

# No-Regret Learning from Partially Observed Data in Repeated Auctions <sup>★</sup>

Orcun Karaca <sup>\*,★★</sup> Pier Giuseppe Sessa <sup>\*,★★</sup> Anna Leidi <sup>\*</sup> and Maryam Kamgarpour <sup>\*</sup>

*\* Automatic Control Laboratory, ETH Zürich, Switzerland.  
(e-mails: {okaraca, sessap, leidia, maryamk}@ethz.ch).*

---

**Abstract:** We study a general class of repeated auctions, such as the ones found in electricity markets, as multi-agent games between the bidders. In such a repeated setting, bidders can adapt their strategies online using no-regret algorithms based on the data observed in the previous auction rounds. Well-studied no-regret algorithms depend on the feedback information available at every round, and can be mainly distinguished as bandit (or payoff-based), and full-information. However, the information structure found in auctions lies in between these two models, since participants can often obtain partial observations of their utilities under different strategies. To this end, we modify existing bandit algorithms to exploit such additional information. Specifically, we utilize the feedback information that bidders can obtain when their bids are not accepted, and build a more accurate estimator of the utility vector. This results in improved regret guarantees compared to standard bandit algorithms. Moreover, we propose a heuristic method for auction settings where the proposed algorithm is not directly applicable. Finally, we demonstrate our findings on case studies based on realistic electricity market models.

*Keywords:* auctions; game theory; no-regret; coarse-correlated equilibrium; electricity markets.

---

## 1. INTRODUCTION

Auctions are effective tools for allocating resources and determining their values among a set of participants. To this end, many complex auction mechanisms have been proposed to buy and sell different resources. Prominent examples include those in spectrum auctions (Bichler and Goeree, 2017), electricity markets (Cramton, 2017), auctions for fish harvesting rights and other natural capitals (Teytelboym, 2019). Moreover, a periodic *recurrence* is intrinsic to several applications. For instance, in electricity markets, the same participants are generally involved in similar market transactions day after day.

Repeated auctions can be studied as multi-agent games among different players, or bidders in the auction, equipped with adaptive and sequential algorithms. In such repeated games, the performance of a player is typically measured in terms of *regret*, which is the utility loss incurred when compared to the best fixed action over a sequence of rounds. Moreover, if all players exhibit no-regret, that is, a diminishing average regret as the number of rounds increases, the game is known to reach a so-called coarse-correlated equilibrium. This equilibrium concept generalizes Nash equilibrium to the case where all players are endowed with a probability distribution over the state of the game, see (Cesa-Bianchi and Lugosi, 2006).

Several no-regret learning algorithms exist, and both the algorithms and their performance crucially depend on the feedback information available at every round of the game.

<sup>\*</sup> This research was gratefully funded by the European Union ERC Starting Grant CONENE, and the Swiss National Science Foundation grant SNSF 200021\_172781.

<sup>★★</sup>These authors contributed equally to this work.

In the *full-information* setting, where the player observes its utility under every action (Freund and Schapire, 1997), such algorithms attain an optimal  $\mathcal{O}(\sqrt{T \log K})$  regret, with  $T$  being the horizon length, and  $K$  being the number of actions. When only *bandit* feedback is available, that is, the player only observes its utility for the action picked (Auer et al., 2002), the optimal regret is  $\mathcal{O}(\sqrt{KT \log K})$  with a significantly worse dependence on the number of actions  $K$ . However, the information structure found in auctions lies in between these two models; on the one hand, full-information feedback is unrealistic since it requires perfect knowledge of opponents' bids and market constraints, and on the other hand, bandit feedback is too restrictive since partial observations of utilities under different actions are often available owing to the auction rules and the additional information released. For instance, in auctions of electrical power marginal prices are announced after every round. Motivated by these particularities, our goal is to extend existing bandit algorithms to account for such additional information.

Our contributions are as follows. We consider no-regret learning in a general class of repeated auctions. Employing standard results from the multi-armed bandit literature, we propose an algorithm that participants could use to update their bidding strategies based on the observed auction data. Specifically, our approach exploits the information available to the bidders when their bids are not accepted by the auctioneer. We show that the proposed algorithm enjoys a no-regret guarantee that can strictly improve upon bandit algorithms. We demonstrate this fact on a market with a simple supply-demand balance constraint and a marginal pricing payment mechanism.

Then, we consider a more general market setting for which the algorithm is not directly applicable since it requires parameters that are unknown in this setting, that is, revelation probabilities for additional information. We propose a heuristic method for this case and demonstrate it outperforms existing bandit algorithms in our experiments.

