the proof of second price auctions. Firstly, random tie-breaking makes the analysis more difficult in first price auctions compared with the one in second price auctions. Secondly,  $b = \lceil \frac{v}{2} \rceil^5$  is not a dominant bidding strategy, we need a more complex time-splitting scheme to make induction procedure works in first price auctions. We defer the proof to Appendix B. We believe our proof for first price auctions is general enough and can be extended to handle more than two bidders setting.

## Multi-Position VCG Auctions

In multi-position auctions, there are  $k$  positions where  $k < n$  and position multipliers  $p_1 \geq \dots \geq p_k \geq 0 = p_{k+1} = \dots = p_n$  (position multipliers determine the relative values or click-through rates of the different positions). In particular, we will also say that  $p_i - p_{i+1} \geq \rho$  for all  $i \leq k$ . In this paper, we run VCG mechanism for multi-position auctions. For more details of multi-position VCG auctions, see Appendix A.

Denote  $z_i^{(k)}$  be the  $k$ -th largest value from the bidders other than bidder  $i$ ,  $G_i^{(k)}$  be the (CDF) distribution of  $z_i^{(k)}$ , and  $g_i^{(k)}$  be the associated PDF. We propose the following thickness assumption for distribution  $G_i^{(k)}$ , for our theoretical purpose.

**Assumption 4.2** (Thickness Assumption of  $G_i^{(k)}$ ). *There exists a constant  $\tau > 0$  (may depend on  $n$ ), s.t.,  $g_i^{(k)}(v) \geq \tau, \forall i \in [n], \forall v \in V$ . Without loss of generality, we assume  $\tau \leq \frac{1}{H^{n-1}}$ .*

It is well-known that the VCG auction is truthful. Applying the same technique, we can show the truthful bid has an advantage compared with all the other bids in the exploration phase, as well as in the mean-based exploitation phase. Similarly, we show the following convergence results of mean-based bidders in multi-position VCG auctions.

<sup>5</sup>  $\lceil x \rceil$  means rounding  $x$  up to the nearest point in  $V$ .

**Theorem 4.3.** *Suppose assumption 4.2 holds and  $T_0$  is large enough, such that  $\exp\left(-\frac{\tau^2 \rho^2 T_0}{32n^2 H^2}\right) \leq \frac{1}{2}$  and  $\gamma_t \leq \frac{\tau \rho}{8nH}, \forall t \geq T_0$ . Then at time  $t > T_0$ , each  $\gamma_t$ -mean-based learner  $i$  will  $b_t = v_{i,t}$  in repeated multi-position VCG auctions with probability at least  $1 - H\gamma_t - 4 \exp\left(-\frac{\tau^2 \rho^2 T_0}{32n^2 H^2}\right)$ , for any fixed  $v_{i,t}$ .*

## 5 Experiments

In this section, we describe the experiments using Contextual Mean-Based Algorithms and Deep-Q Learning agents participating in repeated first price and second price auctions. In these repeated auctions, the private valuations of both players are sampled independently from identical uniform distributions. The observation for each agent is defined by its private valuation and its reward by the auction outcome (with ties broken randomly).

In the first set of experiments, we study the convergence of two independent learning agents following an  $\epsilon$ -Greedy policy in first and second price auction. We use the setting of  $H = 10$  wherein both agents only observe their private valuation and the respective reward as an auction outcome.

In both cases, we observe the bidders converge to the BNE after several time steps. There is a slight gap between the (observed) mean reward and utility under BNE as the value of  $\epsilon$  (randomly exploration probability) has a floor of 0.05. We also observe that in the exploitation phase, the bidding converges completely to the BNE in the contextual bandit setting, which exactly matches our theory in Figures 1a and 2a. More experiments are shown in Appendix C.

