

Smoothed Analysis of Online Non-parametric Auctions

NAVEEN DURVASULA, NIKA HAGHTALAB, and MANOLIS ZAMPETAKIS, University of California, Berkeley, USA

Online learning of revenue-optimal auctions is a fundamental problem in mechanism design without priors. Nevertheless, all the existing positive results assume that the auctioneer optimizes over a parameterized class of auctions, such as pricings and auctions with reserves. This is perhaps not surprising given that natural correlations that occur in online sequences pose a challenge to characterizing a succinct class of revenue-optimal auctions. This has left behind a significant gap in our understanding of online-learnability of general classes of non-parametric auctions.

We provide the first positive results for online learnability of a non-parametric auction class, for *smooth* adversaries and the class of *smooth* auctions. In a nutshell, an online adversary is smooth (in the style of Smoothed analysis [Spielman and Teng, 2004] in online learning [Haghtalab et al., 2021]) if the bid distribution has bounded density at every time step, and an auction is smooth if the level sets of its revenue function have small boundaries. We prove the following fundamental guarantees:

- (1) Revenue maximization in the class of smooth auctions is online-learnable, against smooth adversaries.
- (2) It is impossible to construct a no-regret algorithm even for the class of smooth auctions against worst-case adversaries.
- (3) It is impossible to construct a no-regret algorithm for the class of all incentive-compatible auctions even against smooth adversaries.

This gives a strong characterization of when and which class of non-parametric auctions are online-learnable.

To illustrate the generality of the class of smooth auctions we show that it contains the class of all monotone-revenue auctions, as well as, the class of all competition-monotone auctions. This brings up an interesting observation: while independence across bids leads to the optimal auctions being monotone, significantly weaker assumptions, compared to monotonicity of revenue, are sufficient for learnability.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design**; *Online learning theory*; *Sample complexity and generalization bounds*; *Computational pricing and auctions*.

Additional Key Words and Phrases: Learning in auctions, smoothed analysis, nonparametric auctions, online learning

ACM Reference Format:

Naveen Durvasula, Nika Haghtalab, and Manolis Zampetakis. 2023. Smoothed Analysis of Online Non-parametric Auctions. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, July 9–12, 2023, London, United Kingdom. ACM, New York, NY, USA, 36 pages. <https://doi.org/10.1145/3580507.3597787>

1 INTRODUCTION

Revenue maximization in auctions is a fundamental and well-studied field with numerous applications including internet advertising, real estate sales, and spectrum auctions. Foundational works in this space have contributed revenue-optimal auctions for independent private valuations and even surplus-extracting auctions for some correlated valuations, when the designer knows the distribution of the private values [Cr  mer and McLean, 1988, Myerson, 1981]. One of the main



This work is licensed under a Creative Commons Attribution International 4.0 License.

EC '23, July 9–12, 2023, London, United Kingdom

   2023 Association for Computing Machinery.

ACM ISBN 979-8-4007-0104-7/23/07

<https://doi.org/10.1145/3580507.3597787>

challenges for transitioning the insights afforded by these characterizations to practice is the difficulty (if not impossibility) of knowing the prior distribution and whether private valuations are even generated from fixed distributions. To address these challenges, online auction design has been embraced as a prior-free approach to mechanism design by removing the assumption that private valuations are generated from a distribution, much less, one that is known to the designer.

In the standard model of online auctions, the designer receives a sequence of bidder valuation profiles and, at every round, adaptively designs an auction to immediately and irrevocably allocate the items to the bidders. Prior works have studied *no-regret* algorithms for achieving average revenue that competes with that of the best auction from a target class, e.g., the best auction in the class of anonymous prices, second price auctions with individualized reserves, and level auctions (e.g., [Balcan and Blum, 2006, Balcan et al., 2018, Blum and Hartline, 2005, Bubeck et al., 2017, Cesa-Bianchi et al., 2014, Daskalakis and Syrgkanis, 2022, Dudik et al., 2020, Roughgarden and Wang, 2019]). Notably absent from this literature are online auctions that can compete with the best auction from an *unrestricted* or *non-parametric* class of (possibly DSIC and IR) auctions. In comparison, there is a range of techniques for offline learning of auctions that competes with the truly optimal Myerson auction [Cole and Roughgarden, 2014, Devanur et al., 2016, Guo et al., 2019, Morgenstern and Roughgarden, 2015].

This is perhaps not surprising given that natural correlations that occur in online sequences pose a challenge to characterizing a succinct class of revenue-optimal auctions. These correlations that range in their severity can take place across different bidders and across time. For example, an adaptive adversary can correlate a bidder's value with her or other bidder's valuations in the past; a stochastic sequence may generate bidder valuations from a joint distribution with arbitrary correlations; and even independently distributed bidder valuations, once realized into an empirical set, lose their independence.¹ Indeed, correlated values are barriers to learning revenue-optimal auctions even for a fixed distribution of valuations, where the only positive results require strong assumptions about the nature of the correlations as well as finiteness of support for the value distribution, large minimum probability for each valuation profile [Yang and Bei, 2021], or a small class of candidate distributions [Fu et al., 2014]. It is against the backdrop of these challenges and impossibilities rooted in correlations across time and bidders that we initiate the study of online learning for nonparametric classes of auctions. Correlations across bidders has been explored in the work of Psomas et al. [2019] but only in the context of understanding the approximation guarantees of simple vs optimal auctions.

Beyond auction design, adversarial correlations across time have been an obstacle to online learnability even for parametric classes, such as classification or regression loss for classes with finite Vapnik-Chervonenkis or Pseudo dimension. In such cases, worst-case correlations can lead to linear regret, while, a sequence of independent instances can be learned with small sub-linear regret [Ben-David et al., 2009, Blumer et al., 1989, Littlestone, 1988]. A recent line of work [Gupta and Roughgarden, 2016, Haghtalab et al., 2020, 2021, Rakhlin et al., 2011] proposed *smoothed analysis* as a way forward for bridging that gap by interpolating between a fully adaptive and stochastic adversary. In their model, at every round the adversary is forced to choose an instance (e.g., a valuation profile in our case) from a possibly different distribution with non-negligible anti-concentration property, such as one whose density is upper bounded by $1/\sigma$ times that of the uniform measure. By preserving some randomness across instances, Haghtalab et al. [2021] reduce regret minimization with a smoothed adaptive adversary to the problem of achieving low

¹The last example highlights a subtle but important distinction between online and offline auctions, as captured by the notion of regret. In online auctions, performance is benchmarked against the best revenue on the actual realized sequence (as captured by regret), while in offline auctions performance is against the best expected revenue achievable on the distribution of valuations—quantity that is more closely captured by the notion of pseudo-regret.

regret against a sequence with i.i.d. generated instances from the uniform measure. In this paper, we too consider smoothed adversaries towards establishing online learnability guarantees for non-parametric auction classes.

Online auction design is further complicated by correlations across bidder valuations. In presence of such correlations, no succinct class of auctions are known to characterize revenue-optimal auctions. Furthermore, auction classes that are known to include revenue-optimal auctions in some cases, such as the Crémer McLean condition, demonstrate infinite Pseudo-dimension as well as irregular and complex decision boundaries that render them unlearnable even if the valuation profiles were drawn from a fixed distribution. This shows that to obtain online learnability, we must go beyond our assumption on the smoothness of valuations (and the aforementioned adaptive-to-independent reduction framework). To overcome this challenge, we introduce and work with an abstract class of non-parametric auctions whose only requirement is a form of *smoothness of boundaries* of the level sets of its revenue function. Such auctions, which we name *smooth*, include natural classes of nonparametric auctions, such as monotone-revenue auctions and competition monotone auction (in which allocation of any bidder is monotonically decreasing in the valuations of the other bidders).

1.1 Our Contributions

In this paper, we provide the first online learnability guarantees for non-parametric classes of auctions. Our main result establishes that the class of *smooth* auctions is online learnable, i.e., we achieve vanishing average regret, in the presence of *smooth* adversaries.