Let us contrast our work with the existing works on learning algorithms with partial information. A partial monitoring framework was introduced in (Piccolboni and Schindelhauer, 2001) and was extensively studied (e.g., in Cesa-Bianchi and Lugosi (2006, §6.4)) for online learning with feedback matrices. However, this setting is restricted to learning problems where the utilities are chosen from a finite set, e.g.,  $\{0, 1\}$ . More recently, Mannor and Shamir (2011) introduced a novel framework in which the learning agent is equipped with a sequence of feedback graphs over the action space encoding which additional information is revealed by every action. In case the feedback graph is known before picking the action, Lykouris et al. (2017) obtained small loss regret guarantees that do not depend on the number of actions. In case the feedback graph is revealed only after picking the action, the works of (Alon et al., 2015, 2017) derived regret bounds as a function of the feedback graphs' independence numbers (or sizes of their maximum acyclic subgraphs). In contrast to these works, in our auction framework such graphs may not be known to the bidders at any point during the game. For instance, in the general market problem considered in Section 3.3 revelation probabilities for the additional information originate from a hidden feedback graph, and we have to develop meaningful heuristics to compute these parameters. Feng et al. (2018) study ad auctions where the bidders observe their value for a click only when their ads get clicked. In contrast, in our framework valuations/true costs are known to the bidders, and the allocations are deterministic. Finally, partial information in the context of multi-agent learning have been explored in (Sessa et al., 2019), assuming players can observe their opponents' actions, in addition to the standard bandit feedback. Differently, in this work we do not assume observation of opponents' bids.

Section 2 lays out the preliminaries for the framework and existing algorithms. Our algorithm is proposed in Section 3, followed by two applications. Section 4 presents the case studies based on electricity markets.

## 2. PRELIMINARIES

### 2.1 Repeated Auction Framework

We consider a general reverse (or procurement) auction problem. The set of participants consists of the bidders  $\ell \in \mathcal{N} = \{1, \dots, |\mathcal{N}|\}$ . Let there be  $q \in \mathbb{N}$  types of goods/supplies. For instance, in an electricity market, these types could refer to active and reactive power injections at different nodes, and also control reserves. Goods of the same type from different bidders are interchangeable to the auctioneer. Each bidder has a private true cost function  $c_\ell : \mathbb{X}_\ell \rightarrow \mathbb{R}_+$ ,  $\mathbb{X}_\ell \subseteq \mathbb{R}_+^q$ . We further assume that  $0 \in \mathbb{X}_\ell$  and  $c_\ell(0) = 0$ . This holds for many electricity markets that do not allow shot-down costs. Each bidder  $\ell$  has a finite strategy set  $\mathcal{K}_\ell = \{1, \dots, |\mathcal{K}_\ell|\}$  that consists of its true cost function and bid functions of the form  $b_\ell^k : \mathbb{X}_\ell^k \rightarrow \mathbb{R}_+$ ,

$\mathbb{X}_\ell^k \subseteq \mathbb{R}_+^q$ , where  $k \in \mathcal{K}_\ell$ ,  $0 \in \mathbb{X}_\ell^k \subseteq \mathbb{R}_+^q$ , and  $b_\ell^k(0) = 0$ . Without loss of generality,  $1 \in \mathcal{K}_\ell$  is the true cost function.

Let  $T \in \mathbb{N}$  be the horizon length. We assume that for all rounds the true costs and the strategy set  $\mathcal{K}_\ell$  remain unchanged. Let  $k_\ell(t) \in \mathcal{K}_\ell$  denote the strategy of bidder  $\ell$  at time  $t \leq T$ . Given the strategy profile  $\mathcal{B}(t) = \{b_\ell^{k_\ell(t)}\}_{\ell \in \mathcal{N}}$ , a *mechanism* defines an *allocation rule*  $x_\ell^*(\mathcal{B}(t)) \in \mathbb{X}_\ell^{k_\ell(t)}$ , and a *payment rule*  $p_\ell(\mathcal{B}(t)) \in \mathbb{R}$  for each bidder  $\ell$ . In many auctions, the allocation rule is determined by an optimization problem of the form,

$$J(\mathcal{B}(t)) = \min_{x \in \mathbb{X}(t)} \sum_{\ell \in \mathcal{N}} b_\ell^{k_\ell(t)}(x_\ell) \quad \text{s.t.} \quad \sum_{\ell \in \mathcal{N}} x_\ell \in \mathbb{S}, \quad (1)$$

where  $\mathbb{X}(t) = \prod_{\ell \in \mathcal{N}} \mathbb{X}_\ell^{k_\ell(t)} \subset \mathbb{R}_+^{q|\mathcal{N}|}$ , and the set  $\mathbb{S} \subset \mathbb{R}_+^q$  corresponds to the market constraints. In an electricity market, these constraints may correspond to network balance constraints found in optimal power flow problems (Wu et al., 1996), or probabilistic security requirements found in control reserves markets (Abbaspourtorbati and Zima, 2015). Assume that the market constraints remain unchanged throughout the horizon.

Let the optimal solution of (1) be denoted by  $x^*(\mathcal{B}(t))$ . We assume that in case of multiple optima there is a tie-breaking rule. The utility of bidder  $\ell$  is linear in the payment received;  $u_\ell(t) = u_\ell(\mathcal{B}(t)) = p_\ell(\mathcal{B}(t)) - c_\ell(x_\ell^*(\mathcal{B}(t)))$ . A bidder whose bid is not accepted,  $x_\ell^*(\mathcal{B}(t)) = 0$ , is not paid,  $u_\ell(t) = 0$ , and is referred to as *loser*.

Since the bidders strategize to receive larger payments, there has been many proposals for different payment rules, for instance, locational marginal pricing (Schweppe et al., 1988), Vickrey-Clarke-Groves (VCG) (Ausubel and Milgrom, 2006), and core-selecting mechanisms (Day and Milgrom, 2008). Even though these payment rules are well-discussed in terms of how the bidders should pick their strategies in a Nash equilibrium, in reality bidders are profit maximizing entities that compete under privacy considerations and limited information. Hence, it may not be realistic to assume that they can compute or they would be willing to pick their Nash equilibrium strategy.<sup>1</sup> Instead, a more practical assumption is that the bidders choose their strategies following simple adaptive and sequential algorithms based on observed auction data.