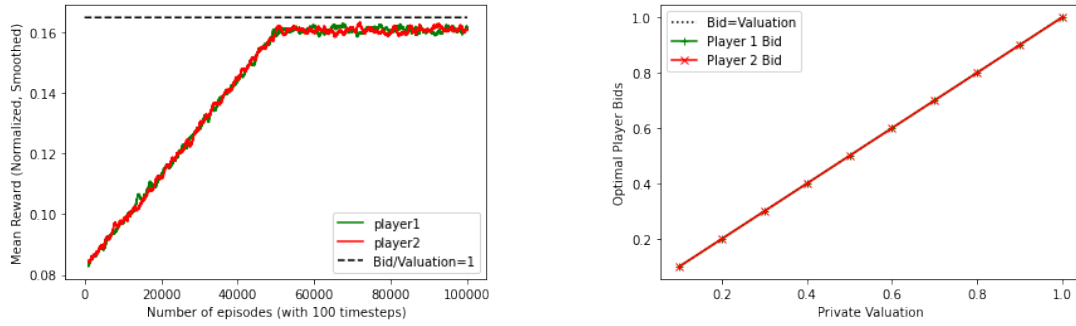
## Extensions to Deep-Q Learning

Contextual Mean Based Algorithms are a broad class of algorithms but can be very expensive to implement if we run a new instance for each possible value and the number of values are large. In line with modern machine learning, one way to mitigate this in practice is to augment it via Deep Q-Learning. To be more concrete, we model the learner by using a deep network with input as the private value and ask it to choose one of many bids. We model this as a reinforcement learning problem where the agents state is input into a deep neural network. The agent's rewards are then observed over time with a chosen discount rate. The details of Deep Q-Learning model and the set of hyperparameters used to train the two Q models are outlined in Appendix C.


We use the setting of  $H = 100$  and consider the observation of the agent as its private valuation. Again, we observe that both agents converge to BNE, shown in Figures 1b and 2b. We also study the model with a wider set of states including previously chosen bids and empirically observe the convergence of independent DQN agents to BNE for both auctions (discussed further in Appendix C).

## Acknowledgments

Zhe Feng is supported by a Google PhD Fellowship. We would like to thank anonymous reviewers for their helpful feedback.

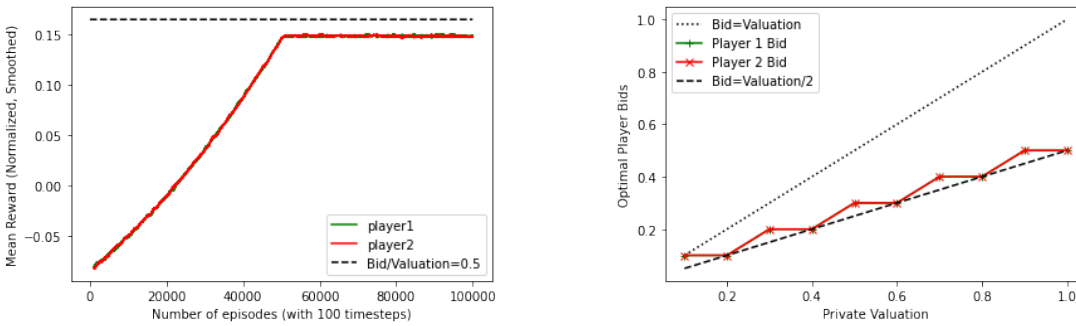


(a) Training curve of mean reward of each bidder (left) and roll-out bidding strategy of each bidder (right) in the exploitation phase of contextual  $\epsilon$ -Greedy algorithm in second price auctions.