In more detail, we consider bidder valuations that are supported on $[0, 1]$. We say that an adversary is σ -smooth if at every round she generates a valuation profile v^t from distribution μ^t whose density is $1/\sigma$ times the density of the Lebesgue measure over $[0, 1]^n$. We consider both *adaptive* and *oblivious* adversaries. An oblivious adversary chooses the sequence μ^1, \dots, μ^T non-adaptively at the beginning of the game, while an *adaptive* adversary can correlate the choice of μ^t with the realized valuation profiles v^1, \dots, v^{t-1} . We also consider the class of *smooth* auctions. These are auctions whose level sets of the revenue function have small boundaries, i.e., as captured by the Lebesgue measure of the boundary sets. We use Γ to refer to the smoothness of the auction class. Our main result is that *smoothness of the adversary* and *smoothness of the auctions* is sufficient to make a class online-learnable. Moreover, neither of these conditions is sufficient by itself.

Upper bounds. We state our regret bounds informally below. To simplify their presentation in this section, we assume that σ , Γ , and n are constants and focus on the dependence of regret on T .

THEOREM 4.1 (INFORMAL). *In the presence of oblivious smooth adversaries, the class of all single-item, (possibly non-deterministic) smooth auctions is online learnable with average regret of $O(T^{\frac{-1}{4(n+1)}})$.*

THEOREM 5.1 (INFORMAL). *In the presence of adaptive smooth adversaries, the class of all single-item, deterministic, DSIC, IR, and smooth auctions is online learnable with average regret of $O(T^{\frac{1}{8(n+1)}})$.*

As first motivated by Spielman and Teng [2004] the smoothness condition on the adversary is a tangible form of anti-concentration promise that reflects random perturbations that occur in practice. On the other hand, smoothness in the auction class may appear as an abstract construct. To elucidate this concept and demonstrate the strength of Theorems 4.1 and 5.1 we will show that several natural classes of auctions are smooth.

CLAIMS 3.10 AND 3.11. *The following two class of monotone auctions are smooth. 1) Revenue monotone auctions, whose revenue function is monotone in the valuation profile, and 2) Competition monotone auctions, where each bidder's allocation is monotonically decreasing in other bidders' bids.*

Interestingly, the Myerson auction is revenue monotone and competition monotone. Moreover, competition monotonicity, which directly considers the allocation function, holds for most parametric auction classes that have been considered in the past, such as non-anonymous reserves and level auctions. More generally, Theorems 4.1 and 5.1 imply that significantly weaker conditions, such as relaxing independence of private valuations and instead embracing just the monotonicity or smoothness properties of auctions, would make the class online learnable.

Lower bounds. Having established online learnability when both the auction class and the adversary are smooth, we next show that neither smoothness condition is sufficient by itself. In both cases, our lower bound holds for $n = 2$ bidders and the weaker oblivious adversary model

THEOREM 6.1 (INFORMAL; SMOOTH AUCTIONS, NON-SMOOTH ADVERSARIES). *For $n = 2$ and any online algorithm, there is an oblivious (but not smooth) adversary, such that the regret with respect to the class of deterministic, DSIC, IR, and competition monotone auctions is $\Omega(T)$.*

COROLLARY 6.3 (INFORMAL; NON-SMOOTH AUCTIONS, SMOOTH ADVERSARIES). *For $n = 2$ bidders and any online algorithm, there is an oblivious $O(1)$ -smooth adversary, such that the regret with respect to the class of all DSIC and IR auctions is $\Omega(T)$.*

Our latter lower bound is indeed a corollary of a stronger lower bound on the impossibility of offline learning revenue optimal auctions even when the prior is smooth, as stated in Theorem 6.2.

1.2 Technical overview

In this section, we give an overview of our technical tools.

Learning Theoretic Tools. We introduce two learning theoretic results about the possibility and the rate of uniform convergence and the size of ϵ -covers when measures and sets are both smooth. We expect that these results will be useful for learning non-parametric functions over smooth distributions more generally.

First, we show that as measured by smooth sets, the empirical distribution of instances drawn from smooth distributions converges to the density of the distribution from which they were drawn (i.e., their mixture). More formally, we introduce a notion of distance $D_\Gamma(\mu, \nu)$, which we call *smooth variation distance*², that equals the maximum gap in the probability assigned by ν and μ to any Γ -smooth set. Our Theorem 4.1 shows that for smooth distributions, the smooth variation distance between the empirical and the true distribution goes to 0, with sufficient number of observations.

Let us note that removing the word *smooth* from both the distribution and the set will simply result in a false statement! Because for a distribution with infinite support we can typically find an irregular set (such as the union of the realized instances) on which the empirical and true distributions largely disagree. Here smoothness helps in two ways. Smoothness of the sets invalidates the use of irregular sets in the smooth variation distance and smoothness of the measures allows us to approximate the event by the more structured nearby events by ensuring that not too much density can be focused in their gap. Together, they enable us to work with measures over lattice boxes where uniform convergence guarantees are known to hold.

Our second learning theoretic tool, Theorem 4.4, gives an ϵ -cover with respect to the *smooth variation distance*, for the set of all smooth distributions. Here again, lattice boxes are handy tools for discretizing a smooth measure, resulting in covers of reasonably small size.

²As opposed to the *total variation distance*

Oblivious Adversaries in Auctions. The above tools help us directly with learning against an oblivious adversary. In particular, uniform convergence of measure over smooth sets allows us to directly show that two different notions of optimality almost coincide: the *expected optimal revenue* achievable on the realized sequence and the *optimal expected revenue* for the mixture distribution. Therefore, bounding the pseudo-regret can be as effective as bounding the regret. To bound the pseudo-regret, we can run any no-regret algorithm over a finite set of auctions that includes the optimal auction for the true mixture distribution. We find this cover by first obtaining a cover for all smooth distributions using Theorem 4.4 and, for each distribution, taking a (near-)optimal auction.

Adaptive Adversaries in Auctions. Our overall approach for adaptive adversaries differs from the case of oblivious adversaries. Nevertheless, approximation of lattice boxes and uniform convergence of measure over smooth sets play a crucial role in this case too. As a first step in dealing with adaptive adversaries, we use a recent approach of Haghtalab et al. [2021] that uses coupled random variables to reduce interactions with an adaptive smooth adversary over T rounds to interactions with an i.i.d sequence over a slightly longer time frame. This shows that we can directly apply a no-regret algorithm (like Hedge [Freund and Schapire, 1997]) to a finite set, if this set met a stronger form of cover guarantee: The revenue of any smooth auction (not just one that is optimal for some measure) must be well-approximated on the cover. This requires a more elaborate and technical treatment of auctions that crucially leverages the payment rule of the deterministic DSIC and IR auctions.

2 PRELIMINARIES

Notation. We use Δ_n to denote the set of probability distributions over $[n]$. We view Δ_n as a subset of \mathbb{R}^n . Let $x \in \mathbb{R}^n$ we use x_{-i} to denote a vector in \mathbb{R}^{n-1} that has all the coordinates of x but the i th. For any two vectors $x, y \in \mathbb{R}^n$ we use $x \geq y$ to denote that for all $i \in [n]$ it holds that $x_i \geq y_i$.

In this section we present the online auctions framework that we use to present our results.

2.1 Auctions

Bidders. In this paper we consider an online environment where one new item is for sale to n bidders every time step $t \in [T]$ where T is the time horizon. At each time step t the bidder i has private valuation $v_i^t \in [0, 1]$ for the t -th item. We denote with $v^t = (v_1^t, \dots, v_n^t)$ the valuation profile at step t . At each time step we assume that the valuation profile is sampled from a distribution μ^t with support on a subset of $[0, 1]^n$. The measure μ^t may include correlations among bidder valuations. Also, as we will see, depending on the model the measure μ^t might or might not depend on the valuation profiles $v^{t'}$ for $t' < t$. We use \mathcal{D}_T to denote the probability distribution with support $[0, 1]^{n \cdot T}$ of all the valuation profiles $V^T = (v^1, \dots, v^T)$. At step $t \in [T]$ each bidder submits a bid $b_i^t \in [0, 1]$. We denote with $b^t = (b_1^t, \dots, b_n^t)$ the bid profile at step t .