### 2.2 No-Regret Learning and Correlated Equilibrium

The performance of a bidder  $\ell$  in a repeated auction can be measured in terms of *regret*, which is a standard notion used in the online learning to measure the performance of a sequential decision making algorithm (see, e.g., (Shalev-Shwartz, 2012)). Define  $\mathcal{B}_{-\ell}(t) = \{b_j^{k_j(t)}\}_{j \in \mathcal{N} \setminus \{\ell\}}$ .

*Definition 1.* (Regret). The *regret* of bidder  $\ell$  at time  $T$  is  $R_\ell(T) = \max_{k \in \mathcal{K}_\ell} \sum_{t=1}^T u_\ell(\{\mathcal{B}_{-\ell}(t), b_\ell^k\}) - \sum_{t=1}^T u_\ell(t)$ .

After  $T$  rounds, hence,  $R_\ell(T)$  quantifies the maximum profit bidder  $\ell$  could have made had she known the sequence of opponents' bids ahead of time, and had she

<sup>1</sup> Under the VCG mechanism, reporting true cost is the dominant strategy of every bidder. However, bidders can still turn to group deviations to maximize their profits (Ausubel and Milgrom, 2006).

---

**Algorithm 1** MWU algorithm for bidder  $l$

---

**Input:** Strategy set  $\mathcal{K}_\ell$  with  $|\mathcal{K}_\ell| = K$ , parameter  $\eta$   
 Initialize mixed strategy  $\mathbf{w}_1 = [1/K, \dots, 1/K]$   
**for**  $t = 1, 2, \dots, T$  **do**  
     Compute estimate  $\tilde{\mathbf{I}}_t$  of the loss vector  $\mathbf{I}_t \in [0, 1]^K$   
     Update mixed strategy:  
          $\mathbf{w}_{t+1}[i] \propto \mathbf{w}_t[i] \cdot \exp(-\eta \tilde{\mathbf{I}}_t[i]), \quad i = 1, \dots, K$   
**end for**

---

chosen the best fixed strategy in  $\mathcal{K}_\ell$ . An algorithm for bidder  $l$  is *no-regret* if  $R_\ell(T)/T \rightarrow 0$  as  $T \rightarrow \infty$ .

In a multi-agent setting like the aforementioned auction problems, the notion of regret has a close connection to equilibria. In fact, it can be shown that if every participant bids according to a no-regret algorithm, the empirical distribution of bids converge to a *coarse-correlated equilibrium* (CCE) of the one-shot game.

*Definition 2.* (CCE). A *coarse-correlated equilibrium* is a distribution  $\sigma$  over  $\prod_{\ell \in \mathcal{N}} \mathcal{K}_\ell$  such that, for each  $l \in \mathcal{N}$ ,  $\mathbb{E}_{\mathcal{B} \sim \sigma} [u_\ell(\mathcal{B})] \geq \mathbb{E}_{\mathcal{B} \sim \sigma} [u_\ell(\{\mathcal{B}_{-\ell}, b_\ell^k\})]$ , for all  $k \in \mathcal{K}_\ell$ .

As a remark, CCEs are the largest (and the weakest) class of equilibria, and include pure and mixed Nash Equilibria (NE). However, computing any NE is PPA-complete (a weaker version of NP-completeness, see (Daskalakis et al., 2009)) and requires full knowledge of the game.

No-regret algorithms for a generic bidder  $l$  can be derived by mapping the repeated auction to an online learning problem faced by bidder  $l$ : At every round  $t$  bidder  $l$  picks an action  $k(t) \in \mathcal{K}_\ell$ , an adversarial environment selects a *loss vector*  $\mathbf{I}_t \in [0, 1]^{|\mathcal{K}_\ell|}$ , and bidder  $l$  incurs loss  $\mathbf{I}_t[k(t)]$ . Since the bidders are in general aware of the range of their utilities, we assume they can map their utilities to losses in the  $[0, 1]$  interval. Thus, we let utility  $u_\ell(t) = u_\ell(\mathcal{B}(t)) \in \mathbb{R}$  correspond to loss  $l_\ell(t) = l_\ell(\mathcal{B}(t)) = 1 - s_\ell(u_\ell(\mathcal{B}(t))) \in [0, 1]$ , where  $s_\ell : \mathbb{R} \rightarrow [0, 1]$  is a suitable monotone map. At every round  $t$ , hence, the corresponding loss vector is:

$$\mathbf{I}_t = [l_\ell(\{\mathcal{B}_{-\ell}(t), b_\ell^1\}), \dots, l_\ell(\{\mathcal{B}_{-\ell}(t), b_\ell^{|\mathcal{K}_\ell|}\})]. \quad (2)$$

Note that the monotonicity of  $s_l$  implies that our regret definition maps to a regret based on such loss formulation.

