

MIT Open Access Articles

Learning in Auctions: Regret is Hard, Envy is Easy

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Daskalakis, Constantinos, and Vasilis Syrgkanis. "Learning in Auctions: Regret Is Hard, Envy Is Easy." 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS) (October 2016).

As Published: <http://dx.doi.org/10.1109/FOCS.2016.31>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/110834>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Learning in Auctions: Regret is Hard, Envy is Easy

Constantinos Daskalakis*
EECS, MIT
costis@csail.mit.edu

Vasilis Syrgkanis†
Microsoft Research, NYC
vasy@microsoft.com

Abstract

A large line of recent work studies the welfare guarantees of simple and prevalent combinatorial auction formats, such as selling m items via simultaneous second price auctions (SiSPAs) [CKS08, BR11, FFGL13]. These guarantees hold even when the auctions are repeatedly executed and the players use no-regret learning algorithms to choose their actions. Unfortunately, off-the-shelf no-regret learning algorithms for these auctions are computationally inefficient as the number of actions available to each player is exponential. We show that this obstacle is insurmountable: there are no polynomial-time no-regret learning algorithms for SiSPAs, unless $\text{RP} \supseteq \text{NP}$, even when the bidders are unit-demand. Our lower bound raises the question of how good outcomes polynomially-bounded bidders may discover in such auctions.

To answer this question, we propose a novel concept of learning in auctions, termed “no-envy learning.” This notion is founded upon Walrasian equilibrium, and we show that it is both efficiently implementable and results in approximately optimal welfare, even when the bidders have valuations from the broad class of fractionally subadditive (XOS) valuations (assuming demand oracle access to the valuations) or coverage valuations (even without demand oracles). No-envy learning outcomes are a relaxation of no-regret learning outcomes, which maintain their approximate welfare optimality while endowing them with computational tractability. Our result for XOS valuations can be viewed as the first instantiation of approximate welfare maximization in combinatorial auctions with XOS valuations, where both the designer and the agents are computationally bounded and agents are strategic. Our positive and negative results extend to many other simple auction formats that have been studied in the literature via the smoothness paradigm.

Our positive results for XOS valuations are enabled by a novel Follow-The-Perturbed-Leader algorithm for settings where the number of experts and states of nature are both infinite, and the payoff function of the learner is non-linear. We show that this algorithm has applications outside of auction settings, establishing big gains in a recent application of no-regret learning in security games. Our efficient learning result for coverage valuations is based on a novel use of convex rounding schemes and a reduction to online convex optimization.

*Supported by a Microsoft Research Faculty Fellowship, and NSF Award CCF-0953960 (CAREER) and CCF-1551875. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing.

†This work was done in part while the author was visiting the Simons Institute for the Theory of Computing.

1 Introduction

A central challenge in Algorithmic Mechanism Design is to understand the effectiveness and limitations of mechanisms to induce economically efficient outcomes in a computationally efficient manner. A practically relevant and most actively studied setting for performing this investigation is that of *combinatorial auctions*.

This setting involves a seller with a set $[m]$ of indivisible items, which he wishes to sell to a set $[n]$ of buyers. Each buyer $i \in [n]$ is characterized by a valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$, assumed monotone, which maps each bundle S_i of items to the buyer's value $v_i(S_i)$ for this bundle. This function is known to the buyer, but is unknown to the seller and the other buyers. The seller's goal is to find a partition $S_1 \sqcup S_2 \sqcup \dots \sqcup S_n = [m]$ of the items together with prices p_1, \dots, p_n so as to maximize the total welfare resulting from allocating bundle S_i to each buyer i and charging him p_i . The total buyer utility from such an allocation would be $\sum_i (v_i(S_i) - p_i)$ and the seller's revenue would be $\sum_i p_i$, so the total welfare from such an allocation would simply be $\sum_i v_i(S_i)$.

Given the seller's uncertainty about the buyer's valuations, she needs to interact with them to select a good allocation. However, the buyers are strategic, aiming to optimize their own utility, $v_i(S_i) - p_i$. Hence, the seller needs to design her allocation and price computation rules carefully so that a good allocation is found despite the agents' strategization in response to these rules. How much of the optimal welfare can the seller guarantee?

A remarkable result in Economics is that welfare can be exactly optimized, as long as we have unbounded computational and communication resources, via the celebrated VCG mechanism [Vic61, Cla71, Gro73]. This mechanism asks bidders to report their valuations, uses their reports at face value to select an optimal partition of the items, and computes payments in a way that it is in the best interest of all bidders to truthfully report their valuations; in particular, it is a *dominant strategy truthful* mechanism, and because of its truthfulness it guarantees that an optimal allocation is truly selected.

Despite its optimality and truthfulness, the VCG mechanism is overly demanding in terms of both computation and communication. Reporting the whole valuation functions is too expensive for the bidders to do for most interesting types of valuations. Moreover, optimizing welfare exactly with respect to the reported valuations is also difficult in many cases. Unfortunately, if we are only able to do it approximately, the truthfulness of the VCG mechanism disappears, and no welfare guarantees can be made. Even with computational concerns set aside, it is widely acknowledged that the VCG mechanism is rarely used in practice [AM06]. At the same time, many practical scenarios involve the allocation of items through simple mechanisms which are often not centrally designed and non-truthful. Take eBay, for example, where several different items are sold simultaneously and sequentially via ascending price and other types of auctions. Or consider sponsored search where several keywords are auctioned simultaneously and sequentially using generalized second price auctions. For most interesting families of valuations such environments induce non truthful behavior, and are thus difficult to study analytically.

The prevalence of such simple decentralized auction environments provides motivation for a quantitative analysis of the quality of outcomes in simple non-truthful mechanisms. A growing volume of research has taken up this challenge, developing tools for studying the welfare guarantees of non-truthful mechanisms; see e.g. [Bik99, CKS08, BR11, HKMN11, FKL12, ST13, FFGL13]. Using the approximation perspective, this literature bounds the Price-of-Anarchy (PoA) of simple non-truthful mechanisms, and has provided remarkable insights into their economic efficiency.

To illustrate these results, let us consider Simultaneous Second Price Auctions, which we will abbreviate to "SiSPAs" in the remainder of this paper. While we focus our attention on these auctions, our results extend to the most common other forms of auctions studied in the PoA

literature; see Section 7 for a discussion. As implied by its name, a SiSPA asks every bidder to bid on each of the items separately and allocates each item using a second price auction based on the bids submitted solely for this item.

Facing a SiSPA, a bidder whose valuation is non-additive is not able to express his complex preferences over bundles of items. It is thus a priori not clear how he will bid, and what the resulting welfare will be. One situation where a prediction can be made is when the bidders have some information about each other, either knowing each other’s valuations, or knowing a distribution from which each others valuations are drawn. In this case, we can study the SiSPA’s Nash or Bayesian Nash equilibrium behavior, computing the welfare in equilibrium. Remarkably, the work on the PoA of mechanisms has shown that the equilibrium welfare of SiSPAs (and of other types of simple auctions) is guaranteed to be within a constant factor of optimum, even when the bidders’ valuations are subadditive [FFGL13].¹ When bidders have no information about each other, the problem becomes ill-posed, as it is impossible for the bidders to form beliefs about each others bids in order to choose their own bid.

A way out of the conundrum comes from the realization that simple mechanisms often occur repeatedly, involving the same set of bidders; think sponsored search. In such a setting it is natural to assume that bidders engage in learning to compute their new bids as a function of their experience so far. One of the most standard types of learning behavior is that of *no-regret learning*. A bidder’s bids over T executions of a SiSPA satisfy the no-regret learning guarantee if the bidder’s cumulative utility over the T executions is within an additive $o(T)$ of the cumulative utility that the bidder would have achieved from the best in hindsight vector of bids b_1, \dots, b_m , if he were to place the same bid b_j on item j in all T executions of the SiSPA. Assuming that bidders use no-regret learning to update their bids in repeated executions of a SiSPA (or other types of simple auctions) the afore-referenced work has shown that the average bidder welfare across the T executions is within a constant factor of the optimal welfare, even when the bidders’ valuations are subadditive [FFGL13].

These guarantees are astounding, especially given the intractability results for dominant strategy truthful mechanisms, which hold even when the bidders have submodular valuations [Dob11, DV12, DV15]—a family of valuations that is smaller than subadditive.² However, moving to simple non-truthful auctions does not come without a cost. Cai and Papadimitriou [CP14] have recently established intractability results for computing Bayesian-Nash equilibria in SiSPAs, even for quite simple types of valuations, namely mixtures of additive and unit-demand [CP14].³ At the same time, implementing no-regret learning in combinatorial auctions is quite tricky as the action space of the bidders explodes. For example, in SiSPAs there is a continuum of possible bid vectors that a bidder may submit and, even if we tried to discretize this set, their number would typically be exponential in the number of items in order to maintain a good approximation from the discretization. Unfortunately, no-regret algorithms typically require in every step computation that is linear in the number of available actions, hence in our case exponential in the number of items.

An important open question in the literature has thus been whether this obstacle can be overcome via specialized no-regret algorithms that only need polynomial computation. Our first result shows that this obstacle is insurmountable. We show that in one of the most basic settings where no-regret learning is non-trivial, it cannot be implemented in polynomial-time unless $\text{RP} \supseteq \text{NP}$.

Theorem 1. *Suppose that a unit-demand bidder whose value for each item $i \in [m]$ is v participates in T executions of a SiSPA. Unless $\text{RP} \supseteq \text{NP}$, there is no learning algorithm running in time*

¹A *subadditive* valuation v is one satisfying $v(S \cup T) \leq v(S) + v(T)$, for all $S, T \subseteq [m]$.

²A *submodular* valuation v is one satisfying $v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$, for all $S, T \subseteq [m]$.

³A *unit-demand* valuation v is one satisfying $v(S) = \max_{i \in S} v(\{i\})$, for all $S \subseteq [m]$.

polynomial in m , v , and T and whose regret is any polynomial in m , v , and $1/T$. The computational hardness holds even when the learner faces *i.i.d.* samples from a fixed distribution of competing bids, and whether or not no-overbidding is required of the bids produced by the learner.

Note that our theorem proves an intractability result even if pseudo-polynomial dependence on the description of v is permitted in the regret bound and the running time. The *no-overbidding assumption* mentioned in the statement of our theorem represents a collection of conditions under which no-regret learning in second-price auctions gives good welfare guarantees [CKS08, FFGL13]. An example of such no-overbidding condition is this: For each subset $S \subseteq [m]$, the sum of bids across items in S does not exceed the bidder’s value for bundle S . Sometimes this condition is only required to hold on average. It will be clear that our hardness easily applies whether or not no-overbidding is imposed on the learner, so we do not dwell on this issue more in this paper.

How can we show the in-existence of computationally efficient no-regret learning algorithms? A crucial (and general) connection that we establish in this paper is that it suffices to prove an inapproximability result for a corresponding offline combinatorial optimization problem. More precisely, we prove Theorem 1 by establishing an inapproximability result for an offline optimization problem related to SiSPAs, together with a “transfer theorem” that transfers in-approximability from the offline problem to intractability for the online problem. The transfer theorem is a generic statement applicable to any online learning setting. In particular, we show the following; see Section 3 for details.

1. In SiSPAs, finding the best response payoff against a polynomial-size supported distribution of opponent bids is strongly NP-hard to additively approximate for a unit-demand bidder. Another way to say this is that one step of a specific learning algorithm, namely Follow-The-Leader (FTL), is inapproximable. See Theorems 18 and 5.
2. In any setting where finding an optimum for an explicitly given distribution of functions over some set \mathcal{F} is hard to additively approximate, no efficient no-regret learner against sequences of functions from \mathcal{F} exists, unless $\text{RP} \supseteq \text{NP}$. This result is generic, saying that whenever one step of FTL is inapproximable, there is no no-regret learner. See Theorem 19.

The intractability result of Theorem 1 casts shadow in the ability of computationally bounded learners to achieve no-regret guarantees in combinatorial auctions where their action space explodes with the number of items and the number of items is large. We have shown this for SiSPAs, but our techniques easily extend to Simultaneous First Price Auctions, and we expect to several other commonly studied mechanisms for which PoA bounds are known. With the absence of efficiently implementable learning algorithms, it is unclear when we should expect computationally bounded bidders to actually converge to approximately efficient outcomes in these auctions.

From a design standpoint it may be interesting to identify conditions for the bidder valuations and the format of the auction under which no-regret learning is both efficiently implementable and leads to approximately optimal outcomes. While this direction is certainly interesting, it would not address the question of what welfare we should expect of SiSPAs and other simple auctions that have been studied in the literature, or how much of the PoA bounds can be salvaged for computationally bounded bidders. Moreover, recent results of Braverman et al. [BMW16] show that for a large class of auction schemes where no-regret algorithms are efficiently computable, no-better than a logarithmic in the number of items welfare guarantee can be achieved (which is achievable by the single-bid auction of [DMSW15]).

We propose an alternative approach to obtaining robust welfare guarantees of simple auctions for computationally bounded players by introducing a new type of learning dynamics, which we call

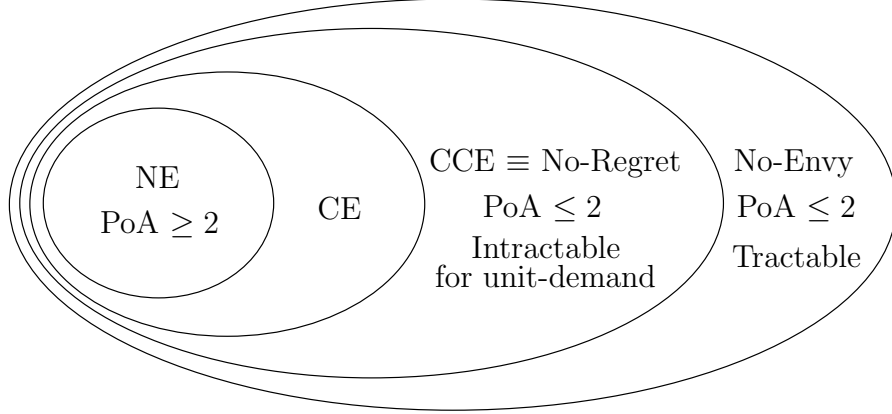


Figure 1: We depict the state of the world for simultaneous second price auctions with XOS bidders. NE denotes the set of Nash equilibria, CE the set of correlated equilibria, CCE the set of coarse correlated equilibria which are equivalent to the limit empirical distributions of no-regret dynamics. Last with No-Envy we denote the limit empirical distributions of no-envy dynamics. PoA refers to the ratio of the optimal welfare over the worst welfare achieved by any solution in each set. Tractability in this figure refers to the existence of polynomial time decentralized algorithms that each player can invoke in the agnostic setting and converge to the no-regret or no-envy condition at a polynomial error rate.

no-envy, and which are founded upon the concept of Walrasian equilibrium. In all our results, no-envy learning outcomes are a super-set of no-regret learning outcomes. We show that this super-set simultaneously achieves two important properties: i) while being a broader set, it still maintains the welfare guarantees of the set of no-regret learning outcomes established via PoA analyses; ii) there exist computationally efficient no-envy learning algorithms; when these algorithms are used by the bidders, their joint behavior converges (in a decentralized manner) to the set of no-envy learning outcomes for a large class of valuations (which includes submodular). Thus no-envy learning provides a way to overcome the computational intractability of no-regret learning in auctions with implicitly given exponential action spaces. We describe our results in the following section. We will focus our attention on SiSPAs but the definition of no-envy learning naturally extends to any mechanism and all our positive results extend to a large class of smooth mechanisms; see Section 7.

1.1 No-Envy Dynamics: Computation and Welfare.

No-envy dynamics is a twist to no-regret dynamics. Recall that in no-regret dynamics the requirement is that the cumulative utility of the bidder after T rounds be within an additive $o(T)$ error of the optimum utility he would have achieved had he played the best fixed bid in hindsight. In no-envy dynamics, we require that the bidder's cumulative utility be within an additive $o(T)$ of the optimum utility that he would have achieved if he was allocated the best in hindsight fixed bundle of items in all rounds and paid the price of this bundle in each round. The guarantee is inspired by Walrasian equilibrium: In auctions, the prices that a bidder faces on each bundle of items is determined by the bids of the other bidders. Viewed as a price-taker, the bidder would want to achieve utility at least as large as the one he would have achieved if he purchased his favorite bundle at its price. No-envy dynamics require that the average utility of the bidder across T rounds is within $o(1)$ of what he would have achieved by purchasing the optimal bundle at its average price in hindsight.

Inspired by Walrasian equilibrium, no-envy learning defines a natural benchmark against which

to evaluate an online sequence of bids. It is easy to see that in SiSPAs the no-regret learning requirement is stronger than the no-envy learning requirement. Indeed, the no-envy requirement is implied by the no-regret requirement against a subset of all possible bid vectors, namely those in $\{0, +\infty\}^m$. So no-envy learning is more permissive than no-regret learning, allowing for a broader set of outcomes. This not true necessarily for other auction formats, but it holds for the types of valuation functions and auctions studied in this paper; (see proof of Lemma 6 and Definition 12). In particular, in all our no-envy learning upper bounds the set of outcomes reachable via no-envy dynamics is always a superset of the outcomes reachable via no-regret dynamics. Moreover, this is true even if the no-envy dynamics are constrained to not overbid.

To summarize, for all types of valuations and auction formats studied in this paper, no-envy learning is a relaxation of no-regret learning, permitting a broader set of outcomes. While no-regret learning outcomes are intractable, we show that this broader set of outcomes is tractable. At the same time, we show that this broader set of outcomes maintains approximate welfare optimality. So we have increased the set of possible outcomes, but maintained their economic efficiency and endowed them with computational efficiency.

We proceed to describe our results for the computational and economic efficiency of no-envy learning. Before proceeding, we should point out that, while in our world the no-envy learning guarantee is a relaxation of the no-regret learning guarantee, the problem of implementing no-envy learning sequences remains similarly challenging. Take SiSPAs, for example. As we have noted no-envy learning is tantamount to requiring the bidder to not have regret against all bid vectors in $\{0, +\infty\}^m$. This set is exponential in the number of items m , so it is unclear how to run an off-the-shelf no-regret learner efficiently. In particular, we are still suffering from the combinatorial explosion in the number of actions, which lead to our lower bound of Theorem 1. Yet the curse of dimensionality is now much more benign. Our upper bounds, discussed next, establish that we can harness big computational savings when we move from competing against any bid vector to competing against bid vectors in $\{0, +\infty\}^m$. Except to do this we still need to develop new general-purpose, no-regret algorithms for online learning settings where the number of experts is exponentially large and the cost/utility functions are arbitrary.