Mechanisms. A mechanism M in this setting consists of two rules: the *allocation rule* $\mathbf{x} : [0, 1]^n \rightarrow \Delta_n$ that takes the bids b^t and outputs the probability $x_i(b^t)$ that each bidder i will receive the item, and the *payment rule* $\mathbf{p}(b^t)$ that takes the bids b^t and outputs the payment of bidder i . Bidder i 's utility at round t is then $u_i^t(b^t) = v_i^t \cdot x_i(b^t) - p_i(b^t)$, where $M^t = (\mathbf{x}^t, \mathbf{p}^t)$ is the mechanism used at round t . The auctioneer is restricted to satisfying the *Dominant Strategy Incentive Compatibility* (DSIC) and the *Individual Rationality* (IR) constraints for all rounds $t \in [T]$:

$$u_i^t(v_i, b_{-i}) \geq u_i(b_i, b_{-i}) \quad \text{for all } v_i, b_i \in [0, 1] \text{ and all } b_{-i} \in [0, 1]^{n-1} \quad (\text{DSIC})$$

$$u_i^t(v_i, b_{-i}) \geq 0 \quad \text{for all } v_i \in [0, 1] \text{ and all } b_{-i} \in [0, 1]^{n-1}. \quad (\text{IR})$$

We consider the setting in which the valuations and the prior distributions are unknown to the auctioneer. Observe that we also assume myopic bidders, i.e., bidders do not strategize across rounds. Therefore, since the mechanisms that we choose are DSIC and IR and the agents myopic, we will assume for the rest of the paper that the bids are equal to the values.

Revenue objective. We define $\text{REV}(M, v)$ to be the revenue that the mechanism $M = (\mathbf{x}, \mathbf{p})$ extracts when the valuations of the bidders are v , i.e., $\text{REV}(M, v) = \sum_{i=1}^n p_i(v)$. The goal of the auctioneer is to find a sequence for mechanisms $(M^t)_{t \in T}$ that maximizes the expected total revenue $\sum_{t=1}^T \mathbb{E}_{v^t \sim \mu^t} [\text{REV}(M^t, v^t)]$. Of course, without the knowledge of μ_t is impossible to solve this optimization problem precisely so we will evaluate the performance of a sequence $(M^t)_{t=1}^T$ using the standard notion of regret and online learnability.

2.2 No-Regret and Online Learnability

We start with the notions of regret and online learnability.

DEFINITION 2.1 (REGRET). *Given a measure \mathcal{D}_T over $\mathbb{R}^{n \cdot T}$, a possibly random sequence of mechanisms $(M_t)_{t \in [T]}$, and a class \mathcal{M} of mechanisms, we define the regret at time T as*

$$\text{AVGREGRET}_T(\mathcal{D}_T, (M_t)_{t \in [T]}, \mathcal{M}) := \mathbb{E} \left[\sup_{M \in \mathcal{M}} \frac{1}{T} \sum_t \text{REV}(M, v^t) - \frac{1}{T} \sum_t \text{REV}(M_t, v^t) \right],$$

where the expectation is over $V^T \sim \mathcal{D}_T$ and over the randomness in the choice and nature of $(M_t)_{t \in [T]}$.

An online algorithm is a map ALG that takes as input the observed valuation profiles v^1, \dots, v^{t-1} and outputs a mechanism $M_t = \text{ALG}(V^{t-1})$. The online algorithm ALG is usually randomized and therefore mechanism M_t is a random variable that depends on the random variables v^1, \dots, v^{t-1} . We use $\text{AVGREGRET}(\mathcal{D}_T, \text{ALG}, \mathcal{M})$ to denote the regret $\text{AVGREGRET}_T(\mathcal{D}_T, (M_t)_{t \in [T]}, \mathcal{M})$ where $M_t = \text{ALG}(V^{t-1})$.

To define the different versions of online learnability that we consider in this paper we need to specify a class of probability measures from which the adversary is allowed to choose. We first define the notion of smooth measures.

DEFINITION 2.2 (SMOOTH MEASURES). *We say that a Borel measure μ over $[0, 1]^n$ is σ -smooth for $\sigma \in (0, 1]$ if for any measurable set $E \subseteq [0, 1]^n$, $\mu(E) \leq \sigma^{-1} \cdot \lambda^n(E)$. Here, λ^n denotes the Lebesgue measure over \mathbb{R}^n . We denote by \mathcal{P}_σ the set of all such Borel measures over $[0, 1]^n$.*

Let \mathcal{P} be a family of distributions that generate sequences $(v^t)_{t \in \mathbb{N}}$, where for each $t \in \mathbb{N}$, $v^t \in [0, 1]^n$. For every $\mathcal{D} \in \mathcal{P}$ we use \mathcal{D}_T to denote the probability distribution of the first T terms of the sequence that is generated by \mathcal{D} . In particular, \mathcal{D}_T has support $[0, 1]^{n \cdot T}$. We define two important families of distributions: \mathcal{P}^O and \mathcal{P}^A . The class \mathcal{P}^O contains all the distributions \mathcal{D} that correspond to fixing distributions μ^1, \dots, μ^T and generating v^t independently from μ^t . The class \mathcal{P}^A contains all the distributions \mathcal{D} for which the distribution of v^t depends on the values of the sequence (v^1, \dots, v^{t-1}) . Similarly we define \mathcal{P}_σ^O to be the set of probability distributions over sequences where the probability distribution of every term is an independent distribution that belongs to \mathcal{P}_σ , and we define \mathcal{P}_σ^A to be the set of probability distributions over sequences where the probability distribution of every term given the values of the previous terms belongs to \mathcal{P}_σ .³

³When μ^t 's are unrestricted, the minmax regret against both is the same [Cesa-Bianchi and Lugosi, 1999]. When μ^t 's are smooth such strong results do not hold, but the recent works of Haghtalab et al. [2020, 2021] show that for parametric classes the minmax regret bounds are almost the same.

DEFINITION 2.3 (NO-REGRET). Given a family \mathcal{P} of measures \mathcal{D} over sequences $(v^t)_{t \in \mathbb{N}}$, where $v^i \in [0, 1]^n$ and a class \mathcal{M} of mechanisms, we say that the online algorithm ALG is a no-regret algorithm for the class \mathcal{M} and the class of distributions \mathcal{P} if for every $\mathcal{D} \in \mathcal{P}$ it holds that

$$\lim_{T \rightarrow \infty} \text{AVGREGET}_T(\mathcal{D}_T, \text{ALG}, \mathcal{M}) = 0.$$

If $\mathcal{P} = \mathcal{P}^O$ (resp. \mathcal{P}_σ^O) then ALG is a no-regret algorithm for the class \mathcal{M} against oblivious (resp. smooth oblivious) adversaries. If $\mathcal{P} = \mathcal{P}^A$ (resp. \mathcal{P}_σ^A) then ALG is a no-regret algorithm for the class \mathcal{M} against adaptive (resp. smooth adaptive) adversaries.

An equivalent notion to no-regret is the online learnability that allows us to understand how fast the average regret converges to 0.

DEFINITION 2.4 (ONLINE LEARNABILITY). Let \mathcal{M} be a class of mechanisms, and let ALG be an online algorithm. We say that ALG learns online the class \mathcal{M} with respect to the family \mathcal{P} of measures over random sequences if there exists a function $T_{\mathcal{P}, \mathcal{M}} : (0, 1) \rightarrow \mathbb{N}$ such that for every $T \geq T_{\mathcal{P}, \mathcal{M}}(\epsilon)$ and every $\mathcal{D} \in \mathcal{P}$ it holds that $\text{AVGREGET}_T(\mathcal{D}_T, \text{ALG}, \mathcal{M}) \leq \epsilon$. We say that the class \mathcal{M} is online learnable with respect to a family of measures \mathcal{P} if there exists an online learning algorithm ALG that online learns \mathcal{M} over \mathcal{P} . Similarly to the definition of no-regret: if $\mathcal{P} = \mathcal{P}^O$ (resp. \mathcal{P}_σ^O) then the class \mathcal{M} is online learnable against oblivious (resp. smooth oblivious) adversaries, and if $\mathcal{P} = \mathcal{P}^A$ (resp. \mathcal{P}_σ^A) then \mathcal{M} is online learnable against adaptive (resp. adaptive smooth) adversaries.

The function $T_{\mathcal{P}, \mathcal{M}}$ represents the number of rounds that ALG needs to observe in order to achieve regret less than ϵ . For this reason we call $T_{\mathcal{P}, \mathcal{M}}(\epsilon)$, the online sample complexity of A for the class \mathcal{M} over the family of measures \mathcal{P} . It is easy to see that no-regret and online learnability are equivalent notions. Traditionally the online learning results show an upper bound on AVGREGET_T that is inverse polynomial in T and hence establish the no-regret property of ALG . Using these upper bounds we also understand the rate at which the regret goes to 0. In this paper, it is notationally simpler and more intuitive to present our results in terms of the online sample complexity $T_{\mathcal{P}, \mathcal{M}}$ of our online algorithms ALG and from this to establish the no-regret property of ALG .