In order to attain no-regret, bidder  $l$  must randomize her actions and bid according to *mixed* strategies, that is, probability distributions over  $\mathcal{K}_\ell$  (Cesa-Bianchi and Lugosi, 2006). Thus, one often reasons about regrets in expectation. Let  $K = |\mathcal{K}_\ell|$  be the number of actions in the strategy set of bidder  $l$  (we drop the dependence on  $l$  for ease of notation) and  $\mathbf{w}_t \in [0, 1]^K$  be the mixed strategy of bidder  $l$  at round  $t$ . A large family of no-regret algorithms are based on a simple, yet effective, Multiplicative Weight Update (MWU) rule to update bidder  $l$ 's mixed strategy (Freund and Schapire, 1997). Such update rule is summarized in Algorithm 1. The performance of these algorithms depend on the chosen estimate  $\tilde{\mathbf{I}}_t$  of the loss vector  $\mathbf{I}_t$ .

Full-information feedback algorithms such as HEDGE (Freund and Schapire, 1997) use the true loss vector  $\mathbf{I}_t$  (of (2)) as the estimate  $\tilde{\mathbf{I}}_t$ . Such algorithms attain an optimal expected regret of  $\mathcal{O}(\sqrt{T \log K})$ . However, full-information

feedback is unrealistic in repeated auctions, since computing  $l_\ell(\{\mathcal{B}_{-\ell}(t), b_\ell^k\})$  for  $k \neq k_\ell(t)$  requires the full knowledge of the bids  $\mathcal{B}_{-\ell}(t)$  and the market constraints to generate solutions of the optimization problem (1).

Bandit algorithms such as EXP3 (Auer et al., 2002) use only the obtained loss  $l_\ell(t)$  to build an estimate as:

$$\tilde{\mathbf{I}}_t = \left[ 0, \dots, 0, \frac{l_\ell(t)}{\mathbf{w}_t[k_\ell(t)]}, 0, \dots, 0 \right]. \quad (3)$$

Although one can show that  $\tilde{\mathbf{I}}_t$  is an *unbiased* estimate of the true loss vector  $\mathbf{I}_t$ , that is,  $\mathbb{E}_{k \sim \mathbf{w}_t} [\tilde{\mathbf{I}}_t[i]] = \mathbf{I}_t[i]$  for all  $i \in \mathcal{K}_\ell$ , the variance of  $\tilde{\mathbf{I}}_t[i]$  grows with the inverse of the squared probability  $\mathbf{w}_t[i]^2$ . This leads to an expected regret of  $\mathcal{O}(\sqrt{KT \log K})$  (matching an algorithm-independent lower bound (Auer et al., 2002) up to log factors), which scales significantly worse with the number of actions  $K$  and often leads to poor performance.

### 3. LEARNING FROM PARTIALLY OBSERVED DATA

While the full-information feedback is unrealistic, the feedback information available in repeated auctions is typically richer than the bandit feedback. In fact, when bidder  $l$  is a *loser* of the auction, not only she observes the loss for the chosen bid, which corresponds to  $l_\ell(0)$ , but often can also infer about other actions that would have led to such losing outcome. This is the key idea that we exploit to improve upon the regret bound of the existing bandit EXP3 algorithm. We will then make it more concrete in Sections 3.2 and 3.3 by considering two specific classes of auctions. As a remark, when bidder  $l$  is a winner of the auction she can often infer about other winning actions; however, she cannot infer the exact utilities. We plan on addressing this fact in our future work.

#### 3.1 Extended EXP3 Algorithm

Consider time  $t$ , and assume bidder  $l$  loses the auction, that is,  $x^*(\mathcal{B}(t)) = 0$ . In this case, we assume bidder  $l$  gets to know of a subset  $\mathcal{L}_t \subseteq \mathcal{K}_\ell$  of losing actions (including  $k_\ell(t)$ ), i.e., actions which would have led to the same losing outcome. In the class of auctions considered in Sections 3.2, for instance, bidder  $l$  finds out about  $\mathcal{L}_t$  from the marginal price set by the auctioneer.

For any action  $k \in \mathcal{K}_\ell$ , we let  $\mathbf{r}_t[k]$  denote the probability that bidder  $l$  discovers the utility of action  $k$ . We refer to  $\mathbf{r}_t[k]$ 's as the *revelation probabilities*. Note that  $\mathbf{r}_t[k]$  is always greater than the probability  $\mathbf{w}_t[k]$  of playing action  $k$ , since when action  $k$  is played, bidder  $l$  directly discovers its utility. However, as explained before, bidder  $l$  can find out about losing actions even when playing actions  $i \neq k$ . We will provide an explicit expression of such probabilities for the auctions considered in Sections 3.2 and 3.3. For now, suppose such revelation probabilities are known to bidder  $l$ . Using this information, bidder  $l$  can construct the following estimator, for any  $k \in \mathcal{K}_\ell$ :

$$\tilde{\mathbf{I}}_t[k] = \begin{cases} l_\ell(t), & \text{if } x_\ell^*(\mathcal{B}(t)) > 0 \text{ and } k = k_\ell(t), \\ \mathbf{r}_t[k], & \text{if } x_\ell^*(\mathcal{B}(t)) > 0 \text{ and } k \neq k_\ell(t), \\ 0, & \text{if } x_\ell^*(\mathcal{B}(t)) = 0 \text{ and } k \in \mathcal{L}_t, \\ \mathbf{r}_t[k], & \text{if } x_\ell^*(\mathcal{B}(t)) = 0 \text{ and } k \notin \mathcal{L}_t, \\ 0, & \end{cases} \quad (4)$$

where  $l_\ell(0) = 1 - s_\ell(0)$  for the sake of brevity, and  $\mathbf{r}_t[k] = \mathbf{w}_t[k]$  for the first case in which the bidder is winning. In words, if action  $k_\ell(t)$  is a winning action, the estimator (4) coincides with the bandit feedback estimator (3). On the other hand, if  $k_\ell(t)$  is a losing action,  $\tilde{\mathbf{I}}_t$  computed above has a larger number of non-zero entries compared to (3).