1.1.1 Efficient No-Envy Learning.

We show that no-envy learning can be efficiently attained for bidders with fractionally subadditive (XOS) valuations. A valuation $v(\cdot)$ belongs to this family if for some collection of vectors $\mathcal{V} = (v^\ell)_\ell$, where each $v^\ell \in \mathbb{R}_+^m$, it satisfies:

$$v(S) = \max_{v^\ell \in \mathcal{V}} \sum_{j \in S} v_j^\ell, \forall S \subseteq [m]. \quad (1)$$

Note that the XOS class is larger than that of submodular valuations. In many applications, the set \mathcal{V} describing an XOS valuation may be large. Thus instead of inputting this set explicitly into our algorithms, we will assume that we are given an oracle, which given S returns the vector $v^\ell \in \mathcal{V}$ such that $v(S) = \sum_{j \in S} v_j^\ell$. Such an oracle is known as an *XOS oracle* [DS06, Fei06]. We will also sometimes assume, as it is customary in Walrasian equilibrium, that we are given access to a *demand oracle*, which given a price vector $p \in \mathbb{R}_+^m$ returns the bundle S maximizing $v(S) - \sum_{j \in S} p_j$. We show the following.

Theorem 2. *Consider a bidder with an XOS valuation $v(\cdot)$ participating in a sequence of SiSPAs. Assuming access to a demand and an XOS oracle for $v(\cdot)$,⁴ there exists a polynomial-time algorithm*

⁴For submodular valuations this is equivalent to assuming access only to demand oracles, as XOS oracles can be simulated in polynomial time assuming demand oracles [DNS10]

for computing the bidder’s bid vector b^t at every time step t such that after T iterations the bidder’s average utility satisfies:

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T u(b^t) \right] \geq \max_S \left(v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - O \left(\frac{m^2(D+H)}{\sqrt{T}} \right), \quad (2)$$

where $\hat{\theta}_j^T$ is the average cost of item j in the T executions of the SiSPA as defined by the bids of the competing bidders, D is an upper bound on the competing bid for any item and H is an upper bound on $\max_S v(S)$. The learning algorithm with the above guarantee also satisfies the no overbidding condition that the sum of bids for any set of items is never larger than the bidder’s value for that set. Moreover, the guarantee holds with no assumption about the behavior of competing bidders. Finally, extensions of this algorithm to other smooth mechanisms are provided in Section 7.

The proof of Theorem 2 is carried out in three steps, of which the first and last are specific to SiSPAs, while the second provides a general-purpose Follow-The-Perturbed-Leader (FTPL) algorithm in online learning settings where the number of experts is exponentially large and the cost/utility functions are arbitrary:

1. The first ingredient is simple, using the XOS oracle to reduce no-envy learning in SiSPAs to no-regret learning in a related “online buyer’s problem,” where the learner’s actions are not bid vectors but instead what bundle to purchase, prior to seeing the prices; see Definition 6. Theorem 6 provides the reduction from no-envy learning to this problem using XOS oracles. This reduction can also be done (albeit starting from approximate no-envy learning) for several mechanisms that have been analyzed through the smoothness framework of [ST13] as we elaborate in Section 7.
2. The second step proposes a FTPL algorithm for general online learning problems where the learner chooses some action $a^t \in A$ and the environment chooses some state $\theta^t \in \Theta$, from possibly infinite, unstructured sets A and Θ , and where the learner’s reward is tied to these choices through some function $u(a^t, \theta^t)$ that need not be linear. Since A need not have finite-dimensional representation and u need not be linear, we cannot efficiently perturb (either explicitly or implicitly) the cumulative rewards of the elements in A as required in each step of FTPL [KV05]; see [BCB12] and its references for an overview of such approaches. Instead of perturbing the cumulative rewards of actions in A directly, our proposal is to do this indirectly by augmenting the history $\theta^1, \dots, \theta^{t-1}$ that the learner has experienced so far with some randomly chosen fake history, and run Follow-The-Leader (FTL) subject to these augmentations. While it is not a priori clear whether our perturbation approach is a useful one, it is clear that our proposed algorithm only needs an offline optimization oracle to be implemented, as each step is an FTL step after the fake history is added. When applying this algorithm to the online buyer’s problem from Step 1, the required offline optimization oracle will conveniently end up being simply a demand oracle.

Our proposed general purpose learner is presented in Section 5. The way our learner accesses function u is via an optimization oracle, which given a finite multiset of elements from Θ outputs an action in A that is optimal against the uniform distribution over the multiset. See Definition 7. In Theorem 8, we bound the regret experienced by our algorithm in terms of u ’s stability. Roughly speaking, the goal of our randomized augmentations of the history in each step of our learning algorithm is to smear the output of the optimization oracle applied to the augmented sequence over A , allowing us to couple the choices of Be-The-Perturbed-Leader and Follow-The-Perturbed-Leader.

3. To apply our general purpose algorithm from Theorem 8 to the online buyer’s problem for SiSPAs from Step 1, we need to bound the stability of the bidder’s utility function subject to a good choice of a history augmentation sampler. This is done in Section 5.1. There turns out to be a simple sampler for our application here, where only one price vector is added to the history, whose prices are independently distributed according to an exponential distribution with mean $O(\sqrt{T})$ and variance $O(T)$.
4. While our motivation comes from mechanism design, our FTPL algorithm from Step 2 is general purpose, and we believe it will find applications in other settings. We provide some relevant discussion in Section F, where we show how our algorithm implies regret bounds independent of $|A|$ when Θ is finite, as well as quantitative improvements on the regret bounds of a recent paper of Balacn et al. for security games [BBHP15].

In the absence of demand oracles, we provide positive results for the subclass of XOS called coverage valuations. To explain these valuations, consider a bidder with k needs, $1, \dots, k$, associated with values a_1, \dots, a_k . There are m available items, each covering a subset of these needs. So we can view each item as a set $\beta_i \subseteq [k]$ of the needs it satisfies. The value that the bidder derives from a set $S \subseteq [m]$ of the items is the total value from the needs that are covered, namely:

$$v(S) = \sum_{\ell \in \cup_j \beta_j} a_\ell. \quad (3)$$

Theorem 3. *Consider a bidder with an explicitly given coverage valuation $v(\cdot)$ participating in a sequence of SiSPAs. There exists a polynomial-time algorithm for computing the bidder’s bid vector b^t at every time step t such that after T iterations the bidder’s utility satisfies:*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T u(b^t) \right] \geq \max_S \left(\left(1 - \frac{1}{e} \right) v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - 3m \frac{H + \sqrt{D}}{\sqrt{T}}, \quad (4)$$

where $\hat{\theta}_j^T$, H and D are as in Theorem 2, and the algorithm satisfies the same no overbidding condition stated in that theorem. There is no assumption about the behavior of the competing bidders, and extensions of this algorithm to other smooth mechanisms are provided in Section 7.

Notice that our no-envy guarantee (4) in Theorem 3 has incurred a loss of a factor of $1 - 1/e$ in front of $v(S)$, compared to the no-envy guarantee (2). This relaxed guarantee is an even broader relaxation of the no-regret guarantee. Still, as we show in the next section this does not affect our approximate welfare guarantees. We prove Theorem 3 via an interesting connection between the online buyer’s problem for coverage valuations and the convex rounding approach for truthful welfare maximization proposed by [DRY11]. In the online buyer’s problem, recall that the buyer needs to decide what set to buy at each step, prior to seeing the prices. It is natural to have the buyer include each item to his set independently, thereby defining an expert for all points $x \in [0, 1]^m$, where x_i is the probability that item i is included. It turns out that the expected utility of the buyer under such distribution x is not necessarily convex, so this choice of experts turns our online learning problem non-convex. Instead we propose to massage each expert $x \in [0, 1]^m$ into a distribution $D(x) \in \Delta(2^{[m]})$ and run online learning on the massaged experts. In Definition 15 we put forth conditions for the massaging operation $D(\cdot)$ under which online learning becomes convex and gives approximate no-regret (Lemma 23). We then instantiate $D(\cdot)$ with the Poisson sampling of [DRY11], establishing Theorem 3. Our approach is summarized in Section 6 and the details can be found in Appendix E.

1.1.2 Welfare Maximization.

Arguably one of the holy grails in Algorithmic Mechanism Design, since its inception, has been to obtain polynomial-time mechanisms optimizing welfare in incomplete information settings. We show that SiSPAs achieve constant factor approximation welfare guarantees for the broad class of XOS valuations at every no-envy or approximate no-envy learning outcome. Thus the relaxation from no-regret to no-envy learning does not degrade the quality of the welfare guarantees, and has the added benefit that no-envy outcomes can be attained by computationally bounded players in a decentralized manner, using our results from the previous section. In Section 7, we show that this property applies to a large class of mechanisms that have been analyzed in the literature via the *smoothness* paradigm [ST13].

Corollary 4. *When each bidder $i \in \{1, \dots, n\}$ participating in a sequence of SiSPAs has an XOS valuation (endowed with a demand and XOS oracle) or an explicitly given coverage valuation $v_i(\cdot)$, there exists a polynomial-time computable learning algorithm such that, if each bidder i employs this algorithm to compute his bids b_i^t at each step t , then after T rounds the average welfare is guaranteed to be at least:*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T SW(b_1^t, \dots, b_n^t) \right] \geq \frac{1}{2} \left(1 - \frac{1}{e} \right) \text{OPT}(v_1, \dots, v_n) - O \left(m^2 \cdot n \cdot \max_{S,i} v_i(S) \sqrt{\frac{1}{T}} \right). \quad (5)$$

If all bidders have XOS valuations with demand and XOS oracles the factor in front of OPT is 1/2.

We regard Corollary 4, in particular our result for XOS valuations with demand queries, as alleviating the intractability of no-regret learning in simple auctions. It also provides a new perspective to mechanism design, namely mechanism design with no-envy bidders. In doing so, it proposes an answer to the question raised by [FGL15] about whether demand oracles can be exploited for welfare maximization with submodular bidders. We show a positive answer for the bigger class of XOS valuations, albeit with a different solution concept. (It still remains open whether there exist poly-time dominant strategy truthful mechanisms for submodular bidders with demand queries.) We believe that no-envy learning is a fruitful new approach to mechanism design, discussing in Section 7 the meaning of the solution concept outside of SiSPAs.

2 Preliminaries

We analyze the online learning problem that a bidder faces when participating in a sequence of repeated executions of a simultaneous second price auction (SiSPA) with m items. While we focus on SiSPAs our results extend to the most commonly studied formats of simple auctions, as discussed in Section 7. A sequence of repeated executions of a SiSPA corresponds to a sequence of repeated executions of a game involving n players (bidders). At each execution t , each player i submits a bid b_{ij}^t on each item j . We denote by b_i^t the vector of bidder i 's bids at time t and by b^t the profile of bids of all players on all items. Given these bids, each item is given to the bidder who bids for it the most and this bidder pays the second highest bid on the item. Ties are broken according to some arbitrary tie-breaking rule. Each player i has some fixed (across executions) valuation $v_i : 2^{[m]} \rightarrow \mathbb{R}_+$ over bundles of items. If at time t he ends up winning a set of items S^t and is asked to pay a price of θ_j^t for each item $j \in S^t$, then his utility is $v_i(S^t) - \sum_{j \in S^t} \theta_j^t$, i.e. his utility is assumed quasi-linear. An important class of valuations that we will consider in this paper is that of XOS valuations, defined in Equation (1), which are a super-set of submodular valuations but a subset of subadditive valuations. We will also consider the class of coverage valuations, defined in Equation 3, which are a subset of XOS. Different results will consider different types of access to

an XOS valuation through an XOS oracle, a demand oracle, or a value oracle, as described in the introduction. For more properties of these oracles see [DNS10].

Online bidding problem. From the perspective of a single player i , all that matters to him to calculate his utility in a SiSPA is the highest bid submitted by the other bidders on each item j , as well as the probability that he wins each item j if he ties with the highest other bid on that item. For simplicity of notation, we will assume throughout the paper that the player always loses an item when he ties first. All our results, both positive and negative, easily extend to the more general case of arbitrary bid-profile dependent tie-breaking. Since we will analyze learning from the perspective of a single player, we will drop the index of player i . For a fixed bid profile of the opponents, we will refer to the highest other bid on item j as the threshold of item j and denote it with θ_j . We also denote with $\theta = (\theta_1, \dots, \theta_m)$. The player wins an item j if he submits a bid $b_j > \theta_j$ and loses the item otherwise. When he wins item j , he pays θ_j . We are interested in learning algorithms that achieve a no-regret guarantee even when the thresholds of the items are decided as it is customary by an adversary. Thus, the online learning problem that a player faces in a simultaneous second price auction is defined as follows:

Definition 1 (Online bidding problem). *At each execution/day/time/step t , the player picks a bid vector b^t and the adversary picks adaptively (based on the history of the player's past bid vectors but not on the bidder's current bid vector b^t) a threshold vector θ^t . The player wins the set $S(b^t, \theta^t) = \{j \in [m] : b_j^t > \theta_j^t\}$ and gets reward:*

$$u(b^t, \theta^t) = v(S(b^t, \theta^t)) - \sum_{j \in S(b^t, \theta^t)} \theta_j^t. \quad (6)$$

We allow a learning algorithm to be randomized, i.e. submit a random bid vector at each step whose distribution may depend on the history of past threshold vectors. We will evaluate a learning algorithm based on its *regret* against the best fixed bid vector in hindsight.

Definition 2 (Regret of Learning Algorithm). *The expected average regret of a randomized online learning algorithm against a sequence $\theta^{1:T} = (\theta^1, \dots, \theta^T)$ of threshold vectors is:*

$$\text{Reg}(\theta^{1:T}) = \mathbb{E}_{b^{1:T}} \left[\sup_{b^*} \frac{1}{T} \sum_{t=1}^T (u(b^*, \theta^t) - u(b^t, \theta^t)) \right], \quad (7)$$

where recall that b^t is random and depends on $\theta^{1:t-1}$, as specified by the online learning algorithm. The regret $r(T)$ of the algorithm against an adaptive adversary is the maximum regret against any adaptively chosen sequence of T threshold vectors. An algorithm has polynomial regret rate if $r(T) = \text{poly}(T^{-1}, m, \max_S v(S))$.

3 Hardness of No-Regret Learning

We will show that there does not exist an efficiently computable learning algorithm with polynomial regret rate for the online bidding problem for SiSPAs unless $\text{RP} \supseteq \text{NP}$, proving a proof of Theorem 1. We first examine a related offline optimization problem which we show is NP-hard to approximate to within a small additive error. We then show how this inapproximability result implies the non-existence of polynomial-time no-regret learning algorithms for SiSPAs unless $\text{RP} \supseteq \text{NP}$. Throughout this section we will consider the following very restricted class of valuations: the player is unit-demand and has a value v for getting any item, i.e. his value for any set of items is given by $v(S) = v \cdot 1\{S \neq \emptyset\}$. Our intractability results are strong intractability results in the sense that they hold even if we assume that v is provided in the input in unary representation.

Optimal Bidding Against An Explicit Threshold Distribution is Hard. We consider the following optimization problem:

Definition 3 (Optimal Bidding Problem). *A distribution D of threshold vectors θ over a set of m items is given explicitly as a list of k vectors, where D is assumed to choose a uniformly random vector from the list. A bidder has a unit-demand valuation with the same value v for each item, given in unary. The problem asks for a bid vector that maximizes the bidder’s expected utility against distribution D . In fact, it only asks to compute the expected value from an optimal bid vector, i.e.*

$$\sup_b \mathbb{E}_{\theta \sim D} [u(b, \theta)] = \sup_b \left\{ v \cdot \Pr [\exists j \in [m] : b_j > \theta_j] - \sum_{j \in [m]} \theta_j \cdot \Pr [b_j > \theta_j] \right\}. \quad (8)$$

We show that the optimal bidding problem is NP-hard via a reduction from r -regular set-cover. In fact we show that it is hard to approximate, up to an additive approximation that is inverse-polynomially related to the input size. This will be useful when using the hardness of this problem to deduce the in-existence of efficiently computable learning algorithms with polynomial regret rates.

Theorem 5 (Hardness of Approximately Optimal Bidding). *The optimal bidding problem is NP-hard to approximate to within an additive ξ even when: the k threshold vectors in the support of (the explicitly given distribution) D take values in $\{1, H\}^m$, $H = k^2 \cdot m^2$, $v = 2 \cdot k \cdot m$ and $\xi = \frac{1}{2k}$.*

An interesting interpretation Theorem 5 is that the Follow-The-Leader (FTL) algorithm is intractable in SiSPAs for unit-demand bidders. Indeed, every step of FTL needs to find a bid vector that is a best response to the empirical distribution of the threshold vectors that have been encountered so far. See Theorem 18 in Appendix A and the discussion around this theorem.

Efficient No-Regret implies Poly-time Approximately Optimal Bidding. Given the hardness of *optimal bidding* in SiSPAs, we are ready to sketch the proof of our main impossibility result (Theorem 1) for *online bidding* in SiSPAs. Our result holds even if the possible threshold vectors that the bidder may see take values in some known discrete finite set. It also holds even if we weaken the regret requirements of the online bidding problem, only requiring that the player achieves no-regret with respect to bids of the form $\{0, v/2m\}^m$, i.e., the bid on each item is either 0 or an $2m$ -th fraction of the player’s value. Notice that any such bid is a non-overbidding bid. Hence, the no-regret requirement that we impose is weaker than achieving no-regret against any fixed bid/any fixed no-overbidding bid. We will refer to the afore-described weaker learning task as the *simplified online bidding problem*. We sketch here how to deduce from the inapproximability of optimal bidding the impossibility of polynomial-time no-regret learning (even for the simplified online bidding problem), deferring full details to Appendix A.2.