3 SMOOTH SETS AND SMOOTH AUCTIONS

In this section we define and prove some basic properties of smooth sets and smooth auctions. The notion of smooth measures is presented in Definition 2.2 and is standard in smoothed analysis (see e.g. [Manthey, 2020]) and has recently bridged learnability gaps in other online applications [Haghtalab et al., 2021]. Beyond smooth probability measures, we introduce the notions of smooth sets and smooth auctions which play a fundamental role in understanding the online learnability of an unrestricted class of auctions. After introducing these notions we show that some general and natural families of auctions satisfy the smoothness condition. Together with the lower bounds that we present in Section 6 this implies that online learning over the class of smooth auctions is a very challenging and very important problem for which none of the known techniques can be applied.

3.1 Smooth Sets

For the definition of smooth sets we first need to discretize the hypercube $[0, 1]^n$ into small cubelets with length $1/k$ for some $k \in \mathbb{N}$.

DEFINITION 3.1 (LATTICE). For any $u, v \in \mathbb{R}^n$ with $u \leq v$ we define the axis aligned box $B_{u,v} = \{x \in \mathbb{R}^n \mid \forall i \in [n] \ u_i \leq x_i \leq v_i\}$. In words, $B_{u,v}$ denotes the axis aligned box with minimal coordinates u and maximal coordinates v . Let also $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^n$ denote the vector of all ones.

We write $\mathcal{L}_k^n := \{B_{u,v} \mid u = \frac{1}{k}(c - \mathbf{1}), v = \frac{1}{k}c \text{ for some } c \in [k]^n\}$ to denote the unit lattice of $[0, 1]^n$ with k^n boxes.

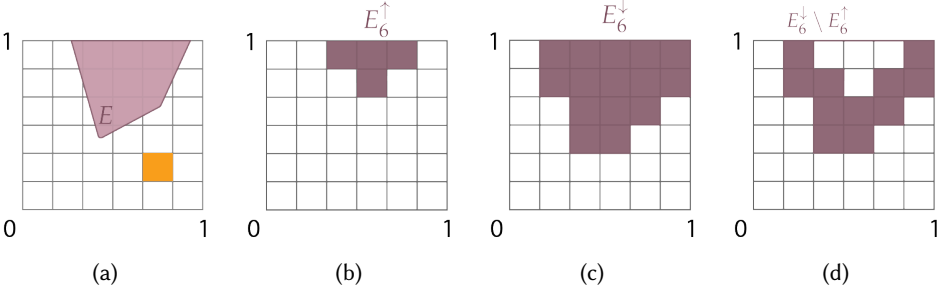


Fig. 1. (a) We illustrate the lattice \mathcal{L}_6^2 . With orange we indicate as an example the box $B_{(\frac{4}{6}, \frac{1}{6}), (\frac{5}{6}, \frac{2}{6})}$ and with purple we indicate a subset E of $[0, 1]^n$. (b) We present the inner approximation E_6^\uparrow of the set E from figure (a) according to Definition 3.2. (c) We present the outer approximation E_6^\downarrow of E according to Definition 3.2. (d) We present the set $E_6^\downarrow \setminus E_6^\uparrow$ which is used in the definition of the boundary measure $\Gamma_6(E)$.

Next, we approximate any subset $E \subseteq [0, 1]$ with subsets of the unit lattice \mathcal{L}_k^n . There are two ways of doing this: the *inner approximation* which does not include the boxes $B_{u,v} \in \mathcal{L}_k^n$ that the boundary of E crosses, and the *outer approximation* that includes the boxes of the boundary as well.

DEFINITION 3.2 (SET APPROXIMATIONS). For any subset $E \subseteq [0, 1]^n$, we define the inner and outer approximation of E by lattice boxes \mathcal{L}_k^n as

$$E_k^\uparrow := \bigcup_{B \in \mathcal{L}_k^n | B \subseteq E} B \quad \text{and} \quad E_k^\downarrow := \bigcup_{B \in \mathcal{L}_k^n | B \cap E \neq \emptyset} B.$$

Based on these we also define $\Gamma_k(E) := k \cdot \lambda^n(E_k^\downarrow \setminus E_k^\uparrow)$ to be the measure of the set boundary of E as captured by its lattice approximations⁴.

We illustrate these notions in Fig. 1.

Given the definitions of inner and outer approximations we are ready to define smooth sets. Intuitively, we call a set $E \subseteq [0, 1]^n$ *smooth* if it has a small boundary. In term of the lattice \mathcal{L}_k^n , this means that the number of boxes that the boundary of E touches should not grow as fast as k^n when k goes to ∞ , but instead should grow as k^{n-1} , which essentially means that the boundary of E is a lower dimensional surface. We make this intuition precise in the following definition.

DEFINITION 3.3 (SMOOTH SET). We say that E is Γ -smooth for some $0 < \Gamma < \infty$ if $\sup_{k \in \mathbb{N}} \Gamma_k(E) \leq \Gamma$. In other words, a set E is Γ -smooth if the boundary of E can be approximated by at most $\Gamma \cdot k^{n-1}$ boxes of the lattice \mathcal{L}_k^n by taking the gap between the corresponding inner and outer approximations. We denote by $C(\Gamma, n)$ the set of all subsets $E \subseteq [0, 1]^n$ that are Γ -smooth.⁵

An important subclass of smooth sets is the class of *monotone sets* that we define below. The fact that all monotone sets are smooth can also be used to illustrate the challenge of solving learning problems over the class of smooth sets. This is due to the fact that the VC-dimension of monotone sets is already infinite even for number of dimensions $n = 2$, as we show in Claim 3.6.

⁴The multiplication by k here is needed so that $\Gamma_k(E)$ converges to a constant as k grows to ∞ . To see why this is needed, observe that the number of cubelets in $E_k^\downarrow \setminus E_k^\uparrow$ is expected to be of order k^{n-1} because these correspond to the cubelets on the boundary of E . The Lebesgue measure of a cubelet on the other hand is $1/k^n$. Hence, $\lambda^n(E_k^\downarrow \setminus E_k^\uparrow)$ should be of order $1/k$ and this is why we multiply by a constant to get the constant that represent the surface area of E .

⁵We note that the parameterization of smooth sets and smooth measures go in opposite directions. That is, an adversary for $\sigma < 1$ becomes more smooth as σ increases and a set for $\Gamma > 1$ becomes more smooth as Γ decreases.

DEFINITION 3.4. Let $E \subseteq [0, 1]^n$, we say that E is monotone if for every $x \in E$ and every $y \in [0, 1]^n$ such that $y \geq x$, it holds that $y \in E$.

LEMMA 3.5. Let $E \subseteq [0, 1]^n$ be a monotone set, then E is n -smooth.

We present the proof of Lemma 3.5 in Appendix A.1.

CLAIM 3.6. The class of all monotone sets that are subsets of $[0, 1]^2$ has infinite VC-dimension. This implies that the set $C(\Gamma, n)$ has infinite VC-dimension even for $n = 2$.

For the proof of Claim 3.6 we refer to the Appendix A.2.

Smooth sets are a key concept in our analysis since they possess strong learning theoretic properties when combined with smooth probability measures as we will show in Section 4. For example, although smooth sets have infinite VC-dimension, we show that they satisfy uniform convergence bounds when the distribution of data is also smooth (see Section 4.1). Such general uniform convergence bounds for unrestricted distributions hold only for classes with finite VC-dimension. For this reason, we believe that these learning theoretic properties of smooth sets under smooth distributions are of independent interest.

Using the definition of smooth sets we can define the following notion of distance between probability measures that captures very well some learning properties of smooth distributions as we will see in Section 4. This distance measure resembles the total variation distance but instead of taking the supremum over all events we take the supremum over all events that correspond to smooth sets.

DEFINITION 3.7 (SMOOTH VARIATION DISTANCE). For any two probability measures μ and ν over $[0, 1]^n$, we define the Γ -smooth variation distance $D_\Gamma(\mu, \nu) := \sup_{E \in C(\Gamma, n)} |\mu(E) - \nu(E)|$. This distance yields a metric for the space of measures over $[0, 1]^n$.