It is not hard to prove the following fact (see (Karaca et al., 2019) for the proof), which is the key to obtaining a sublinear regret bound.

*Fact 1.* The estimator  $\tilde{\mathbf{I}}_t$  computed as in (4) is an unbiased estimate of the true loss vector  $\mathbf{I}_t$  in (2).

Compared to the bandit estimator in (3), the estimator considered above has a strictly smaller variance (since  $\mathbf{r}_t[k] \geq \mathbf{w}_t[k]$  for all  $k$ ) by virtue of the additional information used. To quantify such improvement, we define the quantity  $\alpha_t^k := \frac{\mathbf{r}_t[k]}{\mathbf{w}_t[k]}$ . It measures the additional knowledge available to bidder  $\ell$  about the loss  $\mathbf{I}_t[k]$ , when compared to the standard bandit feedback. Note that  $\alpha_t^k \geq 1$ , and it increases with the information available to bidder. To make our results explicit, we define the following aggregate quantity, called the *average feedback information*.

*Definition 3.* (Average feedback information). The average feedback information available to bidder  $\ell$  over  $T$  auction rounds is  $\alpha_{\text{avg}} := \left( \frac{1}{TK} \sum_{t=1}^T \sum_{k=1}^K \frac{1}{\alpha_t^k} \right)^{-1} \in [1, K]$ .

Bandit feedback corresponds to  $\alpha_{\text{avg}} = 1$ , since  $\mathbf{r}_t[k] = \mathbf{w}_t[k]$  for all  $k$ , while  $\alpha_{\text{avg}} = K$  in case of full-information feedback, since  $\mathbf{r}_t[k] = 1$  for all  $k$ . In repeated auctions, (based on the previous discussions) the average feedback information  $\alpha_{\text{avg}}$  depends on the auction type and also the information available to the participants, and can be computed only after all the auction rounds are completed.

The following theorem bounds the performance of the proposed algorithm, as a function of  $\alpha_{\text{avg}}$ .

*Theorem 1.* Assume bidder  $\ell$  bids according to Algorithm 1 with loss vector estimate  $\tilde{\mathbf{I}}_t$  computed as in (4) and  $\eta = \sqrt{2\alpha_{\text{avg}} \log(K)/(KT)}$ , then  $\mathbb{E}[R_\ell(T)] \leq \sqrt{2(K/\alpha_{\text{avg}})T \log K}$ .

Theorem 1 generalizes the regret bounds of full-information and bandit algorithms, obtained by setting  $\alpha_{\text{avg}} = K$  or  $\alpha_{\text{avg}} = 1$  respectively, to the case where intermediate information is available to bidder  $\ell$ . In a repeated auction, where  $\alpha_{\text{avg}}$  is typically greater than 1, the obtained regret bound strictly improves upon standard bandit guarantees. Our proof is relegated to (Karaca et al., 2019) and it follows from standard online learning arguments by making use of Fact 1 and Definition 3. To optimally select the learning rate  $\eta$ , the value  $\alpha_{\text{avg}}$  needs to be known ahead of time. However, we expect that the bidders could estimate and even update this value using the information from the previous auction rounds. In practice, in our experiments we show that a wide-range of  $\eta$  leads to desirable performance.

In comparison with the works considering feedback graphs (Alon et al., 2015, 2017; Lykouris et al., 2017), Theorem 1 does not depend on any graph-theoretic quantity; instead, it depends on the notion of average feedback information defined above. This makes it applicable to our more

general auction framework in that the way the additional information originates can be unknown to the bidders (see Section 3.3). Comparing with their algorithms and regret bounds, Alon et al. (2017) use an estimator that coincides with (4) where the revelation probabilities are computed from the feedback graph, and obtain the regret

bound  $\sqrt{2 \sum_{t=1}^T m_t \log K}$ , where  $m_t$  is a precomputed upper bound on the maximum size of acyclic subgraphs of the feedback graph at round  $t$ . On the other hand, (Lykouris et al., 2017) requires knowing the feedback graphs ahead of time and relies on a different algorithm which is a freezing modification to HEDGE. Their regret bound holds with high probability and it is of the form  $o(\beta L^*)$ , where  $L^*$  is the loss of the best fixed strategy, and  $\beta$  is an upper bound on the independence number of the feedback graphs at all rounds. Finally, note that these works replace the  $\sqrt{K}$  term of the regret bounds of bandit algorithms with a graph-theoretic quantity, which might still be equal to  $\sqrt{K}$  in the worst-case.