Proof sketch of Theorem 1. We present the structure of our proof and the challenges that arise, leaving details for Section A.2. Consider a hard distribution D for the optimal bidding problem from Theorem 5, and let b^* be the bid vector that optimizes the expected utility of the bidder when a threshold vector is drawn from D . Also, let u^* be the corresponding optimal expected utility. (Theorem 5 says that approximating u^* is NP-hard.) Now let us draw T i.i.d. samples $\theta = (\theta^1, \dots, \theta^T)$ from D . Clearly, if T is large enough, then, with high probability, the expected utility \tilde{u}_θ of b^* against the uniform distribution over $\theta^1, \dots, \theta^T$ is approximately equal to u^* .

Now let us present the sequence $\theta^1, \dots, \theta^T$ to a no-regret learning algorithm. The learning algorithm is potentially randomized so let us call \hat{u}_θ the expected average utility (over the randomness in the algorithm and keeping sequence θ fixed) that the algorithm achieves when facing

the sequence of threshold vectors $\theta^1, \dots, \theta^T$. If the regret of the algorithm is $r(T)$, this means that $\hat{u}_\theta \geq \sup_b (\frac{1}{T} \sum_{t=1}^T u(b, \theta^t)) - r(T) \geq \tilde{u}_\theta - r(T)$. In particular, if $r(T)$ scales polynomially with $1/T$ then, for large enough T , \hat{u}_θ is lower bounded by \tilde{u}_θ (minus some small error), and hence by u^* (minus some small error). Hence, \hat{u}_θ (plus some small error) provides an upper bound to u^* . Moreover, if we run our no-regret learning algorithm a large enough number of times N against the same sequence of threshold vectors and average the average utility achieved by the algorithm in these N executions, we can get a very good estimate of \hat{u}_θ , and hence a very good upper bound for u^* , with high probability. The paragraph “Upper Bound” in Appendix A.2 gives the details of this part.

The challenge that we need to overcome now is that, in principle, the expected average utility \hat{u}_θ of our no-regret learner against sequence $\theta^1, \dots, \theta^T$ could be much larger than $\sup_b (\frac{1}{T} \sum_{t=1}^T u(b, \theta^t))$ and hence \tilde{u}_θ and u^* , as the algorithm is allowed to change its bid vector in every step. We need to argue that this cannot happen. In particular, we would like to upper bound \hat{u}_θ by u^* . We do this via a Martingale argument exploiting the randomness in the choice of the sequence θ . Using Azuma’s inequality, we show that for large enough T , the \hat{u}_θ is upper bounded by u^* plus some small error with high probability. In fact we show something stronger: if T is large enough then, with high probability, u^* plus some small error upper bounds the algorithm’s average utility (not just average expected utility), where now both the threshold and the bid vectors are left random. Hence, we can argue that, with high probability, if we run our algorithm N times over a (long enough) sequence of random threshold vectors and we compute the average (across the N executions) of the average (across the T steps) utility of our algorithm, then this double average is upper bounded by u^* plus some small error. Hence, we get a lower bound on u^* . (One execution would indeed suffice, but we need to argue about the average across N executions given the way we obtain our upper bound in the previous paragraph.) The paragraph “Lower Bound” in Appendix A.2 gives the details of this part.

Overall, if we choose T and N large enough polynomials in the description of the hard instance of the optimal bidding problem from Theorem 5, then all approximation errors can be made arbitrary inverse polynomials, providing any desired (inverse polynomial) approximation to the optimal utility u^* against distribution D , with high probability. Since getting an inverse polynomial approximation is an NP-hard problem, this implies that there cannot exist a polynomial-time no-regret learning algorithm with polynomial regret rate, unless $\text{RP} \supseteq \text{NP}$. ■

4 Walrasian Equilibria and No-Envy Learning in Auctions

The hardness of no-regret learning in simultaneous auctions motivates the investigation of other notions of learning that have rational foundations and at the same time admit efficient implementations. Our inspiration in this paper comes from the study of markets and the well-studied notion of *Walrasian equilibrium*. Recall that an allocation of items to buyers together with a price on each item constitutes a Walrasian equilibrium if no buyer envies some other allocation at the current prices. That is the bundle S allocated to each buyer maximizes the difference $v(S) - p(S)$ of his value $v(S)$ for the bundle minus the cost of the bundle. Implicitly the Walrasian equilibrium postulates some degree of rationality on the buyers: given the prices of the items, each buyer wants a bundle of items such that he has *no-envy* against getting any other bundle at the current prices.

We adapt this *no-envy* requirement to SiSPAs (and other mechanisms in Section 7). In a SiSPA a player is facing a set of prices on the items, which are determined by the bids of the other players and are hence unknown to him when he is choosing his bid vector. In a sequence of repeated executions of a SiSPA, the player needs to choose a bid vector at every time-step. The fact that he

does not know the realizations of the item prices when making his choice turns the problem into a learning problem. We will say that the sequence of actions that he took satisfies the no-envy guarantee, if in the long run he does not regret not buying any fixed set S at its average price.

Definition 4 (No-Envy Learning Algorithm). *An algorithm for the online bidding problem is a no-envy algorithm if, for any adaptively chosen sequence of threshold vectors $\theta^{1:T}$ by an adversary, the bid vectors b^1, \dots, b^T chosen by the algorithm satisfy:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(b^t, \theta^t) \right] \geq \max_{S \subseteq [m]} \left(v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - \epsilon(T) \quad (9)$$

where $\hat{\theta}_j^T = \frac{1}{T} \sum_{t=1}^T \theta_j^t$ and $\epsilon(T) \rightarrow 0$. It has polynomial envy rate if $\epsilon(T) = \text{poly}(T^{-1}, m, \max_S v(S))$.

To allow for even larger classes of settings to have efficiently computable no-envy learning outcomes, we will also define a relaxed notion of no-envy. In this notion the player is guaranteed that his utility is at least some α -fraction of his value for any set S , less the average price of that set. The latter is a more reasonable relaxation in the online learning setting given that, unlike in a market setting, the players do not know the realization of the prices when they make their decision.

Definition 5 (Approximate No-Envy Learning Algorithm). *An algorithm for the online bidding problem is an α -approximate no-envy algorithm if, for any adaptively chosen sequence of threshold vectors $\theta^{1:T}$ by an adversary, the bid vectors b^1, \dots, b^T chosen by the algorithm satisfy:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(b^t, \theta^t) \right] \geq \max_{S \subseteq [m]} \left(\frac{1}{\alpha} v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - \epsilon(T). \quad (10)$$

To gain some intuition about the difference between no-envy and no-regret learning guarantees consider the following. When we compute the utility from a fixed bid vector in hindsight, then in every iteration the set of items that the player would have won is nicely correlated with that round's threshold vector in the sense that the player wins an item in that round only when the item's threshold is low. On the contrary, when evaluating the player's utility had he won a specific set of items in all rounds the player may win and pay for an item even when the price of the item is high. The results of this section imply that for XOS valuations, the no-regret condition is stronger than the no-envy condition. Hence, when we analyze no-envy learning algorithms for XOS bidders we relax the algorithm's benchmark. Correspondingly, if the bidders of a SiSPA are XOS and use no-envy learning algorithms to update their bid vectors, the set of outcomes that they may converge to is broader than the set of no-regret outcomes. So, in comparison to no-regret learning outcomes, our positive results in this section pertain to a broader set of outcomes, endowing them with computational tractability and as we will see also approximate welfare optimality.

Roadmap. In the rest of this section we reduce the no-envy learning problem to a related online learning problem, which we call the *online buyer's problem*. We show that achieving *no-envy* in the *online bidding problem* can be reduced to achieving *no-regret* in the *online buyer's problem*. Similarly, achieving α -approximate no-envy can be reduced to some form of approximate no-regret. Lastly we show that *no-envy learning* implies good welfare: if all players in the simultaneous second-price auction game follow a no-envy learning algorithm then the average welfare of the selected allocations is approximately optimal. In subsequent sections we will provide efficiently computable no-envy or approximate no-envy algorithms for the *online buyer's problem*. Finally, our positive results extend to the most commonly studied mechanisms through the smoothness framework, as we elaborate in Section 7.

4.1 Online Buyer’s Problem

We first show that we can reduce the no-envy learning problem to a related online learning problem, which we call the *online buyer’s problem*.

Definition 6 (Online buyer’s problem). *Imagine a buyer with some valuation $v(\cdot)$ over a set of m items who is asked to request a subset of the items to buy each day before seeing their prices. In particular, at each time-step t an adversary picks a set of thresholds/prices θ_j^t for each item j adaptively based on the past actions of the buyer. Without observing the thresholds at step t , the buyer picks a set S^t of items to buy. His instantaneous reward is:*

$$u(S^t, \theta^t) = v(S^t) - \sum_{j \in S^t} \theta_j^t, \quad (11)$$

i.e., the buyer receives the set S^t and pays the price for each item in the set.

For simplicity, we overload notation and denote by $u(b, \theta)$ the reward in the online bidding problem from a bid vector b and with $u(S, \theta)$ the reward in the online buyer’s problem from a set S . We relate the online buyer’s problem to the online bidding problem in SiSPAs in a black-box way, by showing that when the valuations are XOS, then any algorithm which achieves no-regret or “approximate” no-regret for the *online buyer’s problem* can be turned in a black-box and efficient manner into a no-envy algorithm for the *online bidding problem*, assuming access to an XOS oracle.

Lemma 6 (From buyer to bidder). *Suppose that we are given access to an efficient learning algorithm for the online buyer’s problem which guarantees for any adaptive adversary:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(S^t, \theta^t) \right] \geq \max_S \left(\frac{1}{\alpha} v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - \epsilon(T), \quad (12)$$

where $\hat{\theta}_j^T = \frac{1}{T} \sum_{t=1}^T \theta_j^t$. Then we can construct an efficient α -approximate no-envy algorithm for the online bidding problem, assuming access to XOS value oracles. Moreover, this algorithm never submits an overbidding bid.

A trivial example: efficient no-envy for $O(\log(m))$ -capacitated XOS. Consider a buyer with a d -capacitated XOS valuation, i.e. the valuation is XOS and for any set S : $v(S) = \max_{T \subseteq S: |T| \leq d} v(T)$. If $d = O(\log(m))$, then it suffices for the buyer to achieve no-regret against sets of size d , which are 2^d . This is polynomial if $d = O(\log(m))$. Thus we can simply invoke any off-the-shelf no-regret learning algorithm, such as multiplicative weight updates [ACBFS95], where each set of d is treated as an expert, and apply it to the *online buyer’s problem*. This would be efficiently computable and would lead to a regret rate of $O(\sqrt{T \log(2^d)}) = O(\sqrt{T \cdot d})$. By Lemma 6, we then get an efficiently computable exact no-envy algorithm with the same envy rate.

The challenge addressed by our paper is to remove the bound on d , which we address in the next sections.

4.2 No-Envy Implies Approximately Optimal Welfare

We conclude by showing that if all players in a SiSPA use an α -approximate no-envy learning algorithm, then the average welfare is a 2α -approximation to the optimal welfare, less an additive error term stemming from the envy of the players. In other words the price of anarchy of α -approximate no-envy dynamics is upper bounded by 2α .

Theorem 7. *If n players participating in repeated executions of a SiSPA use an α -approximate no-envy learning algorithm with envy rate $\epsilon(T)$ and which does not overbid, then in T executions of the SiSPA the average bidder welfare is at least $\frac{1}{2\alpha}\text{OPT} - n \cdot \epsilon(T)$, where OPT is the optimal welfare for the input valuation profile $v = (v_1, \dots, v_n)$.*

5 Online Learning with Oracles

In this section we devise novel follow-the-perturbed leader style algorithms for general online learning problems. We then apply these algorithms and their analysis to get no-envy learning algorithms (Section 5.1) for the online bidding problem. In Section F we instantiate our analysis to learning problems where the adversary can only pick one among finitely many parameters and give implications of this setting to no-regret learning algorithms (Section F) for the online bidding problem, with a finite number of possible threshold vectors. In Section F.2, we also give implications to security games [BBHP15].

Consider an online learning problem where at each time-step an adversary picks a parameter $\theta^t \in \Theta$ and the algorithm picks an action $a^t \in A$. The algorithm receives a reward: $u(a^t, \theta^t)$, which could be positive or negative. We will assume that the rewards are uniformly bounded by some function of the parameter θ , for any action $a \in A$, i.e.: $\forall a \in A : u(a, \theta) \in [-f_-(\theta), f_+(\theta)]$. We will denote with $\theta^{1:t}$ a sequence of parameters $\{\theta_1, \theta_2, \dots, \theta_t\}$. Moreover, we denote with: $U(a, \theta^{1:t}) = \sum_{\tau=1}^t u(a, \theta^\tau)$, the cumulative utility of a fixed action $a \in A$ for a sequence of choices $\theta^{1:t}$ of the adversary.

Definition 7 (Optimization oracle). *We will consider the case where we are given oracle access to the following optimization problem: given a sequence of parameters $\theta^{1:t}$ compute some optimal action for this sequence:*

$$M(\theta^{1:t}) = \arg \max_{a \in A} U(a, \theta^{1:t}). \quad (13)$$

We define a new type of perturbed leader algorithms where the perturbation is introduced in the form of extra samples of parameters:

Algorithm 1 (Follow the perturbed leader with sample perturbations). *At each time-step t :*

1. *Draw a random sequence of parameters $\{x\}^t = \{x^1, \dots, \dots, x^k\}^t$ independently and based on some time-independent distribution over sequences. Both the length of the sequence and the parameter $x^i \in \Theta$ at each iteration of the sequence can be random.*
2. *Denote with $\{x\}^t \cup \theta^{1:t-1}$ the augmented sequence of parameters where we append the extra parameter samples $\{x\}^t$ at the beginning of sequence $\theta^{1:t-1}$*
3. *Invoke oracle M and play action:*

$$a^t = M(\{x\}^t \cup \theta^{1:t-1}). \quad (14)$$

Using a reduction of [HP05] (see their Lemma 12) we can show that to bound the regret of Algorithm 1 against adaptive adversaries it suffices to bound the regret against oblivious adversaries (who pick the sequence non-adaptively), of the following algorithm, which only draws the samples once ahead of time (see Appendix C.1). In subsequent sections, we analyze this algorithm and setting.

Algorithm 2 (Follow the perturbed leader with fixed sample perturbations). *Draw a random sequence of parameters $\{x\} = \{x^1, \dots, \dots, x^k\}$ based on some distribution over sequences and at the beginning of time. At each time-step t , invoke oracle M and play action: $a^t = M(\{x\} \cup \theta^{1:t-1})$.*

Perturbed Leader Regret Analysis. We give a general theorem on the regret of a perturbed leader algorithm with sample perturbations. In the sections that follow we will give instances of this analysis in two online learning settings related to no-envy and no-regret dynamics in our bidding problem and provide concrete regret bounds.

Theorem 8. *Suppose that the distribution over sample sequences $\{x\}$, satisfies the stability property that for any sequence of parameters $\theta^{1:T}$ and for any $t \in [1 : T]$:*

$$\mathbb{E}_{\{x\}} [u(M(\{x\} \cup \theta^{1:t}), \theta^t) - u(M(\{x\} \cup \theta^{1:t-1}), \theta^t)] \leq g(t) \quad (15)$$

Then the expected regret of Algorithm 2 against oblivious adversaries is upper bounded by:

$$\sup_{a^* \in A} \mathbb{E}_{\{x\}} \left[\sum_{t=1}^T (u(a^*, \theta^t) - u(a^t, \theta^t)) \right] \leq \sum_{t=1}^T g(t) + \mathbb{E}_{\{x\}} \left[\sum_{x^\tau \in \{x\}} (f_-(x^\tau) + f_+(x^\tau)) \right] \quad (16)$$

Hence, the regret of Algorithm 1 against adaptive adversaries is bounded by the same amount.

5.1 Efficient No-Envy Learning with Demand Oracles

We will apply the perturbed leader approach to the online buyer's problem we defined in Section 4.1. Then using Lemma 6 we can turn any such algorithm to a no-envy learning algorithm for the original bidding problem in second price auctions, when the valuations fall into the XOS class.

In the online buyer's problem the action space is the collection of sets $A = 2^m$, while the parameter set of the adversary is to pick a threshold θ_j for each item j , i.e. $\Theta = \mathbb{R}_+^m$. The reward $u(S, \theta)$, at each round from picking a set S , if the adversary picks a vector $\theta \in \Theta$ is given by Equation (11). We will instantiate Algorithm 2 for this problem and apply the generic approach of the previous section. We will specify the exact distribution over sample sequences that we will use and we will bound the functions $f_-(\cdot)$, $f_+(\cdot)$ and $g(\cdot)$. First, observe that the reward is bounded by a function of the threshold vector: $u(S, \theta) \in [-\|\theta\|_1, H]$, where H is an upper bound on the valuation function, i.e. $v([m]) < H$.

Optimization oracle. It is easy to see that the offline problem for a sequence of parameters $\theta^{1:t}$ is exactly a *demand oracle*, where the price on each item j is its average threshold $\hat{\theta}_j^t$ in hindsight.

Single-sample exponential perturbation. We will use the following sample perturbation: we will only add one sample $x \in \Theta$, where the coordinate x_i of the sample is distributed independently and according to an exponential distribution with parameter ϵ , i.e. for any $k \geq 0$ the density of x_i at k is $f(k) = \frac{1}{2}\epsilon e^{-\epsilon k}$, while it is 0 for $k < 0$.

The most important part of the analysis is proving a stability bound for our algorithm. We provide such a proof in Appendix D.1. Given the stability bound we then apply Theorem 8 to get a bound for Algorithm 2 with a single sample exponential perturbation.

Theorem 9. *Algorithm 2 when applied to the online buyers problem with a single-sample exponential perturbation with parameter $\epsilon = \sqrt{\frac{1}{HDT}}$, where D is the maximum threshold that the adversary can pick and H is the maximum value, runs in randomized polynomial time, assuming a demand oracle and achieves regret:*

$$\sup_{a \in A} \sum_{t=1}^T \mathbb{E} [u(a, \theta^t) - u(a^t, \theta^t)] \leq 2(mD+H)m(\log(T)+1)+4m\sqrt{(mD+H)DT} = O\left(m^2(D+H)\sqrt{T}\right)$$

Theorem 9, Lemma 6 and the reduction from oblivious to adaptive adversaries, imply a polynomial time no-envy algorithm for the online bidding problem assuming access to demand and XOS oracles. If valuations are submodular, then XOS oracles can be simulated in polynomial time via demand oracles [DNS10], thereby only requiring access to demand oracles. Thus we get Theorem 2.