3.2 Smooth Auctions

The definition of smooth sets suggests a natural definition of smooth real-valued functions $f : [0, 1]^n \rightarrow \mathbb{R}$. In particular, we can define f to be smooth when all the level sets of f of the form $\{x \mid f(x) \geq c\}$ are smooth. Applying this idea to auctions we get the following definition of *smooth auctions*.

DEFINITION 3.8 (SMOOTH AUCTION). We say that an auction M is Γ -smooth if the level sets of its revenue function $\{v \in [0, 1]^n \mid \text{REV}(M, v) \geq c\}$ are Γ -smooth for all $c \in [0, 1]$. We denote by \mathcal{M}_Γ the set of all smooth auctions.⁶ We also define \mathcal{M}_Γ^D to be the set of all Γ -smooth auctions that are also deterministic, i.e., their allocation function x takes values in $\{0, 1\}$.

The main purpose of this section is to illustrate the generality of smooth auctions and why optimizing revenue over the set \mathcal{M}_Γ is an important and very challenging problem.

Next, we give two general properties of an auction that makes it a smooth auction.

DEFINITION 3.9 (MONOTONE AUCTIONS). We consider the following types of monotone auctions:

- **Revenue Monotone Auctions.** We say that an auction M is revenue monotone if for every $z, w \in [0, 1]^n$ such that $z \geq w$, it holds that $\text{REV}(M, z) \geq \text{REV}(M, w)$. In other words, M is monotone if and only if the level sets of the revenue function are all monotone.

⁶In this paper, we define smooth sets/auctions as having small Lebesgue boundaries and correspondingly defined smooth measures relative to the Lebesgue measure. With a few technicalities, our techniques should indeed carry over to any similar notion of smoothness relative to an arbitrary base measure, so long as that base measure is the same in the definition of smooth sets and smooth measures.

- **Competition Monotone Auctions.** An auction M with allocation function $\mathbf{x} : [0, 1]^n \rightarrow \{0, 1\}^n$ is said to be competition monotone if for all $i \in [n]$ and any $\mathbf{z}_{-i}, \mathbf{w}_{-i} \in [0, 1]^{n-1}$ such that $\mathbf{z}_{-i} \geq \mathbf{w}_{-i}$, it holds that $x_i(v_i, \mathbf{z}_{-i}) \leq x_i(v_i, \mathbf{w}_{-i})$ for all $v_i \in [0, 1]$.

CLAIM 3.10. All revenue monotone auctions are n -smooth.

PROOF. Let M be a revenue monotone auction, $c \in [0, 1]$ be arbitrary, and E be the corresponding level set of the revenue function $E = \{v \in [0, 1]^n \mid \text{REV}(M, v) \geq c\}$. By the monotonicity of $\text{REV}(M, v)$, it follows that E is a monotone set. Now we can use Lemma 3.5 to conclude that E is n -smooth and hence M is also n -smooth from the definition of smoothness in auctions. ■

CLAIM 3.11. All deterministic, DSIC, and competition monotone auctions are n^2 -smooth.

We present the proof of Claim 3.11 in Appendix A.4. Moreover, we show in Appendix A.3, that the class of *randomized 2-tiered pricings* of Bergemann et al. [2020] is revenue monotone (and therefore n -smooth) in the signal space. This class includes the revenue-optimal auction for the *common valuations model* of Milgrom and Weber [1982].

To illustrate the challenge of learning over the class of Γ -smooth auctions, we show that the complexity of this class is unbounded for natural measures of complexity. In particular, the set of revenue functions of smooth auctions neither has finite pseudo-dimension nor has a useful bound on its scale-sensitive VC dimension. In fact, this is true even for competition monotone auctions that are a subclass of smooth auctions, as we showed above.

CLAIM 3.12. Let \mathcal{R}_{CM} be the set of revenue functions of competition monotone auctions, i.e.,

$$\mathcal{R}_{CM} = \{f : [0, 1]^n \rightarrow [0, 1] \mid f(x) = \text{REV}(M, x), \text{ where } M \text{ is competition monotone and DSIC}\}.$$

The class \mathcal{R}_{CM} has infinite pseudo-dimensions and infinite scale-sensitive VC-dimension for any scale $\delta < 0.5$.

The proof of Claim 3.12 is presented in Appendix A.5. We note that finiteness of the pseudo dimension and small scale-sensitive VC dimension are used for obtaining minmax offline learnability and smooth online learning guarantees for parametric and non-parametric losses [Anthony et al., 1999, Block et al., 2022, Haghtalab et al., 2022, 2021]. Therefore, Claim 3.12 shows that we need new technical innovations to learn these classes.

4 OBLIVIOUS ADVERSARIES

In this section we present our online learning of auctions results for oblivious adversaries. We start with the statement of our main theorem. Then we highlight the barriers of applying existing approaches to get our results and we summarize the main proof ideas which lead to some interesting new learning theoretic results. The full proof of our main theorem is presented in Section 4.1.

THEOREM 4.1. For any $\Gamma \in \mathbb{R}^+$ and $\sigma \in (0, 1]$, the class of all Γ -smooth auctions \mathcal{M}_Γ is online learnable with respect to the family of oblivious σ -smooth measures \mathcal{P}_σ^O with rate

$$T_{\mathcal{P}_\sigma^O, \mathcal{M}_\Gamma}(\epsilon) = \max \left(\tilde{O} \left(\frac{1}{\epsilon^{n+2}} \cdot \left(\frac{\Gamma}{\sigma} \right)^{2n+2} \right), O \left(\frac{1}{\epsilon^{4n+4}} \right) \right).$$

Further, letting Γ , σ , and n be constants, there exists an algorithm ALG such that

$$\sup_{\mathcal{D} \in \mathcal{P}_\sigma^O} \text{AVGREGRET}_T(\mathcal{D}_T, ALG, \mathcal{M}_\Gamma) \leq O \left(T^{\frac{-1}{4(n+1)}} \right).$$

Theorem 4.1 is a surprising result in many ways. First, it is very interesting that the combination of the smoothness of the distributions and the smoothness of auctions suffices to ensure online learnability. As we prove in Section 6, just one of the smoothness properties is not enough (even for $n = 2$, and Γ, σ constants) so the combination of them is necessary. The other reason that Theorem 4.1 is surprising is because the class of revenue functions of auctions that belong to \mathcal{M}_Γ does not have small complexity with respect to any of the known complexity measures of learning theory. This means that it is not even clear how to learn \mathcal{M}_Γ even in the offline setting. In fact, as we show in Section 6, without the smoothness of the prior distributions, learning over \mathcal{M}_Γ is impossible even in the offline setting.

This last observation about \mathcal{M}_Γ leads to the conclusion that to prove Theorem 4.1 we need to provide novel learning theoretic tools that can be applied in settings beyond the classical learning theory that tackle function classes with bounded complexity. We present these learning theoretic results in Sections 4.1 and 4.2. The need for these new learning theoretic tools explains also the absence of online learning results for unrestricted classes of auctions in the existing literature.

Main proof ideas. We next describe the two main steps in proving Theorem 4.1.

To prove Theorem 4.1 we want to find an online algorithm that achieves low average regret, i.e., tries to minimize the following quantity

$$\mathbb{E} \left[\max_{M \in \mathcal{M}_\Gamma} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) - \frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right], \quad (1)$$

where the expectation is over the distributions of v^t that are chosen from the class of oblivious and smooth adversaries \mathcal{P}_σ^O . Naturally, to achieve sublinear expected regret we want to apply the hedge algorithm of Freund and Schapire [1997] which achieves a regret that is sublinear in T and logarithmic in the class size. The class \mathcal{M}_Γ though is not finite and hence we cannot apply hedge directly.