### 3.2 Auctions with Simple Constraints and Convex Bids

We consider a simpler auction problem where the auctioneer has to procure a fixed amount  $Q \in \mathbb{R}_+$  of a single type of good. In addition to the assumptions in Section 2, each bidder is now equipped with a finite strategy set consisting of strongly convex and increasing cost/bid functions over compact intervals. The allocation rule is given by,

$$x^*(\mathcal{B}(t)) = \arg \min_{x \in \mathbb{X}(t)} \sum_{\ell \in \mathcal{N}} b_\ell^{k_\ell(t)}(x_\ell) \text{ s.t. } \sum_{\ell \in \mathcal{N}} x_\ell \geq Q. \quad (5)$$

As a remark, some of the European reserve and day-ahead markets belong to this class since they ignore the network constraints. Let  $\lambda^*(\mathcal{B}(t)) \in \mathbb{R}_+$  denote the Lagrange multiplier associated with the constraint in (5), called the marginal price, and it is announced after each round  $t$  to all bidders. The payment rule is then given by  $p_\ell(\mathcal{B}(t)) = \lambda^*(\mathcal{B}(t)) x_\ell^*(\mathcal{B}(t))$ .

Assume in round  $t$  bidder  $\ell$  is a loser with  $x_\ell^*(\mathcal{B}(t)) = 0$ . From the Karush-Kuhn-Tucker optimality conditions and strong duality of (5), for any  $k \in \mathcal{K}_\ell$ , we have  $\frac{d}{dx} b_\ell^k(x) \Big|_{x=0} \geq \lambda^*(\mathcal{B}(t)) \iff x_\ell^*(\{\mathcal{B}_{-l}(t), b_\ell^k\}) = 0$ . Equivalently, the marginal price information enables a losing bidder to be aware of all the other losing actions in its strategy set (similar arguments can be made for any Lagrange multiplier-based payments under strong duality and strong convexity). Hence, bidder  $\ell$  gets to know the set of all losing actions:  $\mathcal{L}_t = \{k \in \mathcal{K}_\ell \mid \frac{d}{dx} b_\ell^k(x) \Big|_{x=0} \geq \lambda^*(\mathcal{B}(t))\}$ . Notice that the set above is meaningful only when bidder  $\ell$  is a loser, since otherwise marginal price can change whenever bidder  $\ell$  picks another action.

In the partial information setting described above, a losing action  $k$  is revealed when any other losing action is played, since each losing action reveals the full set of losing actions  $\mathcal{L}_t$ . Hence, the revelation probabilities are given by:

$$\mathbf{r}_t[k] = \begin{cases} \sum_{j \in \mathcal{L}_t} \mathbf{w}_t[j], & \text{if } x_\ell^*(\mathcal{B}(t)) = 0 \text{ and } k \in \mathcal{L}_t, \\ \mathbf{w}_t[k], & \text{otherwise.} \end{cases}$$

In the first case, the probabilities are calculated by a sum, since choosing different losing actions constitute independent events. Then, the estimator  $\tilde{\mathbf{I}}_t$  in (4) can be obtained by plugging the probabilities defined above.

### 3.3 General Auctions

We now consider the general auction problem introduced in (1). Assume in round  $t$  bidder  $\ell$  is a loser with  $x_\ell^*(\mathcal{B}(t)) = 0$ . Since the allocation rule is given by an optimization problem, it can be verified that  $x_\ell^*(\mathcal{B}(t)) = 0 \implies x_\ell^*(\{\mathcal{B}_{-\ell}(t), b_\ell^k\}) = 0, \forall k \in \mathcal{L}_t$ , where

$$\mathcal{L}_t = \{k \in \mathcal{K}_\ell \mid \mathbb{X}_\ell^k \subseteq \mathbb{X}_\ell^{k_\ell(t)}, \text{ and } \forall x \in \mathbb{X}_\ell^k, b_\ell^k(x) \geq b_\ell^{k_\ell(t)}(x)\}. \quad (6)$$

If  $b_\ell^{k_\ell(t)} : \mathbb{X}_\ell^{k_\ell(t)} \rightarrow \mathbb{R}_+$  is a losing bid function, also the bid functions that lie in the epigraph of  $b_\ell^{k_\ell(t)}$  are losing. Note that  $\mathbb{X}_\ell^k \subseteq \mathbb{X}_\ell^{k_\ell(t)}$  is required since the function can be assumed to take infinitely large values outside of its domain. The statement above holds regardless of the payment rule, and thus it is applicable to a wide-range of auctions.

In such a framework, a generic action  $i$  is revealed either when action  $i$  is picked, or when a losing action  $j$  is picked such that action  $i$  belongs the epigraph of action  $j$ . Hence, to define the revelation probabilities, we define the set of all losing actions by  $\bar{\mathcal{L}}_t = \{k \in \mathcal{K}_\ell \mid x_\ell^*(\{\mathcal{B}_{-\ell}(t), b_\ell^k\}) = 0\}$ , and we let  $\mathcal{R}_k = \{j \in \mathcal{K}_\ell \mid \mathbb{X}_\ell^j \supseteq \mathbb{X}_\ell^k, \text{ and } \forall x \in \mathbb{X}_\ell^k, b_\ell^j(x) \leq b_\ell^k(x)\}$ , be the set of all actions for which the action  $k \in \mathcal{K}_\ell$  belongs to the epigraph of the corresponding bid function. Then, according to our previous arguments, the revelation probabilities correspond to

$$\mathbf{r}_t[k] = \begin{cases} \sum_{j \in \bar{\mathcal{L}}_t \cap \mathcal{R}_k} \mathbf{w}_t[j], & \text{if } x_\ell^*(\mathcal{B}(t))=0 \text{ and } k \in \mathcal{L}_t, \\ \mathbf{w}_t[k], & \text{otherwise.} \end{cases} \quad (7)$$

Observe that, in the first case, the sum involves all losing actions that can reveal information about action  $k$ .