6 Efficient No-Envy Learning via Convex Rounding

In this section we show how to design efficient approximate no-envy learning algorithms via the use of the convex rounding technique, which has been used in approximation algorithms and in truthful mechanism design, and via online convex optimization applied to an appropriately defined online learning problem in a relaxed convex space. Though our techniques can be phrased more generally, throughout the section we will mostly cope with the concrete case where the valuation of the player is an explicitly given coverage valuation. These valuations have been well-studied in combinatorial auctions [DRY11] and are a subset of submodular valuations. Answering value and XOS queries for such valuations can be done in polynomial time [DRY11, DNS10].

Definition 8 (Coverage valuation). *A coverage valuation is given via the means of a vertex-weighted hyper-graph $\mathcal{G} = (V, E)$. Each item $j \in [m]$ corresponds to a hyper-edge. Each vertex $v \in V$ has a weight $w_v \geq 0$. The value of the player for a set S is the sum of the vertices of the hyper-graph, that is contained in the union of the hyper-edges corresponding to the items in S .*

Proving Theorem 3. Based on Lemma 6, in order to design an α -approximate no-envy algorithm for the online bidding problem, it suffices to design an efficient algorithm for the online buyer's problem with guarantees as described in Lemma 6. In the remainder of the section we will design such an algorithm for the *online buyer's problem* with $\alpha = \frac{e}{e-1}$ and for explicit coverage valuations, thereby proving Theorem 3. Subsequently, by Theorem 7 the latter will imply a price of anarchy guarantee of $\frac{2e}{e-1}$ for such dynamics. The only missing piece in the proof of Theorem 3 is the following lemma, whose full proof we defer to Appendix E.

Lemma 10. *If the bidder's valuation $v(\cdot)$ is an explicitly given coverage valuation, there exists a polynomial-time computable learning algorithm for the online buyer's problem that guarantees that for any adaptively chosen sequence of thresholds $\theta^{1:T}$ with $\theta_j^t \leq K$:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(S^t, \theta^t) \right] \geq \max_S \left(\left(1 - \frac{1}{e} \right) v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - 3m \frac{\max_{j \in [m]} v(\{j\}) + \sqrt{K}}{\sqrt{T}}, \quad (17)$$

Proof sketch. Suppose that the buyer picks a set at each iteration at random from a distribution where each item j is included independently with probability x_j to the set. Then for any vector x , the expected utility of the buyer from such a choice is $\mathbb{E}_{S^t \sim x^t} [u(S^t, \theta^t)] = V(x^t) - \langle \theta^t, x^t \rangle$, where $V(\cdot)$ is the multi-linear extension of $v(\cdot)$ and $\langle x, y \rangle$ is the inner product between vectors x and y . If $V(\cdot)$ was concave we could invoke online convex optimization algorithms, such as the projected gradient descent of [Zin03] and get a regret bound, which would imply a regret bound for the buyers problem. However, $V(\cdot)$ is not concave for most valuation classes. We will instead use a *convex rounding scheme*, which is a mapping from any vector x to a distribution over sets $D(x)$ such that $F(x) = \mathbb{E}_{S \sim D(x)} [v(S)]$ is a concave function of x . We also require that the marginal probability of each item be at most the original probability of that item in x . If the rounding scheme satisfies that for any integral x associated with set S , $F(x) \geq \frac{1}{\alpha} v(S)$, then we can call an online convex optimization algorithm on the concave function $F(x) - \langle \theta, x \rangle$. Then we show that this yields an α -approximate no-envy algorithm for the online buyers problem. ■

7 No-Envy Learning for General Mechanisms

In this section we generalize our approach to most smooth mechanisms [ST13] that have been analyzed in the literature. For ease of exposition we only focus on mechanisms for combinatorial auction settings, even though the approach could be employed for more general mechanism design settings.

A general mechanism M for a combinatorial auction setting is defined via an action space \mathcal{A}_i available to each player i , an allocation function, which maps each action profile $a \in \mathcal{A} \triangleq \mathcal{A}_1 \times \dots \times \mathcal{A}_n$ to a feasible partition $X(a)$ of the items among players, as well as a payment function, which maps an action profile to a vector of payments for $P(a)$ for each player. We denote with $X_i(a)$ and $P_i(a)$ the allocation and payment of player i . These functions could also output randomized allocations and payments, but for simplicity of notation we restrict to deterministic mechanisms.

First and foremost we need to generalize the definition of no-envy to general mechanisms other than simultaneous second price auctions. To achieve this we need to define the equivalent of a threshold vector for a general mechanism. We define the notion of a threshold-payment for a player i and a set S , which will coincide with the sum of thresholds - $\sum_{j \in S} \theta_{ij}$ - for the case of a simultaneous second price auction.

Definition 9 (Threshold Payment). *Given a set S and an action profile a , the threshold payment for player i for set S is the minimum payment he needs to make to win set S , i.e.:*

$$\tau_i(S, a_{-i}) = \inf_{a'_i \in \mathcal{A}_i: X_i(a'_i, a_{-i}) \supseteq S} P_i(a'_i, a_{-i}) \quad (18)$$

The threshold function is additive if:

$$\tau_i(S, a_{-i}) = \sum_{j \in S} \theta_{ij}(a_{-i}) \quad (19)$$

for some item specific functions θ_{ij} , derived based on the auction rules.

The average threshold payment for a set S , takes the role of the average price of the set, in a repeated learning environment. Thus we can analogously define a no-envy learning algorithm for any repeated mechanism setting, where mechanisms M is repeated over time among the same players for T iterations and at each iteration each player picks an action $a_i^t \in \mathcal{A}_i$.

Definition 10 (No-Envy Learning for General Mechanisms). *An algorithm for a repeated mechanism setting is an α -approximate no-envy algorithm if for any adaptively and adversarially chosen sequence of opponent actions $a_{-i}^{1:T}$:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(a_i^t, a_{-i}^t) \right] \geq \max_{S \subseteq [m]} \left(\frac{1}{\alpha} v(S) - \frac{1}{T} \sum_{t \in T} \tau_i(S, a_{-i}^t) \right) - \epsilon(T) \quad (20)$$

where $\epsilon(T) \rightarrow 0$. It has polynomial envy rate if $\epsilon(T) = \text{poly}(T^{-1}, m, \max_S v(S))$.

Sufficient conditions on the mechanism. We now give conditions on the mechanism M , such that it admits efficient no-envy learning dynamics and such that any approximate no-envy outcome is also approximately efficient. Our conditions can be viewed as a stronger version of the smooth mechanism definition of Syrgkanis and Tardos [ST13], as well as a generalization of the value and revenue covering formulation of Hartline et al. [HHT14].

We begin by reminding the reader of the definition of a smooth mechanism [ST13] specialized to a combinatorial auction setting.

Definition 11 ([ST13]). *A mechanism is (λ, μ) -smooth if for any action profile $a \in \mathcal{A}$, there exists for each player i an action a_i^* for each player i , such that:*

$$\sum_{i \in [n]} u_i(a_i^*, a_{-i}) \geq \lambda \text{OPT} - \mu \mathcal{R}(a) \quad (21)$$

where $\mathcal{R}(a) = \sum_{i \in [n]} P_i(a)$ is the revenue of the auctioneer and OPT is the optimal welfare.

To apply our approach we will refine the smoothness definition and require a stronger “smoothness” property, albeit one that holds for almost all mechanisms that have been analyzed via the smooth mechanism framework. Our stronger smoothness version is more inline with the revenue and value covering framework of [HHT14] and can be thought of as an ex-post version of that framework. However, unlike the approach in [HHT14] our definition applies to general multi-dimensional mechanism design environments.

We will follow the terminology of [HHT14] of revenue and value covering. Our definition is stronger than the smooth mechanism definition in two ways. First it requires a deviation inequality for each individual player, rather than on aggregate across players. Moreover, it requires a smoothness inequality not only for the optimal allocation but rather we would require one for every possible allocation. In that respect it is closer to the solution-based smoothness of [LST16] and to the original definition of smooth games of [Rou09]. All of these strengthenings seem essential for our approach on designing no-envy dynamics to work.

Now we are ready to present the definitions of ex-post value and threshold covering, which are a stronger version of the smoothness definition.

Definition 12 (Ex-post λ -value covered). *A mechanism is ex-post λ -value covered if for any feasible allocation profile $x = (S_1, \dots, S_n)$, there exists for each player i an action $a_i^*(S_i) \in \mathcal{A}_i$ such that for any action profile $a \in \mathcal{A}$:*

$$u_i(a_i^*(S_i), a_{-i}) + \tau_i(S_i, a_{-i}) \geq \lambda v_i(S_i) \quad (22)$$

Definition 13 (Ex-post (μ_1, μ_2) -threshold covered). *A mechanism is ex-post (μ_1, μ_2) -threshold covered if for any action profile $a \in \mathcal{A}$ and allocation profile $x = (S_1, \dots, S_n)$:*

$$\sum_{i \in [n]} \tau_i(S_i, a_{-i}) \leq \mu_1 \mathcal{R}(a) + \mu_2 SW(a) \quad (23)$$

where $SW(a) = \sum_{i \in [n]} v_i(X_i(a))$.

It is easy to see that if a mechanism is λ -value covered and $(\mu, 0)$ -threshold covered, then it is (λ, μ) -smooth according to [ST13]. We add the extra welfare term, to enable the analysis of second-price auctions too. This term is related to the weakly (λ, μ_1, μ_2) -smooth mechanisms in [ST13].

No-envy learning and welfare. Now we are ready to give the generalizations of our main theorems for general mechanisms. First we argue that if a mechanism is (μ_1, μ_2) -threshold covered and players use no-envy learning, then the average welfare is approximately optimal. The proof of this theorem follows along very similar lines as in the proof of Theorem 7 and hence we omit the proof.

Theorem 11 (No-Envy Welfare for General Mechanisms). *If a mechanism is ex-post (μ_1, μ_2) -threshold covered and each player invokes an α -approximate no-envy algorithm with envy rate $\epsilon(T)$, then after T iterations the average welfare in the auction is at least $\frac{1}{\alpha(\max\{1, \mu_1\} + \mu_2)} \text{OPT} - n \cdot \epsilon(T)$, where OPT is the optimal welfare for the input valuation profile $v = (v_1, \dots, v_n)$.*

Efficient no-envy algorithms. Next we argue that if a mechanism is λ -value covered, then the existence of an efficient no-envy learning algorithm reduces to the existence of an efficient no-regret algorithm for the natural generalization of the online buyer's problem.

Definition 14 (Online buyer's problem for general mechanisms). *A buyer with some valuation $v(\cdot)$ over a set of m items wants to decide on each day which items to buy. At each time-step t an adversary picks an opponent action profile a_{-i}^t adaptively based on the past actions of the buyer. Without observing a_{-i}^t at step t , the buyer picks a set S^t to buy. His reward is:*

$$u(S^t, a_{-i}^t) = v(S^t) - \tau_i(S^t, a_{-i}^t), \quad (24)$$

i.e., the buyer receives the set S^t and pays the threshold price for the set.

Lemma 12 (From buyer to bidder in general mechanisms). *Suppose that the mechanism is α -value covered and that we are given access to an efficient learning algorithm for the online buyer's problem which guarantees for any adaptive adversary:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(S^t, a_{-i}^t) \right] \geq \max_{S \subseteq [m]} \left(\frac{1}{\alpha} v(S) - \frac{1}{T} \sum_{t=1}^T \tau_i(S^t, a_{-i}^t) \right) - \epsilon(T) \quad (25)$$

Then we can construct an efficient $\frac{\alpha}{\lambda}$ -approximate no-envy algorithm for the online bidding problem, assuming access to XOS value oracles.

The proof of the latter Lemma follows along very similar lines as in the proof of Lemma 6, hence we omit its proof.

Last it is easy to see that when the threshold functions are additive, then the online buyer's problem for general mechanisms is exactly the same as the online buyer's problem for the simultaneous second price auction mechanism. Thus our results in the main sections of the paper, provide an efficient algorithm for the online buyer's problem with $\alpha = 1$ for XOS valuations assuming access to a demand and XOS oracle and with $\alpha = (1 - 1/e)$ for coverage valuations assuming access to a value oracle.

Theorem 13. *Consider a bidder with an XOS valuation $v(\cdot)$ participating in λ -value covered mechanism with additive threshold functions. Assuming access to a demand and an XOS oracle for $v(\cdot)$, there exists a polynomial-time algorithm for computing the bidder's action a_i^t at every time step t such that after T iterations the bidder's average utility satisfies:*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T u(a^t) \right] \geq \max_S \left(\lambda \cdot v(S) - \frac{1}{T} \sum_{t=1}^T \tau_i(S^t, a_{-i}^t) \right) - O \left(\frac{m^2(D+H)}{\sqrt{T}} \right), \quad (26)$$

where D is an upper bound on the threshold function $\theta_{ij}(\cdot)$ for any item and H is an upper bound on $\max_S v(S)$. The guarantee holds with no assumption about the behavior of competing bidders.

Theorem 14. *Consider a bidder with an explicitly given coverage valuation $v(\cdot)$ participating in λ -value covered mechanism with additive threshold functions. There exists a polynomial-time algorithm for computing the bidder's action a_i^t at every time step t such that after T iterations the bidder's utility satisfies:*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T u(a^t) \right] \geq \max_S \left(\lambda \cdot \left(1 - \frac{1}{e} \right) \cdot v(S) - \frac{1}{T} \sum_{t=1}^T \tau_i(S^t, a_{-i}^t) \right) - 3m \frac{H + \sqrt{D}}{\sqrt{T}}, \quad (27)$$

H and D are as in Theorem 13. There is no assumption about the behavior of the competing bidders,

Main result for general mechanisms. Combining the aforementioned discussion and analysis we can draw the following main conclusion of this section.

Corollary 15. *When each bidder $i \in \{1, \dots, n\}$ participating in a sequence of ex-post λ -value covered and (μ_1, μ_2) -threshold covered mechanisms with additive threshold functions, has an XOS valuation (endowed with a demand and XOS oracle) or an explicitly given coverage valuation $v_i(\cdot)$, there exists a polynomial-time computable learning algorithm such that, if each bidder i employs this algorithm to compute his action a_i^t at each step t , then after T rounds the average welfare is guaranteed to be at least:*

$$\frac{1}{T} \mathbb{E} \left[\sum_{t=1}^T SW(a^t) \right] \geq \frac{\lambda}{\max\{1, \mu_1\} + \mu_2} \left(1 - \frac{1}{e} \right) \text{OPT}(v_1, \dots, v_n) - O \left(m^2 \cdot n \cdot \max_{S,i} v_i(S) \sqrt{\frac{1}{T}} \right).$$

If all bidders have XOS valuations with demand and XOS oracles the factor in front of OPT is $\frac{\lambda}{\max\{1, \mu_1\} + \mu_2}$.

We provide below two example applications of the latter theorems:

Application: Simultaneous Second Price Auctions. Revisiting simultaneous second price auctions it is easy to see that the mechanism is 1-value covered and $(0, 1)$ -threshold covered when players actions are restricted to no-overbidding actions and valuations are XOS. The value covering follows from the fact that for any set S , if we use as action $a_i^*(S)$, the bid vector that corresponds to the additive valuation returned by the XOS oracle for set S (see proof of Lemma 6). As is shown in the proof of Lemma 6, this action satisfies that for any opponents action vector:

$$u_i(a_i^*(S), a_{-i}) + \sum_{j \in S} \theta_{ij}(a_{-i}^t) \geq v(S), \quad (28)$$

which is exactly the 1-value covering inequality. The $(0, 1)$ -threshold covering inequality follows from the fact that for any feasible allocation $x = (S_1, \dots, S_n)$, since players do not overbid:

$$\sum_{i \in [n]} \tau_i(S_i, a_{-i}) \leq \sum_{j \in [m]} \max_{i \in [n]} b_{ij} \leq \sum_{i \in [n]} \sum_{j \in X_i(a)} b_{ij} \leq \sum_{i \in [n]} v_i(X_i(a)) \leq SW(a) \quad (29)$$

Thus we can apply the general theorems of this section with $\lambda = 1$, $\mu_1 = 0$ and $\mu_2 = 1$ to recover the main theorems that we derived for SiSPAs in the previous sections.

Application: Simultaneous First Price Auctions. In a simultaneous first price auction at each item the bidder pays his own bid conditional on winning, rather than the second highest bid. Based on the proof of [ST13], that the simultaneous first price auction is $(1 - 1/e, 1)$ -smooth, it is easy to see that the mechanisms is actually $(1 - 1/e)$ -value covered and $(1, 0)$ -threshold covered. Thus we can apply the latter theorems with $\lambda = 1 - 1/e$, $\mu_1 = 1$ and $\mu_2 = 0$.

Application: Simultaneous All-Pay Auctions. In a simultaneous all-pay auction at each item the bidder pays his bid no matter whether he wins or not. Based on the proof of [ST13], that the simultaneous first price auction is $(1/2, 1)$ -smooth, it is easy to see that the mechanisms is actually $1/2$ -value covered and $(1, 0)$ -threshold covered. Thus we can apply the latter theorems with $\lambda = 1/2$, $\mu_1 = 1$ and $\mu_2 = 0$.

Beyond additive threshold functions and combinatorial auctions. Finding efficient algorithms for the online buyer’s problem with general threshold functions is an interesting open problem that we defer to future work. Such an extension will enable an application of the approach presented in this paper to other value and threshold covered mechanisms, such as the greedy combinatorial auctions of [LB10]. Moreover, extending the algorithms for the online buyer’s problem to valuations defined on mechanism design settings that are more general than combinatorial auction settings, such as the lattice valuations defined in [ST13], will also enable the generalization of our approach to compositions of more general mechanisms, such as position auctions. Both of these generalization seem fruitful future directions. Our approach here shows that the problem of efficient learning algorithms that retain welfare guarantees reduces to finding efficient learning algorithms for the online buyer’s problem.