A natural next idea is to try to find a cover of \mathcal{M}_Γ . Observe though that the maximization over \mathcal{M}_Γ appears inside the expectation, i.e., the maximization of revenue happens on the samples. Hence, the simplest way for the cover based approach to be effective is to find an effective cover against any probability distribution. Unfortunately, the existing techniques for constructing such covers against smooth adversaries rely on some complexity upper bound, e.g., bounded pseudo-dimension, of the class of revenue functions of \mathcal{M}_Γ which is still not true in our case as we showed in Claim 3.12. This leads us to our first main idea: *find a way to change the order of max with \mathbb{E} in (1).*

Uniform Concentration of Smooth Events. To change the order of maximum with expectation, we show the following concentration result: if T is large enough then $\frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t)$ concentrates around its expectation uniformly over all auctions $M \in \mathcal{M}_\Gamma$ (see Theorem 4.2). Now once we apply this uniform concentration we can also apply linearity of expectation and get that, with a small additive error, the average regret satisfies the following:

$$(1) \leq \max_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \mathcal{D}} [\text{REV}(M, v)] - \mathbb{E} \left[\frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right] + o(1), \quad (2)$$

where \mathcal{D} is the mixture distribution of the distributions of every step.

Cover of smooth distributions. Equation 2 implies that we are now only interested in the expected revenue achieved by the auction M . This means that we can construct a cover only of the auctions that are optimal for some distribution \mathcal{D} . Hence, it suffices to find a cover \mathcal{A} of smoothed distributions and make sure that the following holds: for every smooth distribution \mathcal{D} , there exists

a distribution $\mathcal{D}' \in \mathcal{A}$ such that $\mathbb{E}_{x \sim \mathcal{D}} [f(x)] \cong \mathbb{E}_{x \sim \mathcal{D}'} [f(x)]$ (see Theorem 4.4). If this is true, then we can define $\bar{\mathcal{M}}$ to be the set of revenue optimal auctions for the distributions inside \mathcal{A} and replace \mathcal{M}_Γ with the finite set $\bar{\mathcal{M}}$ in (2). Combining this with linearity of expectation and the fact that the expected maximum is always larger than the maximum of expectation we get that following expression for the expected regret:

$$\mathbb{E} \left[\max_{M \in \bar{\mathcal{M}}} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) - \frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right]. \quad (3)$$

In this last expression $\bar{\mathcal{M}}$ is finite and hence we can now apply hedge to prove Theorem 4.1.

The rest of this section is organized as follows: in Section 4.1 we present our uniform concentration result for smooth events over smooth distributions, in Section 4.2 we present a way to construct a cover of smooth distribution, and finally in Section 4.3 we present a full proof of Theorem 4.1.

4.1 Uniform Concentration of Smooth Events over Smooth Measures

We restate the definition of smooth variation distance $D_\Gamma(\mu, \nu)$ which is an essential notion to state our uniform concentration result. $D_\Gamma(\mu, \nu)$ is defined as follows:

$$D_\Gamma(\mu, \nu) := \sup_{E \in \mathcal{C}(\Gamma, n)} |\mu(E) - \nu(E)|. \quad (4)$$

As we can see this distance metric resembles the definition of total variation distance, except the supremum is not over all events but only over smooth events. Now the goal of our uniform concentration result is to show that for any smooth measure, the empirical distribution of T samples is close to the true distribution in terms of the smooth variation distance when T is large enough.

THEOREM 4.2. *Take v^t to be a sequence of independent random variables each with corresponding law μ^t , where μ^t is a σ -smooth probability measure over $[0, 1]^n$. Let $\hat{\mu}_T := \frac{1}{T} \sum_{t=1}^T \delta_{v^t}$ denote the sequence of empirical measures, where δ_{v^t} denotes the Dirac measure at v^t . Finally, let $\bar{\mu}_T := \frac{1}{T} \sum_{t=1}^T \mu^t$ denote the mixture measure. Then, for any $\Gamma \in \mathbb{R}^+$, $D_\Gamma(\bar{\mu}_T, \hat{\mu}_T) \xrightarrow{a.s.} 0$. Furthermore, there exist absolute constants c_1 and c_2 such that*

$$\mathbb{P} [D_\Gamma(\bar{\mu}_T, \hat{\mu}_T) > \epsilon] \leq \exp \left((c_1 \cdot \Gamma \cdot \sigma^{-1} \cdot \epsilon^{-1})^n \cdot \log(\Gamma \cdot \sigma^{-1} \cdot \epsilon^{-1}) - c_2 \cdot T \cdot \epsilon^2 \right).$$

The above theorem in one dimension resembles the celebrated DKW inequality [Dvoretzky et al., 1956]. In DKW in the definition (4), instead of the set $\mathcal{C}(\Gamma, n)$ has the set of all intervals and shows the corresponding uniform concentration bound. In higher dimensions many celebrated papers (e.g., Alon et al. [1997], Vapnik and Chervonenkis [1968]) show the same result but again instead of the set $\mathcal{C}(\Gamma, n)$ as we use in definition (4), they work with any set in a class of sets with bounded VC dimension. Vapnik and Chervonenkis [1968] in fact show that having bounded VC dimension is a necessary and sufficient condition for such a uniform concentration result to hold without any additional assumptions on the distribution. Our Theorem 4.2 shows that if we assume that the distribution of data is smooth we can show strong uniform concentration results even for classes of events that have infinite VC dimension.

PROOF OF THEOREM 4.2. In this proof, we make use of a key result from VC theory regarding empirical convergence for sets given by unions of boxes.

FACT 4.3 (UNIFORM CONVERGENCE WITH BOUNDED VC⁷). Let $\mathcal{B}(r, d)$ be the family of sets $E \subseteq \mathbb{R}^d$ such that E is a union of r axis-aligned boxes. Then, for any probability measures μ^1, \dots, μ^T over \mathbb{R}^d it holds that $\sup_{E \in \mathcal{B}(r, d)} |\hat{\mu}_T(E) - \bar{\mu}_T(E)| \xrightarrow{a.s.} 0$, where $\hat{\mu}_T$ and $\bar{\mu}_T$ are the mixture and empirical measures as described in Theorem 4.2. More precisely, for an absolute constant c ,

$$\mathbb{P} \left[\sup_{E \in \mathcal{B}(r, d)} |\hat{\mu}_T(E) - \bar{\mu}_T(E)| > \epsilon \right] \leq \exp(d \cdot r \cdot \log(dr) - c \cdot T \cdot \epsilon^2).$$

We begin by taking E to be an arbitrary Γ -smooth set. For any $k \in \mathbb{N}$, let $E_k^\uparrow \subseteq E \subseteq E_k^\downarrow$ be the inner and outer approximations of E by lattice boxes in \mathcal{L}_k^n as given in Definition 3.3. Notice that $E_k^\downarrow, E_k^\uparrow$, and their set difference all belong to the family $\mathcal{B}(k^n, n)$. At a high level, we bound the difference between $\bar{\mu}_T(E)$ and $\hat{\mu}_T(E)$ in two parts: first, as the gap between the empirical and true distribution on the interior of the set E (which is a union of boxes and therefore abides by uniform convergence laws) and their gap outside of this interior set (which by smoothness has low measure to start with). We can therefore write

$$\begin{aligned} |\bar{\mu}_T(E) - \hat{\mu}_T(E)| &\leq |\bar{\mu}_T(E \setminus E_k^\uparrow) - \hat{\mu}_T(E \setminus E_k^\uparrow)| + |\bar{\mu}_T(E_k^\uparrow) - \hat{\mu}_T(E_k^\uparrow)| \\ &\leq \max(\bar{\mu}_T(E \setminus E_k^\uparrow), \hat{\mu}_T(E \setminus E_k^\uparrow)) + |\bar{\mu}_T(E_k^\uparrow) - \hat{\mu}_T(E_k^\uparrow)| \\ &\leq \max(\bar{\mu}_T(E_k^\downarrow \setminus E_k^\uparrow), \hat{\mu}_T(E_k^\downarrow \setminus E_k^\uparrow)) + |\bar{\mu}_T(E_k^\uparrow) - \hat{\mu}_T(E_k^\uparrow)| \\ &\leq \bar{\mu}_T(E_k^\downarrow \setminus E_k^\uparrow) + |\bar{\mu}_T(E_k^\downarrow \setminus E_k^\uparrow) - \hat{\mu}_T(E_k^\downarrow \setminus E_k^\uparrow)| + |\bar{\mu}_T(E_k^\uparrow) - \hat{\mu}_T(E_k^\uparrow)| \\ &\leq \bar{\mu}_T(E_k^\downarrow \setminus E_k^\uparrow) + 2 \sup_{F \in \mathcal{B}(k^n, n)} |\bar{\mu}_T(F) - \hat{\mu}_T(F)| \\ &\leq \sigma^{-1} \cdot \lambda^n(E_k^\downarrow \setminus E_k^\uparrow) + 2 \sup_{F \in \mathcal{B}(k^n, n)} |\bar{\mu}_T(F) - \hat{\mu}_T(F)| \\ &= \sigma^{-1} \cdot k^{-1} \cdot \Gamma_k(E) + 2 \sup_{F \in \mathcal{B}(k^n, n)} |\bar{\mu}_T(F) - \hat{\mu}_T(F)| \\ &\leq \sigma^{-1} \cdot k^{-1} \cdot \Gamma + 2 \sup_{F \in \mathcal{B}(k^n, n)} |\bar{\mu}_T(F) - \hat{\mu}_T(F)|, \end{aligned}$$

where the first inequality follows from triangle inequality, the third inequality from the fact that $E \subseteq E_k^\downarrow$ as shown in Fig. 1, the fourth inequality from triangle inequality, the fifth inequality follows from the fact that the sets $E_k^\downarrow \setminus E_k^\uparrow$ and E_k^\uparrow are both unions of at most k^n axis aligned boxes, the sixth inequality follows from the smoothness of the measure $\bar{\mu}_T$ and the last two inequalities from the fact that E is Γ -smooth.