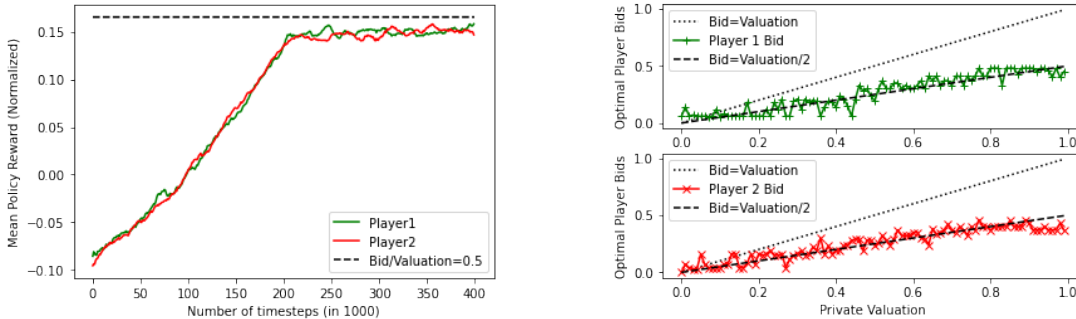


(b) Training curve of mean reward of each bidder (left) and roll-out bidding strategy of each bidder (right) in the exploitation phase of Deep-Q Learning algorithm in second price auctions.

Figure 1: Simulation results for second price auctions



(a) Training curve of mean reward of each bidder (left) and roll-out bidding strategy of each bidder (right) in the exploitation phase of contextual  $\epsilon$ -Greedy algorithm in first price auctions.



(b) Training curve of mean reward of each bidder (left) and roll-out bidding strategy of each bidder (right) in the exploitation phase of Deep-Q Learning algorithm in first price auctions.

Figure 2: Simulation Results for first price auctions

## References

- Auer, P.; Cesa-Bianchi, N.; Freund, Y.; and Schapire, R. E. 1995. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of IEEE 36th Annual Foundations of Computer Science*, 322–331.
- Bailey, J. P.; and Piliouras, G. 2018. Multiplicative Weights Update in Zero-Sum Games. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 321–338. ACM.
- Balseiro, S.; Golrezaei, N.; Mahdian, M.; Mirrokni, V.; and Schneider, J. 2019. Contextual Bandits with Cross-Learning. In *Advances in Neural Information Processing Systems 32*, 9679–9688.
- Blum, A.; and Mansour, Y. 2007. Learning, Regret Minimization, and Equilibria. In *Algorithmic Game Theory*, 79–102. Cambridge University Press.
- Braverman, M.; Mao, J.; Schneider, J.; and Weinberg, M. 2018. Selling to a No-Regret Buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 523–538. ACM.
- Cesa-Bianchi, N.; and Lugosi, G. 2006. *Prediction, learning, and games*. Cambridge University Press.
- Daskalakis, C.; and Panageas, I. 2018. Last-Iterate Convergence: Zero-Sum Games and Constrained Min-Max Optimization. In *10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Feldman, M.; Lucier, B.; and Nisan, N. 2016. Correlated and Coarse equilibria of Single-item auctions. In *International Conference on Web and Internet Economics*, 131–144. Springer.
- Feng, Z.; Podimata, C.; and Syrgkanis, V. 2018. Learning to Bid Without Knowing Your Value. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 505–522.
- Fudenberg, D.; and Levine, D. K. 1998. *The theory of learning in games*. MIT Press.
- Hartline, J.; Syrgkanis, V.; and Tardos, E. 2015. No-Regret Learning in Bayesian Games. In *Advances in Neural Information Processing Systems*, volume 28, 3061–3069. Curran Associates, Inc.
- Krishna, V. 2002. *Auction Theory*. Elsevier Science.
- Mertikopoulos, P.; Zenati, H.; Lecouat, B.; Foo, C.-S.; Chandrasekhar, V.; and Piliouras, G. 2019. Optimistic Mirror Descent in Saddle-Point Problems: Going the Extra (Gradient) Mile. In *ICLR '19 - International Conference on Learning Representations*, 1–23. New Orleans, United States.
- Nekipelov, D.; Syrgkanis, V.; and Tardos, E. 2015. Econometrics for learning agents. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, 1–18. ACM.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic game theory*. Cambridge University Press.
- Papadimitriou, C.; and Piliouras, G. 2019. Game Dynamics As the Meaning of a Game. *SIGecom Exch.* 16(2): 53–63. ISSN 1551-9031.
- Weed, J.; Perchet, V.; and Rigollet, P. 2016. Online learning in repeated auctions. In *29th Annual Conference on Learning Theory*, 1562–1583.