8 Further Related Work

Closer to our intractability results is the work of Cai and Papadimitriou [CP14], who show intractability of computing Bayesian-Nash equilibrium, as well as certain notions of Bayesian no-regret learning, in SiSPAs. In the Bayesian model each player’s valuation is not fixed, but drawn from some distribution independently. They show that both computing best responses and a Bayesian-Nash equilibrium in such a setting are PP-hard. They also show that Bayesian coarse correlated equilibria are NP-hard, and hence a certain type of Bayesian no-regret learning (namely when bidders re-sample their type in every round) is intractable. There are two important differences of their hardness results compared to ours:

- First, the hardness of best response in their setting is driven by the fact that the opponent bids implicitly define a distribution of exponential support. In contrast, our inapproximability of best response is shown for an explicitly given opponent bid distribution.
- Theirs is a setting where Bayesian coarse correlated equilibria are already hard, implying in particular that no-regret learning (with resampling of types in every round) is intractable. In contrast, in our setting [CKS08] has provided a centralized polynomial-time algorithm for computing a pure Nash equilibrium in complete information SiSPAs with submodular bidders. Moreover, for some special cases of combinatorial auctions with submodular bidders, [DFK15], show that computing an equilibrium with good welfare is as easy as the algorithmic problem, ignoring incentives. The centralized nature of the algorithms in these papers and the complete information assumption make this result inherently different from the setting that we want to analyze, which is the *agnostic* setting where players don’t know anything about the game and behave in a decentralized manner. In particular, in our setting, the intractability comes from the distributed nature of the computation and the incomplete, non-Bayesian, information that the bidders have.

There is a large body of work on price of anarchy in auctions, in the incomplete information Bayesian/non-Bayesian setting and under no-regret learning behavior. We cannot do justice to the vast literature but here are some example papers: [CKS08, BR11, HKMN11, FFGL13, MT12, dKMST13, LB10]. The price of anarchy of no-regret learning outcomes was first analyzed by [BHLR08] in the context of routing games and was generalized to many games in [Rou09] and to many mechanisms in [ST13], via the notion of smoothness. There is a strong connection between the smoothness framework and no-envy dynamics. In particular, the no-envy guarantee directly implies the lower bounds on the bidder’s utility, which needed for the smoothness proof to go through. This is the main reason why no-envy implies price of anarchy guarantees.

Another major stream of work in algorithmic mechanism design addresses the design of computationally efficient dominant strategy truthful mechanisms [Dob11, DV12, DL13, DRY11, DV15, Dob16]. For instance, [Dob11] shows that with only value queries, no distribution over deterministic truthful mechanisms can achieve better than polynomial approximations for submodular bidders. With demand queries [DL13] shows that no truthful in expectation mechanism can achieve better than $1 - 1/2e$ -approximation. For coverage valuations [DRY11] gives a $1 - 1/e$ -approximation, truthful in expectation randomized mechanism. For submodular bidders with demand queries the best truthful mechanism was recently given by [Dob16] achieving $O(\sqrt{\log(m)})$ -approximation. In contrast, our result shows that for no-envy XOS bidders with demand oracles, simultaneous item auctions achieve constant factor approximations.

Moreover, several papers address only the algorithmic problem of welfare maximization in combinatorial auctions with complement-free valuations. For instance, [Fei06] provides a 2-approximation for combinatorial auctions with sub-additive bidders and a $(1 - 1/e)$ -approximation for XOS bidders, with access to demand oracles, improving upon prior work of [DS06] which also required XOS oracles. Our work can also be viewed as providing a simple and distributed algorithm for welfare maximization with XOS bidders, with a $(1 - 1/e)$ -approximation guarantee: simply run our no-envy algorithms in a simultaneous first price auction game and then pick the best solution after a sufficient number of iterations.

There is a large body of work on online learning and online convex optimization to which we cannot possibly do justice. We refer the reader to two recent surveys [BCB12, SS12]. There is also a large body of work on online linear optimization where the number of experts is exponentially large, but the utility is linear in some low dimensional space. This setting was initiated by [KV05] and spurred a long line of work. We refer the reader to the relevant section of [BCB12]. Our results on perturbed leader algorithms generalize these results beyond the linear setting and we have provided some example applications beyond SiSPAs in Sections F and F.2.

Our work is also related to the recent work of [HK15] on the power of best-response oracles in online learning. This paper gives query complexity lower bounds for the general online learning problem. In contrast, our approach defines sufficient conditions (the stability) under which best-response oracles are sufficient for efficient learning and hence optimization is equivalent to online learning. Therefore, we provide a positive counterpart to these negative results.

References

- [ACBFS95] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. Gambling in a rigged casino: The adversarial multi-armed bandit problem. In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science, FOCS '95*, pages 322–, Washington, DC, USA, 1995. IEEE Computer Society.
- [AM06] Lawrence M Ausubel and Paul Milgrom. The lovely but lonely vickrey auction. *Combinatorial auctions*, 17:22–26, 2006.
- [BBHP15] Maria-Florina Balcan, Avrim Blum, Nika Haghtalab, and Ariel D. Procaccia. Commitment without regrets: Online learning in stackelberg security games. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15*, pages 61–78, New York, NY, USA, 2015. ACM.
- [BCB12] Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and non-stochastic multi-armed bandit problems. *Machine Learning*, 5(1):1–122, 2012.

- [BHLR08] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 373–382, New York, NY, USA, 2008. ACM.
- [Bik99] Sushil Bikhchandani. Auctions of heterogeneous objects. *Games and Economic Behavior*, 26(2):193 – 220, 1999.
- [BMW16] Mark Braverman, Jieming Mao, and S. Matthew Weinberg. Interpolating between truthful and non-truthful mechanisms for combinatorial auctions. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '16, pages 1444–1457. SIAM, 2016.
- [BR11] Kshipra Bhawalkar and Tim Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '11, pages 700–709. SIAM, 2011.
- [CKS08] George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th international colloquium on Automata, Languages and Programming, Part I*, ICALP '08, pages 820–832, Berlin, Heidelberg, 2008. Springer-Verlag.
- [Cla71] EdwardH. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- [CP14] Yang Cai and Christos Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 895–910, New York, NY, USA, 2014. ACM.
- [DFK15] Shahar Dobzinski, Hu Fu, and Robert Kleinberg. On the complexity of computing an equilibrium in combinatorial auctions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 110–122. SIAM, 2015.
- [dKMST13] Bart de Keijzer, Evangelos Markakis, Guido Schfer, and Orestis Telelis. Inefficiency of standard multi-unit auctions. In HansL. Bodlaender and GiuseppeF. Italiano, editors, *Algorithms ESA 2013*, volume 8125 of *Lecture Notes in Computer Science*, pages 385–396. Springer Berlin Heidelberg, 2013.
- [DL13] Shahar Dobzinski and Renato Paes Leme. Efficiency guarantees in auctions with budgets. *CoRR*, abs/1304.7048, 2013.
- [DMSW15] Nikhil Devanur, Jamie Morgenstern, Vasilis Syrgkanis, and S. Matthew Weinberg. Simple auctions with simple strategies. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 305–322, New York, NY, USA, 2015. ACM.
- [DNS10] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, February 2010.
- [Dob11] Shahar Dobzinski. An impossibility result for truthful combinatorial auctions with submodular valuations. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 139–148, New York, NY, USA, 2011. ACM.

- [Dob16] Shahar Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. *STOC'16*, 2016.
- [DRY11] Shaddin Dughmi, Tim Roughgarden, and Qiqi Yan. From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions. In *Proceedings of the Forty-third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 149–158, New York, NY, USA, 2011. ACM.
- [DS06] Shahar Dobzinski and Michael Schapira. An improved approximation algorithm for combinatorial auctions with submodular bidders. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithm*, SODA '06, pages 1064–1073, Philadelphia, PA, USA, 2006. Society for Industrial and Applied Mathematics.
- [DV12] Shahar Dobzinski and Jan Vondrak. The computational complexity of truthfulness in combinatorial auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 405–422, New York, NY, USA, 2012. ACM.
- [DV15] Shaddin Dughmi and Jan Vondrák. Limitations of randomized mechanisms for combinatorial auctions. *Games and Economic Behavior*, 92:370 – 400, 2015.
- [Fei98] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, July 1998.
- [Fei06] Uriel Feige. On maximizing welfare when utility functions are subadditive. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 41–50, New York, NY, USA, 2006. ACM.
- [FFGL13] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. Simultaneous auctions are (almost) efficient. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 201–210, New York, NY, USA, 2013. ACM.
- [FGL15] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 123–135. SIAM, 2015.
- [FKL12] Hu Fu, Robert Kleinberg, and Ron Lavi. Conditional equilibrium outcomes via ascending price processes with applications to combinatorial auctions with item bidding. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 586–586. ACM, 2012.
- [Gro73] Theodore Groves. Incentives in teams. *Econometrica*, 41(4):pp. 617–631, 1973.
- [Haz06] Elad Hazan. *Efficient Algorithms for Online Convex Optimization and Their Applications*. PhD thesis, Princeton, NJ, USA, 2006. AAI3223851.
- [HHT14] Jason Hartline, Darrell Hoy, and Sam Taggart. Price of anarchy for auction revenue. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 693–710, New York, NY, USA, 2014. ACM.
- [HK15] Elad Hazan and Tomer Koren. The computational power of optimization in online learning. *CoRR*, abs/1504.02089, 2015.

- [HKMN11] A. Hassidim, Haim Kaplan, Yishay Mansour, and Noam Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM conference on Electronic commerce, EC '11*, pages 295–296, New York, NY, USA, 2011. ACM.
- [HP05] Marcus Hutter and Jan Poland. Adaptive online prediction by following the perturbed leader. *J. Mach. Learn. Res.*, 6:639–660, December 2005.
- [KV05] Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291 – 307, 2005. Learning Theory 2003 Learning Theory 2003.
- [LB10] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '10*, pages 537–553, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
- [LCK16] Yuqian Li, Vincent Conitzer, and Dmytro Korzhyk. Catcher-evader games. *CoRR*, abs/1602.01896, 2016.
- [LST16] Thodoris Lykouris, Vasilis Syrgkanis, and Éva Tardos. Learning and efficiency in games with dynamic population. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '16*, pages 120–129. SIAM, 2016.
- [MT12] Evangelos Markakis and Orestis Telelis. Uniform price auctions: Equilibria and efficiency. In Maria Serna, editor, *Algorithmic Game Theory*, Lecture Notes in Computer Science, pages 227–238. Springer Berlin Heidelberg, 2012.
- [Rou09] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st annual ACM symposium on Theory of computing, STOC '09*, pages 513–522, New York, NY, USA, 2009. ACM.
- [SS12] Shai Shalev-Shwartz. Online learning and online convex optimization. *Found. Trends Mach. Learn.*, 4(2):107–194, February 2012.
- [ST13] Vasilis Syrgkanis and Éva Tardos. Composable and efficient mechanisms. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 211–220, New York, NY, USA, 2013. ACM.
- [Tre01] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *Proceedings of the Thirty-third Annual ACM Symposium on Theory of Computing, STOC '01*, pages 453–461, New York, NY, USA, 2001. ACM.
- [Vic61] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [Zin03] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Machine Learning, Proceedings of the Twentieth International Conference (ICML 2003), August 21-24, 2003, Washington, DC, USA*, pages 928–936, 2003.

A Omitted Proofs from Section 3

A.1 Proof of Theorem 5

Theorem 5 *The optimal bidding problem is NP-hard to approximate to within an additive ξ even when: the threshold vectors in the support of (the explicitly described distribution) D take values in $\{1, H\}^m$, where $H = k^2 \cdot m^2$, $v = 2 \cdot k \cdot m$ and $\xi = \frac{1}{2k}$.*

Proof. We break the proof in two Lemmas. In the first we show NP-hardness of the exact problem and then we show hardness of the additive approximation problem.

Lemma 16 (Hardness of Optimal Bidding). *The optimal bidding problem is NP-hard even if the threshold vectors in the support of D take values in $\{1, H\}^m$, where $H = k^2 \cdot m^2$ and $v = 2 \cdot k \cdot m$.*

Proof. Before we move to the reduction we introduce some notation that is useful for the special case when thresholds are in $\{1, H\}^m$. First each threshold vector $p \in \{1, H\}^m$ can be uniquely represented by a set T , which corresponds to the items on which the threshold is 1. Hence, the bid distribution D can be therefore described by a collection of sets $\mathcal{T} = \{T_1, \dots, T_k\}$, such that each set T_i arises with probability $1/k$.

Moreover, observe that in the optimization problem we might as well only consider strategies where the player a bid vector in $\{0, 2\}^m$. Bidding any bid in $[0, 1]$ is equivalent to bidding 0. Bidding anything in $(1, H)$ is equivalent to bidding 2. Moreover, bidding in $[H, \infty)$ is dominated by bidding 2. The reason is that bidding H increases your probability of winning only in the cases when the threshold is $H = k^2 \cdot m^2$. But in those cases your utility is negative since $v = 2 \cdot k \cdot m$. Thus it is always optimal to remove those winning cases.

Thus any bidding strategy is also uniquely characterized by a set S , which is the set of items on which the player bids 2. If a bidder chooses a set S , then he loses all items in S only if a set T_i arises, such that $T_i \cap S = \emptyset$, since then all items in S have a threshold of H . Thus the probability that he wins some item is equal to:

$$\Pr[\text{win with } S] = 1 - \frac{|\{T_i \in \mathcal{T} : T_i \cap S = \emptyset\}|}{k} \quad (30)$$

Moreover, he pays only for the items for which he bids 2 and only when the threshold of the item is also 1. Thus his expected payment when he chooses S is $\frac{1}{k} \sum_{T_i \in \mathcal{T}} |T_i \cap S|$.

Therefore, the optimal bidding problem boils down to finding the set S that maximizes:

$$\max_S v \cdot \left(1 - \frac{|\{T_i \in \mathcal{T} : T_i \cap S = \emptyset\}|}{k}\right) - \frac{1}{k} \sum_{T_i \in \mathcal{T}} |T_i \cap S| \quad (31)$$

Equivalently it is the problem of minimizing the negative part (since $v = 2 \cdot k \cdot m$ is a constant), multiplied by k :

$$\min_S v \cdot |\{T_i \in \mathcal{T} : T_i \cap S = \emptyset\}| + \sum_{T_i \in \mathcal{T}} |T_i \cap S| \quad (32)$$

Moreover, we will only consider the case where \mathcal{T} does not contain the empty set of items $[m]$ (i.e. some item always has a threshold of 1). In this case observe that by picking the whole set of items the first part of the above objective vanishes, and the second part is at most $k \cdot m$. Thus the value of the optimal objective is at most $k \cdot m$.

Observe that if the player picks a set S , such that $\exists T_i \in \mathcal{T} : T_i \cap S = \emptyset$, then the first term is at least $v = 2 \cdot k \cdot m$. Thus any such solution S must be suboptimal. Thereby at any optimal solution

the first term in the objective vanishes to zero. Hence, the optimization problem for $v = 2 \cdot k \cdot m$, is equivalent to the problem:

$$\min_{S: \nexists T_i \in \mathcal{T}: T_i \cap S = \emptyset} \sum_{T_i \in \mathcal{T}} |T_i \cap S| \quad (33)$$

We will reduce the set cover problem on a regular hypergraph to the latter equivalent form of the optimal bidding problem, which we will refer to as the simplified bidding problem. We consider a set cover instance with k elements $\{t_1, \dots, t_k\}$ and m sets $\{C_1, \dots, C_m\}$. All sets C_j are of equal size r .

For each set C_j in the set cover problem we create an item c_j in the bidding problem. For each element t_i in the set cover problem we create a set T_i in the simplified bidding problem defined as follows:

$$T_i = \{c_j \in [m] : t_i \in C_j\} \quad (34)$$

Thus set T_i contains all the items c_j , corresponding to sets C_j that contain element t_i .

First observe that for the specific instance of the simplified bidding problem that we created we can simplify the objective, because all sets C_j have the same size r :

$$\sum_{T_i \in \mathcal{T}} |T_i \cap S| = \sum_{c_j \in S} |\{T_i : c_j \in T_i\}| = \sum_{c_j \in S} |\{t_i : t_i \in C_j\}| = \sum_{c_j \in S} r = r|S| \quad (35)$$

Thus the simplified optimal bidding problem for the instance we created boils down to finding the set S of minimum cardinality, which satisfies the condition $\{\nexists T_i \in \mathcal{T} : T_i \cap S = \emptyset\}$. We will refer to any such feasible set of items S as a winning item-set (since it guarantees that the player always wins some item independent of the threshold vector).

We will now show that there is a one-to-one correspondence between winning item-sets in the simplified bidding problem instance and set covers in the original set cover instance. From this we will then conclude that finding the minimum cardinality winning item-set problem will imply finding the minimum cardinality set cover problem and would complete the reduction.

Consider a set cover $Q \subseteq \mathcal{T}$ of the set cover instance. We claim that the set of corresponding items $S = \{c_j : C_j \in Q\}$ is a winning item-set. Suppose that there exists a set $T_i \in \mathcal{T}$ such that $T_i \cap S = \emptyset$. Now consider the element t_i . Since t_i was covered by Q , it means that there exists a set $C_j \in Q$ such that $t_i \in C_j$. By construction of T_i we know that T_i contains item c_j . Since $c_j \in S$, we get that the difference $S \cap T_i$ is non-empty, a contradiction.

Consider a winning item-set S and let Q denote the corresponding collection of sets in the set cover instance. Suppose that Q is not a set cover. Thus there exists an element t_i that is not covered. This means that all the collections $C_j \in Q$ do not contain element t_i . Since all such collections c_j are not part of item set T_i , we have that $T_i \cap S = \emptyset$. Thus S cannot be a winning item-set. ■

We now show the stronger version of the hardness result. This will be useful when using the hardness of the optimal bidding problem to imply the impossibility of efficiently computable learning algorithms with polynomial regret rates.

Lemma 17 (Hardness of Approximately Optimal Bidding). *The optimal bidding problem is NP-hard to approximate to within an additive ξ even when: the threshold vectors in the support of D take values in $\{1, H\}^m$, where $H = k^2 \cdot m^2$, $v = 2 \cdot k \cdot m$ and $\xi = \frac{1}{2k}$.*

Proof. We show that getting the value of the instance of the optimal bidding created in the proof of Theorem 5 to within an additive error $\xi = \frac{1}{2k}$, will imply a 3-approximation to the original

r -regular set cover instance for any r . It is known that it is NP-hard to approximate the optimal r -regular set cover to within a $\log(r) - O(\log \log(r))$ factor [Tre01, Fei98].