Computing these probabilities requires knowing the set of all losing actions  $\bar{\mathcal{L}}_t$ , which is in general unknown to the bidders. To apply our proposed algorithm in such a general framework, we propose the following heuristic choice:

$$\hat{\mathbf{r}}_t[k] = \begin{cases} \mathbf{w}_t[k] + \sum_{j \in \mathcal{R}_k \setminus \{\mathcal{W}_t \cup k\}} \mathbf{w}_t[j], & \text{if } x_\ell^*(\mathcal{B}(t))=0, \\ \mathbf{w}_t[k], & \text{otherwise,} \end{cases}$$

where  $\mathcal{W}_t \subset \mathcal{K}_\ell$  is the set of actions that lead to a non-zero allocation at every instance up to time  $t$ . Such set can be computed by bidder  $\ell$  based on the previous auction rounds. Note that the above heuristic approximates the true  $\mathbf{r}_t[k]$  by assuming that the actions that always lead to non-zero allocations ( $\mathcal{W}_t$ ) are the ones that would lead to non-zero allocations in the current auction round. If  $\mathcal{W}_t$  is equivalent to the latter set, the heuristic coincides with (7). When such revelation probabilities are plugged in (4), the resulting estimator  $\tilde{\mathbf{I}}_t$  may not be unbiased and hence the result of Theorem 1 is not directly applicable. Nevertheless, in our numerical case studies we will showcase the performance of the proposed heuristic in realistic electricity market models.




Fig. 1. Average regret (in \$ and averaged over 50 runs and 3 bidders) of different algorithms: Shaded areas represent  $\pm$  one standard deviation.

## 4. NUMERICAL STUDIES

First, we consider an instance of Section 3.2. Then, we illustrate our approach for the general setting on the Swiss reserve auction. (Karaca et al., 2019) provides additional experiments on an IEEE 14-bus test system.

### 4.1 Auction with Simple Constraints

Consider the set-up of Section 3.2, with  $Q = 15$  MW and three bidders with quadratic polynomial cost functions  $c_\ell(x_\ell) = a_\ell x_\ell^2 + d_\ell x_\ell$ , with  $a_1 = 0.1, a_2 = 0.095, a_3 = 0.105, d_1 = 8, d_2 = 9, d_3 = 10$ , and production limits of  $\bar{X}_\ell = 10$  MW for  $l = 1, 2, 3$ . Each bidder chooses among 15 actions, obtained by perturbing the linear terms of their cost functions as  $d_\ell + \delta$ , where  $\delta$  is chosen from a uniform distribution in the interval  $[-6, 30]$ . Let  $T = 600$  be the horizon. For EXP3 and HEDGE algorithms, we picked the optimal learning rates  $\eta = \sqrt{2 \log(K)/(KT)}$  and  $\eta = \sqrt{8 \log(K)/T}$ , respectively, see (Bubeck and Cesa-Bianchi, 2012). For our algorithm, we picked  $\eta = \sqrt{2\hat{\alpha} \log(K)/(KT)}$ , by picking  $\hat{\alpha}$  from  $\{8.5, 11, 13.5\}$ , since the true value of  $\alpha_{\text{avg}}$  is not known a-priori and it is required to lie in the interval  $[1, K] = [1, 15]$ . These values are chosen to be the same for all bidders. Recall that the full-information HEDGE algorithm is unrealistic and hence we run it to upper bound the achievable performance.

Figure 1 illustrates the average regret as in Definition 1 over 50 runs and 3 bidders for the cases under which all bidders implement HEDGE, EXP3, and the extended EXP3 algorithms. Observe that the extended EXP3 achieves a performance significantly better than EXP3, and also close to HEDGE for a wide-range of learning rates  $\eta$ . Figure 2 shows the average feedback information  $\alpha_{\text{avg}}$  in simulations based on different number of actions, again averaged over 50 runs. We see that  $\alpha_{\text{avg}}$  is close to  $K$  in all cases, which explains the improvement in the bound obtained in Theorem 1. Additional information helps Bidder 2, who has the lowest cost, to discard losing actions quickly, as shown in (Karaca et al., 2019, Table 1).

### 4.2 Swiss Reserve Market

The following simulations are based on the bids placed in the 46<sup>th</sup> weekly Swiss reserve auction of 2014 (Abaspourtorbati and Zima, 2015). This auction involves 21 plants bidding for secondary reserves, 25 for positive tertiary reserves and 21 for negative tertiary reserves. Since the problem is big, we picked six of the largest




Fig. 2.  $\alpha_{\text{avg}}$  as a function of the number of actions  $K$ .




Fig. 3. Average regret (in CHF) for Swiss market.

Table 1. Average social cost (in CHF).