Hence, it holds that $D_\Gamma(\bar{\mu}_T, \hat{\mu}_T) \xrightarrow{a.s.} 0$, as we may make the first term arbitrarily small by taking $k \rightarrow \infty$, and we may make the second term arbitrarily small almost surely by taking $T \rightarrow \infty$ as per Fact 4.3. Furthermore if we apply Fact 4.3 with $r = k^n$, we have that for any $k > \Gamma \cdot \sigma^{-1} \cdot \epsilon^{-1}$,

$$\mathbb{P}[D_\Gamma(\bar{\mu}_T, \hat{\mu}_T) > \epsilon] \leq \exp(n^2 \cdot k^n \cdot \log(kn) - c \cdot T \cdot (\epsilon - \Gamma \cdot \sigma^{-1} \cdot k^{-1})^2)$$

Choosing $k = 2 \cdot \Gamma \cdot \sigma^{-1} \cdot \epsilon^{-1}$, we obtain the desired result. \blacksquare

⁷This bound is typically applied to i.i.d random variable, but it is folklore that it holds for independent (not necessarily identically distributed) random variable as well. Refer to Haghtalab [2018, Lemma 7.3.3] for the proof of this fact.

4.2 A Cover of Smooth Measures

In this section we show how to construct a cover of all smooth distributions with respect to the smooth variation distance D_Γ . As usual, this covering result suggests the existence of a distribution learning algorithm for smooth distributions in terms of the smooth variation distance D_Γ .

THEOREM 4.4. *For all $\Gamma, \epsilon \in \mathbb{R}^+$, there exists a finite set of measures $\mathcal{A}_\Gamma(\sigma, \epsilon)$ such that for all σ -smooth probability measures μ , $\min_{\nu \in \mathcal{A}_\Gamma(\sigma, \epsilon)} D_\Gamma(\mu, \nu) < \epsilon$. Furthermore,*

$$\log |\mathcal{A}_\Gamma(\sigma, \epsilon)| < n \cdot O(\Gamma \cdot \sigma^{-1} \cdot \epsilon^{-1})^n \cdot \log(\Gamma \cdot \sigma^{-1} \cdot \epsilon^{-2}).$$

The proof of Theorem 4.4 is deferred to Appendix B.1. We are now ready to prove Theorem 4.1.

4.3 Proof of Theorem 4.1

Our previous results relate to the learnability of smooth distributions under the D_Γ metric. To link these results to statements about auctions, we make use of the following Lemma.

LEMMA 4.5. *Let $M \in \mathcal{M}_\Gamma$ be a Γ -smooth auction. For any two Borel probability measures μ and ν over $[0, 1]^n$,*

$$|\mathbb{E}_{v \sim \mu} [\text{REV}(M, v)] - \mathbb{E}_{v \sim \nu} [\text{REV}(M, v)]| \leq D_\Gamma(\mu, \nu).$$

PROOF. Note that for any $N \in \mathbb{N}$, we may approximate any $\text{REV}(M, v)$ up to an error of $\frac{1}{N}$ using a linear combination of indicator functions of the form $\mathbb{1}_{\text{REV}(M, v) \geq c}$ for $c \in [0, 1]$. That is, $|\text{REV}(M, v) - \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\text{REV}(M, v) \geq \frac{i}{N}}| \leq \frac{1}{N}$. It follows by the triangle inequality that

$$\begin{aligned} |\mathbb{E}_{v \sim \mu} [\text{REV}(M, v)] - \mathbb{E}_{v \sim \nu} [\text{REV}(M, v)]| &\leq \frac{1}{N} \left| \sum_{i=1}^N \mathbb{E}_{v \sim \mu} \left[\mathbb{1}_{\text{REV}(M, v) \geq \frac{i}{N}} \right] - \sum_{i=1}^N \mathbb{E}_{v \sim \nu} \left[\mathbb{1}_{\text{REV}(M, v) \geq \frac{i}{N}} \right] \right| + \frac{2}{N} \\ &\leq \frac{1}{N} \sum_{i=1}^N \left| \mu \left\{ \text{REV}(M, v) \geq \frac{i}{N} \right\} - \nu \left\{ \text{REV}(M, v) \geq \frac{i}{N} \right\} \right| + \frac{2}{N} \leq D_\Gamma(\mu, \nu) + \frac{2}{N}, \end{aligned}$$

where the last transition is taken by noting that the summand is the difference between measures assigned to smooth sets. Taking $N \rightarrow \infty$, we obtain the desired result. \blacksquare

We now use Theorem 4.2 to show the average regret in our setting is close to the pseudoregret. The proof of the next corollary (Corollary 4.6) is in Appendix B.2.

COROLLARY 4.6. *For any $T \geq O\left(\frac{\Gamma}{\sigma}\right)^{2n+2}$ it holds that*

$$\mathbb{E}_{v^t \sim \mu^t} \left[\sup_{M \in \mathcal{M}_\Gamma} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) \right] \leq \sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M, v)] + O\left(T^{-\frac{1}{4(n+1)}}\right).$$

Next, we apply Theorem 4.4 to show that the pseudoregret over the nonparametric class of auction mechanisms \mathcal{M}_Γ is close to the corresponding pseudoregret over a finite net of auction mechanisms of suitable size.

COROLLARY 4.7. *For all $\epsilon > 0$, there exists a finite set of auction mechanisms $\mathcal{M}_\Gamma(\sigma, \epsilon)$ such that*

$$\sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M, v)] \leq \max_{M \in \mathcal{M}_\Gamma(\sigma, \epsilon)} \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{v^t \sim \mu^t} [\text{REV}(M, v^t)] + \epsilon,$$

where $|\log \mathcal{M}_\Gamma(\sigma, \epsilon)| \leq n \cdot O(\Gamma \cdot \sigma^{-1} \cdot \epsilon^{-1})^n \cdot \log(\Gamma \cdot \sigma^{-1} \cdot \epsilon^{-2})$.

PROOF. Let $\mathcal{A}_\Gamma(\sigma, \epsilon/2)$ be as given in Theorem 4.4. We define

$$M_\Gamma(\sigma, \epsilon) := \{M_\nu^* \mid \nu \in \mathcal{A}_\Gamma(\sigma, \epsilon),\} \quad (5)$$

where $M_\nu^* \in \mathcal{M}_\Gamma$ is any Γ -smooth auction such that $\sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \nu} [\text{REV}(M, v)] - \mathbb{E}_{v \sim \nu} [\text{REV}(M_\nu^*, v)] \leq \frac{\epsilon}{2}$. That is, $M_\Gamma(\sigma, \epsilon)$ is a collection of (near-)optimal auctions for the net distributions. By Theorem 4.4, it follows that the size of $M_\Gamma(\sigma, \epsilon)$ is as desired. Further, by Theorem 4.4, there exists some distribution $\nu^* \in \mathcal{A}_\Gamma(\sigma, \epsilon/2)$ such that $D_\Gamma(\bar{\mu}_T, \nu) \leq \epsilon/4$. We can then apply Lemma 4.5 to find that

$$\begin{aligned} \sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M, v)] &\leq \sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \nu^*} [\text{REV}(M, v)] + \sup_{M \in \mathcal{M}_\Gamma} |\mathbb{E}_{v \sim \nu^*} [\text{REV}(M, v)] - \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M, v)]| \\ &\leq \sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \nu^*} [\text{REV}(M, v)] + D_\Gamma(\bar{\mu}_T, \nu^*) \leq \mathbb{E}_{v \sim \nu^*} [\text{REV}(M_\nu^*, v)] + \frac{\epsilon}{4} + \frac{\epsilon}{2} \\ &\leq \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M_\nu^*, v)] + \frac{3\epsilon}{4} + D_\Gamma(\bar{\mu}_T, \nu^*) \leq \max_{M \in \mathcal{M}_\Gamma(\sigma, \epsilon)} \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M, v)] + \epsilon \\ &= \max_{M \in \mathcal{M}_\Gamma(\sigma, \epsilon)} \frac{1}{T} \sum_t \mathbb{E}_{v^t \sim \mu^t} [\text{REV}(M, v^t)] + \epsilon. \end{aligned}$$

■

Finally, we apply the Hedge algorithm to choose mechanisms M_t from the finite net $\mathcal{M}_\Gamma(\sigma, \epsilon)$.