Let APX be a ξ -additive approximation to the reduced instance of the optimal bidding problem and OPT the optimal value of the optimal bidding problem, i.e.:

$$OPT \in [APX - \xi, APX + \xi] \quad (36)$$

Let OPT_c be the optimal value to the set cover instance. We already argued in the proof of Theorem 5 that:

$$OPT = v - OPT_c \frac{r}{k} \quad (37)$$

Thus we get:

$$(v - APX) - \xi \leq OPT_c \frac{r}{k} \leq (v - APX) + \xi \quad (38)$$

Now we know that the bidder in the bidding problem must always pay at least 1 to win. Thus $APX \leq v - 1 \Leftrightarrow v - APX \geq 1$. Hence, we get that if:

$$\xi \leq \frac{1}{2k} \leq \frac{v - APX}{2k} \quad (39)$$

then:

$$\left(1 - \frac{1}{2k}\right) (v - APX) \leq (v - APX) - \xi \leq OPT_c \frac{r}{k} \leq (v - APX) + \xi \leq \left(1 + \frac{1}{2k}\right) (v - APX) \quad (40)$$

Which subsequently implies:

$$\frac{2k-1}{2 \cdot r} (v - APX) \leq OPT_c \leq \frac{2k+1}{2 \cdot r} (v - APX) \quad (41)$$

Thus the value: $Q = \frac{2k-1}{2 \cdot r} (v - APX)$ satisfies that:

$$Q \leq OPT_c \leq \frac{2k+1}{2k-1} Q \leq 3 \cdot Q \quad (42)$$

i.e. it is a 3-approximation to the value of the set cover instance. ■

■

A.2 Proof of Theorem 1

Proof of Theorem 1. We have already provided a high-level sketch of the approach together with a discussion of the challenges that arise in Section 3. So we proceed directly with the technical details of the proof. The proof is via a contradiction. Towards this, we will suppose that there exists a no-regret learning algorithm for the online bidding problem that we are considering.

Consider the instances of the *optimal bidding problem* obtained in the proof of Theorem 5. In these instances, we have $v = 2 \cdot k \cdot m$ and a distribution D of threshold vectors with support k . The threshold vectors in the support of the distribution take values in $\{1, k^2 m^2\}^m$.

Now imagine that our no-regret learning algorithm faces threshold vectors that are guaranteed to lie in $\{1, k^2 m^2\}^m$. In such a setting, every bid vector that our algorithm may submit is utility equivalent to or strictly dominated by some bid vector in $\{0, v/2m\}^m = \{0, k\}^m$. Thus we can

assume that our no-regret learning algorithm only submits bid vectors from this set. (This is because, given a learning algorithm that does not satisfy this property, we can easily correct it to a better one that satisfies this property.) Hence, given the history of encountered threshold vectors until time step t , our no-regret learning algorithm will compute and submit a (potentially random) bid vector $b^t \in \{0, k\}^m$. Moreover, the no-regret learning property guarantees that, facing any sequence of threshold vectors $\{\theta^t\}_{t=1}^T$ from the set $\{1, k^2 m^2\}^m$, the bid vectors submitted by our algorithm will satisfy:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E} [u(b^t, \theta^t)] \geq \max_{b \in \{0, k\}^m} \frac{1}{T} \sum_{t=1}^T u(b, \theta^t) - \epsilon(T) \quad (43)$$

where $\epsilon(T) = \text{poly}(T^{-1}, m, v) = \text{poly}(T^{-1}, m, k)$. We will show how to use such no-regret learning guarantee to compute an approximate solution to an arbitrary instance of the optimal bidding problem from Theorem 5.

We will achieve this as follows: Get T i.i.d. samples of threshold vectors from distribution D . Denote this sequence of samples $\{\theta^t\}$. Run the no-regret learning algorithm against this sequence N independent times.

We denote by b_z^t the bid vector submitted by the algorithm at time-step t of the z -th execution. Let $V_z^T(\{\theta^t\})$ be the average realized utility of the algorithm in the z -th execution, i.e.

$$V_z^T(\{\theta^t\}) = \frac{1}{T} \sum_{t=1}^T u(b_z^t, \theta^t) \quad (44)$$

We will show that:

$$\left| \frac{1}{N} \sum_{z=1}^N V_z^T(\{\theta^t\}) - \max_b \mathbb{E}_{p \sim D} [u(b, p)] \right| \leq c(T, \delta, N, k, m) \quad (45)$$

with probability at least $1 - \delta$ for some appropriately defined function $c(T, \delta, N, k, m)$.

Upper bound. Observe that V_z^T is drawn i.i.d. across the z runs and its expectation is equal to:

$$V^T(\{\theta^t\}) = \mathbb{E}_{\{b^t\}} \left[\frac{1}{T} \sum_{t=1}^T u(b^t, \theta^t) \right] \quad (46)$$

By Chernoff-Hoeffding bounds and since the average utility is upper bounded by v :

$$\Pr \left[\left| \frac{1}{N} \sum_{z=1}^N V_z^T(\{\theta^t\}) - V^T(\{\theta^t\}) \right| \geq \frac{q}{N} \right] \leq 2e^{-\frac{2q^2}{N \cdot v^2}} \quad (47)$$

Since $v = 2km$, if we set $c(N, \delta) = k \cdot m \cdot \sqrt{\frac{2 \log(2/\delta)}{N}}$, then we know that with probability $1 - \delta$:

$$\left| \frac{1}{N} \sum_{z=1}^N V_z^T(\{\theta^t\}) - V^T(\{\theta^t\}) \right| \leq c(N, \delta) \quad (48)$$

Let $b^* = \arg \max_{b \in \{0, k\}^m} \mathbb{E}_{p \sim D} [u(b, p)]$. By the no-regret assumption we know that:

$$V^T(\{\theta^t\}) \geq \max_b \frac{1}{T} \sum_{t=1}^T u(b, \theta^t) - \epsilon(T) \geq \frac{1}{T} \sum_{t=1}^T u(b^*, \theta^t) - \epsilon(T) \quad (49)$$

Since θ^t are drawn i.i.d. from D , we know by Chernoff-Hoeffding bounds that:

$$\Pr \left[\left| \frac{1}{T} \sum_{t=1}^T u(b^*, \theta^t) - \mathbb{E}_{p \sim D} [u(b^*, p)] \right| \geq \frac{q}{T} \right] \leq 2e^{-\frac{2q^2}{T \cdot v^2}} \quad (50)$$

Thus with probability $1 - \delta$:

$$\left| \frac{1}{T} \sum_{t=1}^T u(b^*, \theta^t) - \mathbb{E}_{p \sim D} [u(b^*, p)] \right| \leq c(T, \delta) \quad (51)$$

Combining the above we get that with probability $1 - 2\delta$:

$$\begin{aligned} \frac{1}{N} \sum_{z=1}^N V_z^T(\{\theta^t\}) &\geq V^T(\{\theta^t\}) - c(T, \delta) \\ &\geq \frac{1}{T} \sum_{t=1}^T u(b^*, \theta^t) - \epsilon(T) - c(T, \delta) \\ &\geq \mathbb{E}_{p \sim D} [u(b^*, \theta^t)] - \epsilon(T) - c(T, \delta) - c(N, \delta) \end{aligned}$$

Lower bound. Consider the random variable $X_z^t = u(b_z^t, \theta^t) - \mathbb{E}_{p \sim D} [u(b_z^t, p)]$. Let \mathcal{F}_{t-1} denote the filtration of all the information observed by the algorithm up till time-step $t - 1$ in the z -th execution. This is the thresholds and bids in the past steps. Observe that the conditional expectation (over bids and thresholds) of the latter variable is:

$$\mathbb{E}[X_z^t \mid \mathcal{F}_{t-1}] = \mathbb{E}[u(b_z^t, \theta^t) \mid \mathcal{F}_{t-1}] - \mathbb{E}[\mathbb{E}_{p \sim D} [u(b_z^t, p)] \mid \mathcal{F}_{t-1}] \quad (52)$$

Since θ^t are drawn i.i.d. at each time-step and are not observed by the algorithm before deciding b^t , we have that:

$$\mathbb{E}[u(b_z^t, \theta^t) \mid \mathcal{F}_{t-1}] = \mathbb{E}[\mathbb{E}_{p \sim D} [u(b_z^t, p)] \mid \mathcal{F}_{t-1}] \quad (53)$$

Thus we have that $\mathbb{E}[X_z^t \mid \mathcal{F}_{t-1}] = 0$ and thereby $\{X_z^t\}_t$ is a bounded martingale difference sequence, with $|X_z^t| \leq v$. Hence, by Hoeffding-Azuma inequality:

$$\Pr \left[\left| \sum_{t=1}^T X_z^t \right| \geq q \right] \leq 2e^{-\frac{2q^2}{T \cdot v^2}} \quad (54)$$

The latter implies that with probability $1 - \delta$:

$$\left| \frac{1}{T} \sum_{t=1}^T u(b_z^t, \theta^t) - \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{p \sim D} [u(b_z^t, p)] \right| \leq c(T, \delta) \quad (55)$$

Hence, we also get that for each z with probability $1 - \delta$:

$$\begin{aligned} V_z^T(\{\theta^t\}) &= \frac{1}{T} \sum_{t=1}^T u(b_z^t, \theta^t) \leq \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{p \sim D} [u(b_z^t, p)] + c(T, \delta) \\ &\leq \frac{1}{T} \sum_{t=1}^T \max_b \mathbb{E}_{p \sim D} [u(b, p)] + c(T, \delta) \\ &= \max_b \mathbb{E}_{p \sim D} [u(b, p)] + c(T, \delta) \end{aligned}$$

The latter holds for all z with probability at least $1 - N \cdot \delta$. Therefore, with probability at least $1 - N\delta$, the average across the N runs will satisfy the above bound.

Concluding. Hence we can conclude that with probability at least $1 - (N + 2)\delta$:

$$- \epsilon(T) - c(N, \delta) - c(T, \delta) \leq \left(\frac{1}{N} \sum_{z=1}^N V_z^T(\{\theta^t\}) - \max_b \mathbb{E}_{p \sim D} [u(b, p)] \right) \leq c(T, \delta) \quad (56)$$

Picking $\delta = \frac{\zeta}{N+2}$ and $N = T$, we get that with probability $1 - \zeta$ we have:

$$\left| \frac{1}{N} \sum_{z=1}^N V_z^T(\{\theta^t\}) - \max_b \mathbb{E}_{p \sim D} [u(b, p)] \right| \leq \epsilon(T) + 2 \cdot k \cdot m \cdot \sqrt{\frac{2 \log(2(T+2)/\zeta)}{T}} \quad (57)$$

Thus by doing T executions against a random sequence of threshold vectors from D of length T and averaging the average utilities gives us an approximation to the value of the offline *optimal bidding* problem, $\max_b \mathbb{E}_{p \sim D} [u(b, p)]$. Specifically, if $\epsilon(T) = \text{poly}(T^{-1}, k, m)$ then by setting $T = \text{poly}(\frac{1}{\xi}, k, m)$ we can get the value of the latter problem to within an additive error of ξ . Thus we cannot possibly have $\epsilon(T) = \text{poly}(T^{-1}, k, m)$, since then $\text{poly}(k, m)$ accesses to the polynomial time no-regret algorithm would give us a $\frac{1}{2k}$ additive approximation to the optimal bidding problem in polynomial time and with high probability. ■

A.3 Intepretation of Theorem 1

Theorem 1 can be viewed as a corollary of two results, of which one is specific to SiSPAs and the other is a general claim about online learning.

- Theorem 5 is equivalent to saying that it is NP-hard to compute one step of the Follow-The-Leader (FTL) algorithm in SiSPAs, even for a unit-demand bidder with the same value for all items, and even when this value is given in unary representation. Every step of FTL needs to compute an optimal bid vector against the empirical distribution of threshold bids that the algorithm has already encountered. Theorem 5 implies that this is NP-hard.

Theorem 18 (Corollary of Theorem 5). *Computing one step of FTL in SiSPAs is NP-hard, even for a unit-demand bidder with the same value v for all items that is given in unary. In fact, it is NP-hard to even compute the expected utility resulting from the optimum bid against the distribution of past opponent bids to within an additive m/v .*

- Even though FTL is not itself a no-regret learning algorithm, and even though we only established that one step of it is intractable, this sufficed to actually show (see Section A.2) that there are no polynomial-time implementable no-regret learning algorithms in SiSPAs. Here, we comment that this is a general phenomenon, applicable to any online learning setting: namely, in any setting where one step of FTL is inapproximable, there is no polynomial-time no-regret learning algorithm either. This is summarized in the following theorem.

Theorem 19. *Consider a family \mathcal{F} of functions $f : \mathcal{X} \rightarrow [0, v]$, and suppose that b bits suffice to index each function in \mathcal{F} and element in \mathcal{X} under some encoding. Suppose also that, given an explicit description of a distribution \mathcal{D} over \mathcal{F}^5 as well as v in unary representation, it is NP-hard to find some $x \in \mathcal{X}$ whose expected value $\mathbb{E}_{f \sim \mathcal{D}} [f(x)]$ is within an additive $O(1/|D|^c)$ of the optimum for some constant $c > 0$, where $|D|$ is the size of D 's support. Then, unless $\text{RP} \supseteq \text{NP}$, there is no learning algorithm running in time polynomial in b , v , and T and whose regret after T steps is any polynomial in b , v , and $1/T$.*

⁵The size of the description is the total number of bits needed to index the functions in its support and the probabilities assigned to them.

Observe that the hypothesis of Theorem 19 is tantamount to the problem facing FTL in a learning environment with cost functions \mathcal{F} , which need not be linear or convex. Our statement is also careful to cater to pseudo-polynomial dependence of the running time and regret bound on the diameter v of the range of our cost functions, as the typical dependence of no-regret learning algorithms on the diameter is pseudo-polynomial, and we are seeking to get lower bounds for such learners.

Proof sketch of Theorem 19. We simply observe that in the proof of Theorem 1 in Section A.2 we did not use the structure of the actual learning problem other than the fact that it is NP-hard to find an approximately optimal response to an explicitly given distribution, which guarantees an additive error that is inverse-polynomial in the input. ■

B Omitted Proofs from Section 4

Proof of Lemma 6. Let \mathcal{A} be the algorithm for the online buyer’s problem. At every iteration t , we will query algorithm \mathcal{A} for a set S^t . We will then call the XOS oracle on S^t to get an additive valuation $\ell \in \mathcal{L}$ such that $v(S^t) = \sum_{j \in S^t} a_j^\ell$. We will submit a bid $b_j^t = a_j^t \cdot 1\{j \in S^t\}$ on each item j . Then after seeing the threshold vector θ^t , we feed it to algorithm \mathcal{A} as a threshold vector for the buyer’s problem. We now show that this algorithm is an α -approximate no-envy algorithm for the *online bidding problem*.

First we argue that, based on the latter construction for any realization of the price vector θ^t :⁶

$$u(b^t, \theta^t) \geq u(S^t, \theta^t) = v(S^t) - \sum_{j \in S^t} \theta_j^t$$

Consider an arbitrary threshold vector θ^t and let $X^t = \{j : b_j^t > \theta_j^t\}$ be the subset of items of S^t that the player won. By the definition of the XOS valuation we have that, $v(X^t) \geq \sum_{j \in X^t} a_j^\ell$. Thus:

$$\begin{aligned} u(b^t, \theta^t) &= v(X^t) - \sum_{j \in X^t} \theta_j^t \geq \sum_{j \in X^t} a_j^\ell - \sum_{j \in X^t} \theta_j^t = \sum_{j \in S^t} (a_j^\ell - \theta_j^t) \cdot 1\{a_j^t > \theta_j^t\} \geq \sum_{j \in S^t} (a_j^\ell - \theta_j^t) \\ &= v(S^t) - \sum_{j \in S^t} \theta_j^t \end{aligned}$$

Thus we have that the expected reward of the online bidding algorithm we constructed satisfies that for any adaptively chosen sequence of thresholds:

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(b^t, \theta^t) \right] \geq \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(S^t, \theta^t) \right], \quad (58)$$

where S^t is the random set that the *online buyer’s algorithm* would have chosen at time-step t , for the same sequence of adaptively chosen thresholds. By the guarantee of the *online buyer’s algorithm* \mathcal{A} , from Equation (12), we can then conclude that the algorithm we created for the bidding problem is α -approximate no-envy.

⁶It is interesting to remark that the latter property has strong connection to the fact that the simultaneous second price auction is a smooth mechanism as defined in [ST13] and the proof is similar to the proof of showing that the auction is smooth.

It remains to argue that our algorithm never submits an overbidding bid. This follows from the definition of XOS valuations. If ℓ is the index of the additive valuation that we picked at a time-step t , then for any set X , we have: $v(X) \geq \sum_{j \in X} a_j^\ell \geq \sum_{j \in X} b_j^t$. ■

Proof of Theorem 7. Let $X^* = (X_1^*, \dots, X_n^*)$ be the optimal allocation of items to bidders. Moreover, for each item j let $B_j = \max_{i \in [n]} b_{ij}$. By the no-envy property of each player we get:

$$\mathbb{E} \left[\frac{1}{T} \sum_{i=1}^T u_i(b^t) \right] \geq \frac{1}{\alpha} v_i(X_i^*) - \mathbb{E} \left[\sum_{j \in X_i^*} \hat{\theta}_{ij} \right] - \epsilon(T)$$

Summing over all players we get:

$$\begin{aligned} \mathbb{E} \left[\frac{1}{T} \sum_{i=1}^T \sum_{i \in [n]} u_i(b^t) \right] &\geq \frac{1}{\alpha} \sum_{i \in [n]} v_i(X_i^*) - \mathbb{E} \left[\sum_{i \in [n]} \sum_{j \in X_i^*} \hat{\theta}_{ij} \right] - n \cdot \epsilon(T) \\ &\geq \frac{1}{\alpha} \text{OPT} - \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{i \in [n]} \sum_{j \in X_i^*} B_j^t \right] - n \cdot \epsilon(T) \\ &\geq \frac{1}{\alpha} \text{OPT} - \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{j \in [m]} B_j^t \right] - n \cdot \epsilon(T) \end{aligned}$$

Let X_i^t denote the random set that player i acquired at time-step t . Since the utility of the player is at most his value, we have: $v_i(X_i^t) \geq u_i(b^t)$. Moreover, since players do not overbid, we have that for any realization of the randomness: $\sum_{i \in [n]} v_i(X_i^t) \geq \sum_{j \in [m]} B_j^t$. Hence:

$$\mathbb{E} \left[\frac{1}{T} \sum_{i=1}^T \sum_{i \in [n]} v_i(X_i^t) \right] \geq \frac{1}{\alpha} \text{OPT} - \mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T \sum_{i \in [n]} v_i(X_i^t) \right] - n \cdot \epsilon(T)$$

By re-arranging we get the theorem. ■

C Omitted Proofs from Section 5

C.1 From adaptive to oblivious adversaries

We will utilize a generic reduction provided in Lemma 12 of [HP05], which states that given that in Algorithm 1 we draw independent randomization at each iteration, it suffices to provide a regret bound only for oblivious adversaries, i.e., the adversary picks a fixed sequence $\theta^{1:T}$ ahead of time without observing the actions of the player. Moreover, for any such fixed sequence of an oblivious adversary, the expected utility of the algorithm can be easily shown to be equal to the expected utility if we draw a single random sequence $\{x\}$ ahead of time and use the same random vector all the time.