TRUTHFUL BIDDING	EXP3	Extended EXP3	HEDGE
2,615,800	2,652,300	2,637,400	2,635,500

secondary reserve providers as the learning agents. The bids are discrete, that is, they are given by sets of reserve size and price pairs. Strategy sets are created by inflating the bid prices with multipliers from the set  $\{1, 1.05, 1.15, 1.25, 1.35, 1.45\}$ . The market involves complex constraints arising from nonlinear cumulative distribution functions. These constraints imply that the deficit of reserves cannot occur with a probability higher than 0.2%. The payment rule is  $p_\ell(\mathcal{B}(t)) = b_\ell^{k_\ell(t)}(x_\ell^*(\mathcal{B}(t)))$ .

Let  $T = 600$  be the horizon. For EXP3 and HEDGE algorithms, we again picked the optimal learning rates  $\eta$ . For our algorithm, we picked  $\eta = \sqrt{2\hat{\alpha} \log(K)/(KT)}$  with  $\hat{\alpha} = 6$ . Figure 3 shows the average regrets over 50 runs and over 6 bidders for the cases under which all bidders implement HEDGE, EXP3, and the extended EXP3 algorithm. Initially, the regrets are large considering that the average true cost of a bidder is 39,429 CHF. The extended EXP3 algorithm attains a better performance than the EXP3 with both the true revelation probabilities (computed by  $\mathcal{R}_k$ ) and our heuristic choice. In Table 1 we show the resulting average social cost: the participants' total production cost. We will explore mechanism design with the goal of achieving lower social cost as a future work.

## 5. CONCLUSION

In this paper, we have considered online learning in a general class of repeated auctions. For such problems, we have shown that the information structure lies in between the well-studied full-information and bandit settings. Exploiting the additional information acquired by the bidders whose bids are not accepted, we have formulated an extension to the standard bandit algorithms involving a more accurate utility estimation. We have shown that the regret guarantee of this algorithm improves upon the bandit algorithms. Our results were verified in several case studies based on realistic electricity market models.

## REFERENCES

- Abbaspourtorbati, F. and Zima, M. (2015). The Swiss reserve market: Stochastic programming in practice. *IEEE Tran. on Power Syst.*, 31(2), 1188–1194.
- Alon, N., Cesa-Bianchi, N., Dekel, O., and Koren, T. (2015). Online learning with feedback graphs: Beyond bandits. In *Annual Conf. on Learn. Theory*, volume 40.
- Alon, N., Cesa-Bianchi, N., Gentile, C., Mannor, S., Mansour, Y., and Shamir, O. (2017). Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM J. on Computing*, 46(6), 1785–1826.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R.E. (2002). The nonstochastic multiarmed bandit problem. *SIAM J. on Computing*, 32(1), 48–77.
- Ausubel, L.M. and Milgrom, P. (2006). The lovely but lonely Vickrey auction. In *Combinatorial Auctions*, chapter 1. MIT press.
- Bichler, M. and Goeree, J.K. (2017). *Handbook of spectrum auction design*. Cambridge University Press.
- Bubeck, S. and Cesa-Bianchi, N. (2012). Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Found. and Tr. in Mach. Learn.*, 5(1), 1–122.
- Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, learning, and games*. Cambridge university press.
- Cramton, P. (2017). Electricity market design. *Oxford Review of Econ. Policy*, 33(4), 589–612.
- Daskalakis, C., Goldberg, P.W., and Papadimitriou, C.H. (2009). The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2), 89–97.
- Day, R. and Milgrom, P. (2008). Core-selecting package auctions. *Int. J. of Game Theory*, 36(3-4), 393–407.
- Feng, Z., Podimata, C., and Syrgkanis, V. (2018). Learning to bid without knowing your value. In *ACM Conf. on Economics and Comp.*, 505–522.
- Freund, Y. and Schapire, R.E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. of Computer and Syst. Sci.*, 55(1), 119–139.
- Karaca, O., Sessa, P.G., Leidi, A., and Kamgarpour, M. (2019). No-regret learning from partially observed data in repeated auctions. *arXiv preprint arXiv:1912.09905*.
- Lykouris, T., Sridharan, K., and Tardos, É. (2017). Small-loss bounds for online learning with partial information. *arXiv preprint arXiv:1711.03639*.
- Mannor, S. and Shamir, O. (2011). From bandits to experts: On the value of side-observations. In *Advances in Neur. Inf. Proc. Syst.*, 684–692.
- Piccolboni, A. and Schindelhauer, C. (2001). Discrete prediction games with arbitrary feedback and loss. In D. Helmbold and B. Williamson (eds.), *Computational Learn. Theory*, 208–223. Springer.
- Schweppe, F.C., Caramanis, M.C., Tabors, R.D., and Bohn, R.E. (1988). *Spot pricing of electricity*. Springer Science & Business Media.
- Sessa, P.G., Bogunovic, I., Kamgarpour, M., and Krause, A. (2019). No-regret learning in unknown games with correlated payoffs. *Neur. Inf. Proc. Syst.*
- Shalev-Shwartz, S. (2012). Online learning and online convex optimization. *Found. and Tr. in Mach. Learn.*, 4(2), 107–194.
- Teytelboym, A. (2019). Natural capital market design. *Oxford Review of Economic Policy*, 35(1), 138–161.
- Wu, F., Varaiya, P., Spiller, P., and Oren, S. (1996). Folk theorems on transmission access: Proofs and counterexamples. *J. of Regulatory Economics*, 10(1), 5–23.