The proof is as follows: by linearity of expectation and the fact that each sequence $\{x\}^t$ drawn

at each time-step t is identically distributed:

$$\begin{aligned}
\mathbb{E}_{\{x\}^1, \dots, \{x\}^t} \left[\sum_{t=1}^T u(M(\{x\}^t \cup \theta^{1:t-1}), \theta^t) \right] &= \sum_{t=1}^T \mathbb{E}_{\{x\}^t} [u(M(\{x\}^t \cup \theta^{1:t-1}), \theta^t)] \\
&= \sum_{t=1}^T \mathbb{E}_{\{x\}^1} [u(M(\{x\}^1 \cup \theta^{1:t-1}), \theta^t)] \\
&= \mathbb{E}_{\{x\}^1} \left[\sum_{t=1}^T u(M(\{x\}^1 \cup \theta^{1:t-1}), \theta^t) \right]
\end{aligned}$$

The latter is equivalent to the expected reward if we draw a single random sequence $\{x\}$ ahead of time and use the same random vector all the time. Thus it is sufficient to upper bound the regret of this modified algorithm, which draws randomness only once.

C.2 Proof of Theorem 8

Proof of Theorem 8. To prove the theorem it suffices to show that the regret of the algorithm which gets to observe θ^t ahead of time and at each time-step plays action $\tilde{a}^t = M(\{x\} \cup \theta^{1:t})$ is upper bounded by the second term in the bound. Then the theorem easily follows by observing that the expected reward of Algorithm 2 is close to the reward of this fore-sight algorithm by an error which is equal to the first term in the bound.

Thus in the remainder of the section we will analyze this algorithm with fore-sight, and bound its regret in a sequence of two Lemmas.

Lemma 20 (Be-the-leader lemma). *Suppose that we actually learned θ^t ahead of time-step t and we played according to action $\tilde{a}^t = M(\theta^{1:t})$. Then for any sequence $\theta^{1:T}$, this algorithm has no-regret against any action a^* , i.e.:*

$$\sum_{t=1}^T u(M(\theta^{1:t}), \theta^t) \geq \sum_{t=1}^T u(M(\theta^{1:T}), \theta^t) \tag{59}$$

Proof. We show it by induction. The induction hypothesis is that for any $t \in [1 : T]$:

$$\sum_{\tau=1}^t u(M(\theta^{1:\tau}), \theta^\tau) \geq U(M(\theta^{1:t}), \theta^{1:t})$$

This trivially holds for $t = 1$. Assume it holds for t and consider the case of $t + 1$:

$$\begin{aligned}
\sum_{\tau=1}^{t+1} u(M(\theta^{1:\tau}), \theta^\tau) &= \sum_{\tau=1}^t u(M(\theta^{1:\tau}), \theta^\tau) + u(M(\theta^{1:t+1}), \theta^{t+1}) \\
&\geq U(M(\theta^{1:t}), \theta^{1:t}) + u(M(\theta^{1:t+1}), \theta^{t+1}) \\
&\geq U(M(\theta^{1:t+1}), \theta^{1:t}) + u(M(\theta^{1:t+1}), \theta^{t+1}) \\
&= U(M(\theta^{1:t+1}), \theta^{1:t+1})
\end{aligned}$$

which concludes the induction. ■

Lemma 21 (Be-the-leader with fixed sample perturbations). *Suppose that we actually learned θ^t ahead of time-step t and we played according to action $\tilde{a}^t = M(\{x\} \cup \theta^{1:t})$, i.e. including θ^t in the frequency vector for some fixed sequence $\{x\} = \{x^1, \dots, x^k\}$ of k parameters. Then, this algorithm achieves regret against any action a^* :*

$$\sum_{t=1}^T (u(a^*, \theta^t) - u(\tilde{a}^t, \theta^t)) \leq \sum_{\tau=1}^k (f_+(x^\tau) + f_-(x^\tau)) \quad (60)$$

Proof. Denote with k the length of sequence $\{x\}$. Consider the sequence $\{x\} \cup \theta^{1:T}$ and apply the *be-the-leader* Lemma 20 to this sequence. We get that for any action a^* :

$$\sum_{\tau=1}^k u(M(x^{1:\tau}), x^\tau) + \sum_{t=1}^T u(M(\{x\} \cup \theta^{1:t}), \theta^t) \geq \sum_{\tau=1}^k u(a^*, x^\tau) + \sum_{t=1}^T u(a^*, \theta^t)$$

Since for all $a \in A$: $-f_-(\theta) \leq u(a, \theta) \leq f_+(\theta)$ for any θ , and since $\tilde{a}^t = M(\{x\} \cup \theta^{1:t})$, by re-arranging the latter inequality we get:

$$\sum_{t=1}^T (u(a^*, \theta^t) - u(\tilde{a}^t, \theta^t)) \leq \sum_{\tau=1}^k (u(M(x^{1:\tau}), x^\tau) - u(a^*, x^\tau)) \leq \sum_{\tau=1}^k (f_+(x^\tau) + f_-(x^\tau))$$

■
■

D Omitted Proofs from Section 5.1

D.1 Proof of Theorem 9

Stability bound. The most important part of the analysis is providing the stability bound $g(t)$, for each time-step t , given the perturbation we used. We will show it in the following lemma.

Lemma 22 (Stability lemma). *For any sequence $\theta^{1:T}$ such that $0 \leq \theta_j^t \leq D$ for all $j \in [m]$, $t \in [T]$, the stability of Algorithm 2 when applied to the online buyer's problem and with the single-sample exponential perturbation is upper bounded by:*

$$\mathbb{E}_x [u(M(x \cup \theta^{1:t}), \theta^t) - u(M(x \cup \theta^{1:t-1}), \theta^t)] \leq (mD + H) \left(\frac{m}{t} + 3\epsilon mD \right) \quad (61)$$

Proof. We consider a specific sequence $\theta^{1:T}$ and a specific time-step t . First observe that it suffices to prove the lemma for the translated utilities $u'(a, \theta) = u(a, \theta) + mD$, since such a translation preserves utility differences for any two actions and does not alter the optimizing actions in hindsight. These translated utilities lie in $[0, mD + H]$, when $\theta_j \leq D$. Thus we will prove the lemma for utilities $u(a, \theta)$ that lie in $[0, mD + H]$. For succinctness we will denote with:

$$\begin{aligned} \text{FTPL}^t &= \mathbb{E}_x [u(M(x \cup \theta^{1:t-1}), \theta^t)] \\ \text{BTPL}^t &= \mathbb{E}_x [u(M(x \cup \theta^{1:t}), \theta^t)] \end{aligned}$$

the expected reward of *follow the perturbed leader* and *be the perturbed leader*, correspondingly, at time-step t .

We will construct a mapping $\mu : \Theta \rightarrow \Theta$, such that for any random sample x , we will have that:

$$M(x \cup \theta^{1:t}) = M(\mu(x) \cup \theta^{1:t-1}) \quad (62)$$

Observe that the two maximization problems are the same if we have that the average threshold vector ends up being the same, i.e.,

$$\frac{\sum_{\tau=1}^t \theta^\tau + x}{t+1} = \frac{\sum_{\tau=1}^{t-1} \theta^\tau + \mu(x)}{t} \quad (63)$$

By re-arranging we get:

$$\mu(x) = \frac{t}{t+1}x - \frac{1}{t+1} \sum_{\tau=1}^{t-1} \theta^\tau + \frac{t}{t+1} \theta^t \quad (64)$$

Observe that $\mu(x)$ is a bijection from $\Theta^+ \triangleq \{x \in \Theta : \mu(x) \geq 0\}$ to Θ (where by $\mu(x) \geq 0$ we mean a coordinate-wise comparison). Thus we can write:

$$\begin{aligned} \text{FTPL}^t &= \int_{x \in \Theta} u(M(x \cup \theta^{1:t-1}), \theta^t) f(x) dx \\ &= \int_{y \in \Theta^+} u(M(\mu(y) \cup \theta^{1:t-1}), \theta^t) f(\mu(y)) |det(\nabla \mu)(y)| dy \\ &= \int_{y \in \Theta^+} u(M(y \cup \theta^{1:t}), \theta^t) f(\mu(y)) \left(\frac{t}{t+1}\right)^m dy \\ &= \int_{x \in \Theta^+} u(M(x \cup \theta^{1:t}), \theta^t) f(\mu(x)) \left(\frac{t}{t+1}\right)^m dx \end{aligned}$$

Now observe that for any $x \in \Theta^+$:

$$\begin{aligned} f(\mu(x)) &= \exp\{-\epsilon(\|\mu(x)\|_1 - \|x\|_1)\} f(x) \\ &= \exp\left\{-\epsilon\left(\left\|\frac{t}{t+1}x - \frac{1}{t+1} \sum_{\tau=1}^{t-1} \theta^\tau + \frac{t}{t+1} \theta^t\right\|_1 - \|x\|_1\right)\right\} f(x) \\ &\geq \exp\left\{-\epsilon\left(\frac{t}{t+1}\|x\|_1 + \frac{1}{t+1} \left\|\sum_{\tau=1}^{t-1} \theta^\tau\right\|_1 + \frac{t}{t+1} \|\theta^t\|_1 - \|x\|_1\right)\right\} f(x) \\ &\geq \exp\{-\epsilon 2mD\} f(x) \\ &\geq (1 - 2\epsilon mD) f(x) \end{aligned}$$

Moreover, it is easy to check that $\left(\frac{t}{t+1}\right)^m \geq 1 - \frac{m}{t}$ and that for any two non-negative numbers x, y , $(1-x)(1-y) \geq 1-x-y$. Last, we remind that $0 \leq u(a, \theta^t) \leq mD + H$ for any $a \in A$. Plugging the above lower bounds in the integral we get:

$$\begin{aligned} \text{FTPL}^t &\geq \int_{x \in \Theta^+} u(M(x \cup \theta^{1:t}), \theta^t) (1 - 2\epsilon mD) \left(1 - \frac{m}{t}\right) f(x) dx \\ &\geq \int_{x \in \Theta^+} u(M(x \cup \theta^{1:t}), \theta^t) \left(1 - 2\epsilon mD - \frac{m}{t}\right) f(x) dx \\ &\geq \int_{x \in \Theta} u(M(x \cup \theta^{1:t}), \theta^t) \left(1 - 2\epsilon mD - \frac{m}{t}\right) f(x) dx - (mD + H) \cdot \Pr[x \notin \Theta^+] \\ &\geq \int_{x \in \Theta} u(M(x \cup \theta^{1:t}), \theta^t) f(x) dx - (mD + H) \left(2\epsilon mD + \frac{m}{t} + \Pr[x \notin \Theta^+]\right) \\ &= \text{BTPL}^t - (mD + H) \left(2\epsilon mD + \frac{m}{t} + \Pr[x \notin \Theta^+]\right) \quad (65) \end{aligned}$$

Now, let us examine the quantity $\Pr[x \notin \Theta^+] = 1 - \Pr[\mu(x) \geq 0]$. The inequality $\mu(x) \geq 0$ is equivalent to:

$$tx - \sum_{\tau=1}^{t-1} \theta^\tau + t\theta^t \geq 0 \Leftrightarrow x \geq \frac{1}{t} \sum_{\tau=1}^{t-1} \theta^\tau - \theta^t$$

Observe that: $D \geq \frac{1}{t} \sum_{\tau=1}^{t-1} \theta^\tau - \theta^t$, thereby, the condition $\{x \geq D\}$ coordinate-wise, implies the condition $\{\mu(x) \geq 0\}$. Thus:

$$\Pr[\mu(x) \geq 0] \geq \Pr[x \geq D] = \prod_{j=1}^m \Pr[x_j \geq D] = e^{-\epsilon D m} \geq 1 - \epsilon D m.$$

Thus we get:

$$\Pr[x \notin \Theta^+] = 1 - \Pr[\mu(x) \geq 0] \leq \epsilon D m \tag{66}$$

Plugging the bound from Equation (66) into Equation (65) we get:

$$\text{FTPL}^t \geq \text{BTPL}^t - (mD + H) \left(3\epsilon m D + \frac{m}{t} \right)$$

This concludes the proof of the stability property. ■

Now we can apply Theorem 8 to get a concrete bound for the follow-the-perturbed leader algorithm with the single sample exponential perturbation.

Proof of Theorem 9. By applying Theorem 8 and Lemma 22, we get that the regret is upper bounded by:

$$\begin{aligned} \sum_{t=1}^T g(t) + \mathbb{E}_x [f_+(x) + f_-(x)] &\leq \sum_{t=1}^T (mD + H) \left(\frac{m}{t} + 3\epsilon m D \right) + \mathbb{E}_x [H + \|x\|_1] \\ &\leq (mD + H) \sum_{t=1}^T \frac{m}{t} + 3(mD + H)\epsilon m D T + H + \frac{m}{\epsilon} \\ &\leq (mD + H)m(\log(T) + 1) + 3(mD + H)\epsilon m D T + H + \frac{m}{\epsilon} \end{aligned}$$

Picking $\epsilon = \sqrt{\frac{1}{(mD+H)DT}}$, we get:

$$\sum_{t=1}^T g(t) + \mathbb{E}_x [f_-(x) + f_+(x)] \leq (mD + H)m(\log(T) + 1) + H + 4m\sqrt{(mD + H)DT}$$
■

E Omitted Proofs from Section 6

E.1 Proof of Lemma 10

From the Online Buyer's Problem to Online Convex Optimization. Suppose that the buyer picks a set at each iteration at random from a distribution where each item j is included

independently with probability x_j to the set. Any such distribution is characterized by its vector of marginals $x \in [0, 1]^m$. For succinctness we will be denoting with $\mathcal{P} = [0, 1]^m$. The expected utility of the buyer at time-step t , from picking a set S^t from such a distribution with marginals x^t can be written as:

$$\mathbb{E}_{S^t \sim x^t} [u(S^t, \theta^t)] = \mathbb{E}_{S^t \sim x^t} [v(S^t)] - \langle \theta^t, x^t \rangle, \quad (67)$$

where we denote with $\langle x, y \rangle$ the inner product between vectors x and y . The function $V(x) \equiv \mathbb{E}_{S \sim x} [v(S)]$ is what is called the *multi-linear extension of set function* v . Thus we can write the expected reward of the buyer from distribution x^t as:

$$\mathbb{E}_{S^t \sim x^t} [u(S^t, \theta^t)] = V(x^t) - \langle \theta^t, x^t \rangle. \quad (68)$$

When $V : \mathcal{P} \rightarrow \mathbb{R}_+$ is a concave function, we can solve the online problem that the buyer faces by invoking some online convex optimization algorithm such as the online gradient descent with projections of [Zin03]. All we need in order to invoke such an algorithm in this case is that we can compute the gradient of $V(x)$. If both conditions were satisfied, it would be easy to get a polynomial-time learning algorithm for computing the distribution of sets x^t at each time-step so as to achieve polynomial regret rates for the online buyer's problem.

However, the multi-linear extension is not a concave function for most natural classes of valuation functions. This same issue arises when designing truthful-in-expectation mechanisms for welfare maximization, such as in [DRY11]. We will exploit the approach taken in that paper to solve the problem we face here. Instead of set-distributions that are independent across items, the idea is to use correlated set-distributions that both convexify the expected utility of the buyer, and maintain the polynomial-time computability of its gradient.

Convex Rounding. For any item-independent distribution x^t we will construct a distribution $D(x^t)$ and then draw a set S^t from $D(x^t)$. The distribution that we will construct will have the following properties:

Definition 15 (Convex rounding scheme.). *A mapping $D : [0, 1]^m \rightarrow \Delta(2^m)$ is an α -approximate convex rounding scheme if:*

1. *The function $F(x) = \mathbb{E}_{S \sim D(x)} [v(S)]$ is a concave function of x .*
2. *The gradient of F is polynomial-time computable.*
3. *If we denote by $y_j(x)$ the probability that item j is included in S when S is drawn from distribution $D(x)$, then $y_j(x) \leq x_j$.*
4. *For any integral x , which corresponds to a set S : $F(x) \geq \frac{1}{\alpha} v(S)$.*

If we can find such an α -approximate convex rounding scheme, then we show that we can construct an α -approximate no-envy algorithm for the online buyer's problem.

Lemma 23. *If the valuation function $v(\cdot)$ admits an α -approximate convex rounding scheme with $\sup_x \|\nabla F(x)\|_2 \leq H$, then there exists a polynomial-time computable learning algorithm for the online buyer's problem which guarantees:*

$$\mathbb{E} \left[\frac{1}{T} \sum_{t=1}^T u(S^t, \theta^t) \right] \geq \max_S \left(\frac{1}{\alpha} v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - 3(H + \sqrt{mK}) \sqrt{\frac{m}{T}}. \quad (69)$$

where $K = \max_{t \in [T], j \in [m]} \theta_j^t$.

Proof. Our online buyer algorithm will invoke the projected gradient descent algorithm of [Zin03] (see also [Haz06] for a slightly improved analysis) applied to the following convex optimization problem: At each iteration the algorithm picks a vector $x^t \in [0, 1]^m$ and receives a reward of:

$$f^t(x^t) = F(x^t) - \langle \theta^t, x^t \rangle \quad (70)$$

At each iteration, our online buyer's algorithm calls projected gradient descent and receives a prediction x^t . Then our algorithm draws a set $S^t \sim D(x^t)$. Then it feeds back to projected gradient descent the parameter vector θ^t that it observed or equivalently the function f^t that it observed.

We first analyze how projected gradient works, so as to show that we can run the algorithm in polynomial time. Then we state the algorithms regret guarantees and show how these guarantees translate to guarantees about the online buyer's problem.

Projected gradient descent works as follows: Let $\Pi_{\mathcal{P}}(y) = \arg \min_{x \in \mathcal{P}} \|x - y\|_2$ denote the L_2 projection of y onto \mathcal{P} . At each iteration the projected gradient descent plays:

$$x^t = \Pi_{\mathcal{P}}(x^{t-1} + \eta_t \nabla f^{t-1}(x^{t-1})) = \Pi_{\mathcal{P}}(x^{t-1} + \eta_t \nabla F(x^{t-1}) - \eta_t \theta^t). \quad (71)$$

Our polytope $\mathcal{P} = [0, 1]^m$ is so simple that the L_2 projection step takes the closed form:

$$\Pi_{\mathcal{P}}^j(y) = \max\{0, \min\{1, y_j\}\} \quad (72)$$

and thereby is efficiently computable. The gradient step is also efficiently computable if we can compute the gradient of F at any point x in polynomial time.

[Zin03] provided the first bounds on projected gradient descent. We borrow the slightly improved bounds of [Haz06]. If we denote with $D = \max_{x, y \in \mathcal{P}} \|x - y\|_2$ and with $G = \sup_{x \in \mathcal{P}, t \in [T]} \|\nabla f^t(x)\|_2$, then we get that the average regret of projected gradient descent with $\eta_t = \frac{G}{D\sqrt{t}}$ is upper bounded by $3GD\sqrt{\frac{1}{T}}$, i.e. for any sequence of thresholds $\theta^{1:T}$:

$$\max_{x^* \in \mathcal{P}} \frac{1}{T} \sum_{t=1}^T (f^t(x^*) - f^t(x^t)) \leq 3GD\sqrt{\frac{1}{T}} \quad (73)$$

In our setting, we have that: $D = \max_{x, y \in [0, 1]^m} \|x - y\|_2 \leq \sqrt{m}$ and we have that:

$$\|\nabla f^t(x)\|_2 = \|\nabla F(x) - \theta^t\|_2 \leq \|\nabla F(x)\|_2 + \|\theta^t\|_2$$

Thus if we denote with $H = \sup_x \|\nabla F(x)\|_2$ and with $K = \max_{t \in [T], j \in [m]} \theta_j^t$, then we get that:

$$\max_{x^* \in \mathcal{P}} \frac{1}{T} \sum_{t=1}^T (f^t(x^*) - f^t(x^t)) \leq 3(H + \sqrt{mK})\sqrt{\frac{m}{T}} \quad (74)$$

Finally, we show how this regret guarantee of projected gradient descent maps back to a guarantee for our algorithm for the online buyer's problem. Observe that since by the definition of a convex rounding scheme, $y_j(x) \leq x_j$, we have that at each time-step t :

$$\mathbb{E}_{S^t \sim D(x^t)} [u(S^t, \theta^t)] = F(x^t) - \sum_{j \in [m]} \theta_j^t \cdot y_j(x^t) \geq F(x^t) - \sum_{j \in [m]} \theta_j^t \cdot x_j^t = F(x^t) - \langle \theta^t, x^t \rangle \quad (75)$$

Combining the latter with the regret guarantees of projected gradient descent we get that for any sequence of adaptively chosen $\theta^{1:T}$:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{S^t \sim D(x^t)} [u(S^t, \theta^t)] &\geq \frac{1}{T} \sum_{t=1}^T (F(x^t) - \langle \theta^t, x^t \rangle) \\ &\geq \max_{x \in \mathcal{P}} \left(F(x) - \left\langle \frac{1}{T} \sum_{t=1}^T \theta^t, x \right\rangle \right) - 3(H + \sqrt{mK}) \sqrt{\frac{m}{T}} \end{aligned}$$

Invoking the α -approximate property of the convex rounding scheme we then conclude:

$$\begin{aligned} \frac{1}{T} \sum_{t=1}^T \mathbb{E}_{S^t \sim D(x^t)} [u(S^t, \theta^t)] &\geq \max_{x \in \mathcal{P}} \left(F(x) - \left\langle \frac{1}{T} \sum_{t=1}^T \theta^t, x \right\rangle \right) - 3(H + \sqrt{mK}) \sqrt{\frac{m}{T}} \\ &\geq \max_S \left(\frac{1}{\alpha} v(S) - \sum_{j \in S} \hat{\theta}_j^T \right) - 3(H + \sqrt{mK}) \sqrt{\frac{m}{T}} \end{aligned}$$

■

[DRY11] gave a $(1 - \frac{1}{e})$ -approximate convex-rounding scheme for the case of coverage valuations, through the process that they call Poisson rounding.

Definition 16 (Poisson rounding). *Given an input x define $D(x)$ as follows: for each item j independently include j in the output set with probability: $1 - \exp\{-x_j\}$.*

We show that it satisfies all the properties that we need for the case of explicitly given coverage valuations. Some of these properties were not shown in [DRY11] as they were not needed, hence we provide a complete proof.

Lemma 24. *Poisson rounding is a $\frac{e}{e-1}$ -approximate convex rounding scheme for explicitly given coverage valuations. Moreover, $\sup_{x \in \mathcal{P}} \|\nabla F(x)\|_2 \leq \max_{j \in [m]} v(\{j\}) \sqrt{m}$.*

Proof. The following properties were shown by [DRY11], hence we re-direct the reader to that paper: i) the function $F(x) = \mathbb{E}_{S \sim D(x)}[v(S)]$ is concave, ii) it has a polynomially computable gradient given the hyper-graph representation of a coverage valuation as input, iii) the scheme satisfies that for any integral x associated with a set S , $F(x) \geq (1 - \frac{1}{e}) v(S)$. Last, it is easy to see that $y_j(x) = 1 - \exp\{-x_j\} \leq x_j$. Thus all four properties of a convex rounding scheme are satisfied.

It remains to upper bound the quantity $\sup_{x \in \mathcal{P}} \|\nabla F(x)\|_2$. For an explicitly given weighted hypergraph representation $\mathcal{G} = (V, E)$ of a coverage valuation we have:

$$F(x) = \sum_{v \in V} w_v \left(1 - \prod_{j \in E: v \in j} \exp\{-x_j\} \right) = \sum_{v \in V} w_v \left(1 - \exp \left\{ - \sum_{j \in E: v \in j} x_j \right\} \right) \quad (76)$$

Thus we have that:

$$\nabla_j F(x) = \sum_{v \in V: v \in j} w_v \exp \left\{ - \sum_{j' \in E: v \in j'} x_{j'} \right\} \leq \sum_{v \in V: v \in j} w_v = v(\{j\}) \quad (77)$$

Hence:

$$\|\nabla F(x)\|_2 \leq \sqrt{\sum_{j \in [m]} v(\{j\})^2} \leq \max_{j \in [m]} v(\{j\}) \cdot \sqrt{m} \quad (78)$$

■

Combining Lemma 23 and Lemma 24 we get Lemma 10.

F Oracle Learning and Finite Parameter Space

In this section we examine a special case of the general online learning problem, defined in Section 5, with the restriction that the adversary can pick among only a small set of d parameters. Then we show how this general results implies specific results for no-regret learning in the online bidding problem and no-regret learning in the security games model of [BBHP15]. Specifically, we show that in any online learning problem where the adversary has only d choices to pick from on each day, i.e. $|\Theta| = d$, then with an access to an offline optimization problem an instantiation of Algorithm 2, can achieve regret of $2H\sqrt{d/T}$, assuming utilities are bounded in $[0, H]$. This regret is independent of the number of available actions! Our result could be of independent interest and nicely complements recent impossibility results of [HK15].

Regret Analysis. We will assume that $u(a, \theta) \in [0, H]$, which is satisfied in our online bidding problem for any non-overbidding action a and when the value of the player is at most H . Observe that when the parameter space is finite, then we can map each sequence $\theta^{1:t}$ to a vector $\phi^t \in \mathbb{N}_+^d$, representing how many times each parameter $\theta \in \Theta$ has arrived in the first t time-steps. Then we will denote with ϕ^t this frequency vector at time-step t . Thus we can equivalently write the cumulative utility from a fixed action in hindsight in terms of this frequency vector as:

$$U(a, \phi^t) = \sum_{\theta \in \Theta} \phi_\theta^t \cdot u(a, \theta) \quad (79)$$

We will interchangeably use this notation and the old $U(a, \theta^{1:t})$ notation in this section. Equivalently, we will be writing for any frequency vector $\phi \in \mathbb{N}_+^d$:

$$M(\phi) = \arg \max_{a \in A} U(a, \phi) \quad (80)$$

for the optimization oracle as a function of the frequency vector, rather than the actual sequence and we will assume that we have access to such an oracle M .

We will show that, given such an optimization oracle, there exists an efficient no-regret algorithm which achieves regret $O(H\sqrt{dT})$ against adaptive adversaries. We consider the instantiation of Algorithm 2 with the following sample perturbation.

Definition 17 (Geometric sample perturbations). *We construct the random sequence $\{x\}$ as follows: we go over each possible parameter $\theta \in \Theta$. For each possible parameter vector we flip a coin with probability p of heads. If it comes up tails then we add an extra fake observation of that parameter to the sequence $\{x\}$ and continue flipping the coin adding an extra parameter at each time we see tails. If the coin ever comes up heads, then we stop adding vectors from that parameter and we move on to the next parameter.*

Typically we will be setting the parameter p of the order of $1/\sqrt{T}$, so that in expectation we add \sqrt{T} extra occurrences to each possible parameter. An equivalent way of viewing our sample perturbations is as if we are adding a random vector z to the frequency vector ϕ^t , where each coordinate z_θ of z is drawn independently from a geometric distribution with parameter p and supported on $\{0, 1, 2, \dots\}$. Under this interpretation at each time-step t the algorithm is picking action $a^t = M(\phi^{t-1} + z)$. We will interchangeably use this interpretation of the perturbations in the following proof.

Theorem 25 (Efficient oracle-based no-regret learning). *Follow the perturbed leader with geometric sample perturbations is efficiently computable, given access to an optimization oracle, and achieves regret against any adaptive adversary:*

$$\max_{a \in A} \sum_{t=1}^T \mathbb{E} [u(a, \theta^t) - u(a^t, \theta^t)] \leq 2H\sqrt{dT} \quad (81)$$

F.1 Application to No-Regret in Online Bidding

In this section we examine whether we can actually achieve the stronger no-regret property, rather than the no-envy property, in the online bidding problem against adversaries, assuming that we are given an offline bidding optimization oracle.

An easy corollary of our general theorem in the previous section is that if the adversary can pick one among d possible threshold vectors on each day, then there exists an efficient no-regret algorithm for the online bidding problem against adversaries, assuming that we have oracle access to the offline optimization problem. Observe that this problem is highly non-trivial, as without the optimization oracle, our impossibility result covers such settings, since in the set cover reduction, the number of different price vectors used was at most $r \cdot m$, which is finite and polynomial in the description of the instance. Yet, without an oracle there does not exist any polynomial time no-regret algorithm. We will show that in such settings an optimization oracle bypasses this impossibility.

Corollary 26. *Consider a bidder with any valuation $v(\cdot) \leq H$ participating in a sequence of simultaneous second-price auctions and assume that the number of different threshold vectors that can arise each day is at most d . Assuming access to a bidding oracle against explicitly given distributions, there exists a polynomial-time computable learning algorithm for computing the bidder's bid vector b^t at every time step t such that after T iterations the regret of the bidder is bounded by:*

$$\mathbb{E}_{b^{1:T}} \left[\frac{1}{T} \sum_{t=1}^T (u(b^*, \theta^t) - u(b^t, \theta^t)) \right] \leq 2H\sqrt{\frac{d}{T}}, \quad (82)$$

F.2 Application to Security Games

Our general online learning results in Section F find a nice application in the context of security games and specifically in the repeated Stackelberg model of [BBHP15]. In this model a defender repeatedly plays every day against a sequence of attackers. On each day he commits to a mixed strategy $a \in \Delta(S)$ on protecting a set of n targets, i.e. $S = 2^{[n]}$. The attacker that arrives each day is characterized by a type θ , which comes from a finite type space Θ , with $|\Theta| = k$. Based on his type, the attacker best responds to the mixed strategy a of the defender, through some arbitrary best-response function. Then the utility of the defender is some function of the best-response action of the attacker and the action that he committed. Observe that the action space of the defender is infinite, as it is the space of mixed strategies. [BBHP15] show that only a finite subset of the actions make sense and then invoke standard learning algorithms, resulting in a regret rate of $O\left(\sqrt{\frac{n^2 k \log(nk)}{T}}\right)$, assuming utilities of the defender are bounded in $[0, 1]$. Their algorithm however is not polynomially computable as it invokes standard learning algorithms over an exponential action space.

The crucial property is that the utility of the defender on each day can be viewed as only a function of the action a that he committed to and the type of the attacker θ , i.e. $u(a, \theta)$. This function is a complicated non-linear function, but our results in Section F do not require any structure from this function. Hence, by applying our Theorem 25 we get the following theorem:

Corollary 27. *The perturbed leader algorithm with geometric sample perturbations achieves regret rate of $2\sqrt{\frac{k}{T}}$, when applied to the Stackelberg security game of [BBHP15]. Moreover, it runs in polynomial time assuming an offline best-commitment oracle against an explicitly given distribution of attacker types.*

In other words, we get a much better regret bound and more importantly a bound that is independent of the number of targets n !

Moreover, to apply our algorithm we do not need a pre-processing step, as in [BBHP15], to make the action space of the defender finite. The latter is handled within the offline oracle. One way to construct an offline oracle is to utilize the pre-processing step of [BBHP15] and then go through all actions in the reduced action space. This would be a finite yet exponential running time, which is also the running time of the algorithms in [BBHP15]. Recently, Li et al. [LCK16], has shown that in its general form this online problem is NP-hard, thereby the exponential computation is insurmountable in the general form of the problem. However, our result allows for the immediate use of any better polynomial time offline algorithms for specific settings, in a black-box manner.

F.3 Omitted Proofs

Proof of Theorem 25. First by applying Theorem 8, we get that the regret of the algorithm is bounded by:

$$\mathbb{E}_{\{x\}} \left[\sum_{t=1}^T (u(a^*, \theta^t) - u(a^t, \theta^t)) \right] \leq \sum_{t=1}^T g(t) + H \cdot \mathbb{E}[\|z\|_1] = \sum_{t=1}^T g(t) + H \cdot \frac{d}{p} \quad (83)$$

where we used the fact that $u(a, \theta) \in [0, H]$ and thereby, $f_-(\theta) = 0$ and $f_+(\theta) = H$ and the fact that $\|z\|_1$ is the sum of d independent geometrically distributed random variables each with parameter p , thereby $\mathbb{E}[\|z\|_1] = \frac{d(1-p)}{p} \leq \frac{d}{p}$.

Now it suffices to bound the stability $g(t)$ of the algorithm, which is the main technical difficulty and which we will argue in the following stability lemma:

Lemma 28 (Stability lemma). *For any sequence $\theta^{1:T}$ and for any time-step t :*

$$\mathbb{E}_z [u(M(\phi^{t-1} + z), \theta^t)] \geq (1-p) \cdot \mathbb{E}_z [u(M(\phi^t + z), \theta^t)] \quad (84)$$

Hence, $\mathbb{E}_z [u(M(\phi^t + z), \theta^t) - u(M(\phi^{t-1} + z), \theta^t)] \leq p \cdot H = g(t)$.

Proof. Let $y^t \in \mathbb{N}^d$ denote a vector which has 1 on the coordinate associated with θ^t and 0 on every other coordinate. We can write:

$$\begin{aligned} \mathbb{E}_z [u(M(\phi^t + z), \theta^t)] &= \sum_{k \in \mathbb{N}_{\geq 0}^d} \Pr[z = k] \cdot u(M(\phi^t + k), \theta^t) \\ &= \sum_{k \in \mathbb{N}_{\geq 0}^d} \Pr[z = k] \cdot u(M(\phi^{t-1} + k + y^t), \theta^t) \\ &= \sum_{\tilde{k} \in \mathbb{N}_{\geq 0}^d, \tilde{k}_{\theta^t} \geq 1} \Pr[z = \tilde{k} - y^t] \cdot u(M(\phi^{t-1} + \tilde{k}), \theta^t) \end{aligned}$$

Observe that when each coordinate of z is an independent geometric distribution then for any $\tilde{k} \in \mathbb{N}_{\geq 0}^d$ such that $\tilde{k}_{\theta^t} \geq 1$:

$$\Pr[z = \tilde{k} - y^t] = p^n (1-p)^{\|\tilde{k} - y^t\|_1} = (1-p)^{-1} p^n (1-p)^{\|\tilde{k}\|_1} = (1-p)^{-1} \Pr[z = \tilde{k}] \quad (85)$$

Thus we get:

$$\begin{aligned}
\mathbb{E}_z [u(M(\phi^t + z), \theta^t)] &= (1 - p)^{-1} \sum_{\tilde{k} \in \mathbb{N}_{\geq 0}^d, \tilde{k}_{\theta^t} \geq 1} \Pr[z = \tilde{k}] \cdot u(M(\phi^{t-1} + \tilde{k}), \theta^t) \\
&\leq (1 - p)^{-1} \sum_{\tilde{k} \in \mathbb{N}_{\geq 0}^d} \Pr[z = \tilde{k}] \cdot u(M(\phi^{t-1} + \tilde{k}), \theta^t) \\
&= (1 - p)^{-1} \mathbb{E}_z [u(M(\phi^{t-1} + z), \theta^t)]
\end{aligned}$$

■

Thus we can conclude that the regret of the algorithm is bounded by:

$$\mathbb{E}_{\{x\}} \left[\sum_{t=1}^T (u(a^*, \theta^t) - u(a^t, \theta^t)) \right] \leq T \cdot H \cdot p + H \cdot \frac{d}{p} \tag{86}$$

Setting $p = \sqrt{\frac{d}{T}}$, yields regret $2H\sqrt{dT}$, which completes the proof.

■