FACT 4.8 (HEDGE [FREUND AND SCHAPIRE, 1997]). *The hedge algorithm can select a mechanism $M_t \in \mathcal{M}_\Gamma(\sigma, \epsilon)$ at each time iteration t such that*

$$\mathbb{E}_{v^1, \dots, v^t \sim \mathcal{D}_T} \left[\max_{M \in \mathcal{M}_\Gamma(\sigma, \epsilon)} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) - \frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right] \leq T^{-1/2} \sqrt{\log |\mathcal{M}_\Gamma(\sigma, \epsilon)|}$$

Putting everything together, we find that for any $\epsilon > 0$, $T \geq O\left(\frac{\Gamma}{\sigma}\right)^{2n}$ and $\mathcal{D}_T \in \mathcal{P}_\sigma^O$,

$$\begin{aligned} \text{AVGREGRET}(\mathcal{D}_T, (M_t), \mathcal{M}_\Gamma) &= \mathbb{E}_{v^t \sim \mu^t} \left[\sup_{M \in \mathcal{M}_\Gamma} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) - \frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right] \\ &\leq \sup_{M \in \mathcal{M}_\Gamma} \mathbb{E}_{v \sim \bar{\mu}_T} [\text{REV}(M, v)] + O\left(T^{-\frac{1}{4(n+1)}}\right) - \mathbb{E}_{v^t \sim \mu^t} \left[\frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right] \quad (\text{Cor. 4.6}) \\ &\leq \max_{M \in \mathcal{M}_\Gamma(\sigma, \epsilon)} \frac{1}{T} \sum_{t \in [T]} \mathbb{E}_{v^t \sim \mu^t} [\text{REV}(M, v^t)] - \mathbb{E}_{v^t \sim \mu^t} \left[\frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right] + O\left(T^{-\frac{1}{4(n+1)}}\right) + \epsilon \quad (\text{Cor. 4.7}) \\ &\leq T^{-1/2} \sqrt{\log |\mathcal{M}_\Gamma(\sigma, \epsilon)|} + O\left(T^{-\frac{1}{4(n+1)}}\right) + \epsilon \quad (\text{by Fact 4.8}) \\ &\leq T^{-1/2} \sqrt{n \cdot O\left(\frac{\Gamma}{\epsilon\sigma}\right)^n \cdot \log \frac{\Gamma}{\epsilon\sigma}} + O\left(T^{-\frac{1}{4(n+1)}}\right) + \epsilon \end{aligned}$$

We now set $\epsilon = T^{-\frac{1}{4(n+1)}}$ to obtain the desired result.

5 ADAPTIVE ADVERSARIES

Adaptive adversaries introduce many additional challenges and the technique that we presented in Section 4 does not apply. In this section we show how to overcome these challenges and still prove online learnability when we focus on deterministic auctions. Our main result is the following.

THEOREM 5.1. *For any $\Gamma \in \mathbb{R}^+$ and $\sigma \in (0, 1]$, the class of all Γ -smooth, deterministic, DSIC, and IR auctions \mathcal{M}_Γ^D is online learnable with respect to the family of adaptive σ -smooth sequences \mathcal{P}_σ^A with*

$$T_{\mathcal{P}_\sigma^A, \mathcal{M}_\Gamma^D}(\epsilon) = \max \left(\left(\frac{n}{\sigma} \cdot \log \frac{1}{\epsilon} \right)^{O(n)}, O \left(\frac{\Gamma^{2n+2} \cdot \log(n)^{4/3}}{\epsilon^{8/3}} \right), O \left(\frac{1}{\epsilon^{8(n+1)}} \right) \right).$$

Further, letting Γ , σ , and n be constant, there exists an algorithm ALG such that

$$\sup_{\mathcal{D} \in \mathcal{P}_\sigma^A} \text{AVGREGRET}(\mathcal{D}_T, ALG, \mathcal{M}_\Gamma^D) \leq O \left(T^{-\frac{1}{8(n+1)}} \right).$$

Main proof ideas. The main step from Section 4 that fails against adaptive adversaries is that we cannot directly apply our uniform concentration bound of Theorem 4.2 because the distributions across rounds are not independent. For this reason we follow a different approach that leads us to an expression where we can apply Theorem 4.2. Our goal is again to apply Hedge to get a sub-constant bound in the expected regret

$$\mathbb{E} \left[\max_{M \in \mathcal{M}_\Gamma^D} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) - \frac{1}{T} \sum_{t \in [T]} \text{REV}(M_t, v^t) \right]. \quad (6)$$

Now instead of first changing the order of the expectation and the maximum and then providing a cover that approximates the optimal smooth auctions, we try to find a good cover from the beginning. Let $\overline{\mathcal{M}}$ be any class of auctions, then it is not hard to see that the first term of (6) is upper bounded by

$$\mathbb{E} \left[\max_{M \in \overline{\mathcal{M}}} \frac{1}{T} \sum_{t \in [T]} \text{REV}(M, v^t) \right] + \mathbb{E} \left[\sup_{M \in \mathcal{M}_\Gamma^D} \min_{M' \in \overline{\mathcal{M}}} \left| \frac{1}{T} \sum_{t \in [T]} (\text{REV}(M, v^t) - \text{REV}(M', v^t)) \right| \right]. \quad (7)$$

Hence it suffices to find a finite set $\overline{\mathcal{M}}$ such that the following quantity, which is larger than the second term of (7), is small. Then we can apply the hedge algorithm over the set of auctions $\overline{\mathcal{M}}$ and prove our online learnability result. The second term in Equation 7 is upper bounded by:

$$\begin{aligned} & \mathbb{E} \left[\sup_{M \in \mathcal{M}_\Gamma^D} \min_{M' \in \overline{\mathcal{M}}} \frac{1}{T} \sum_{t \in [T]} |\text{REV}(M, v^t) - \text{REV}(M', v^t)| \right] \\ & \leq \mathbb{E} \left[\sup_{M \in \mathcal{M}_\Gamma^D} \min_{M' \in \overline{\mathcal{M}}} \frac{1}{T} \sum_{t \in [T]} \frac{1}{N} \sum_{i=1}^N \left| \mathbb{1} \left[\text{REV}(M, v^t) \geq \frac{i}{N} \right] - \mathbb{1} \left[\text{REV}(M', v^t) \geq \frac{i}{N} \right] \right| \right] + \frac{2}{N} \quad (8) \end{aligned}$$

The latter upper bound is given by approximating the revenue functions by sums of indicator functions in a similar manner to Lemma 4.5. This bound makes it easier to control the smoothness of the difference of indicator functions. The first step towards showing that (8) is small in the adaptive adversary case is to use the following powerful coupling lemma of Haghtalab et al. [2021].

LEMMA 5.2 ([HAGHTALAB ET AL., 2021]). *Let $v^1, \dots, v^T \sim \mathcal{D}_T$ be an adaptive σ -smooth sequence. Then, for each $z > 0$ there exists a coupling Π such that $(v^1, x_1^1, \dots, x_z^1, \dots, v^T, x_1^T, \dots, x_z^T) \sim \Pi$ and*

- (1) v^1, \dots, v^T are jointly distributed as \mathcal{D}_T
- (2) x_j^t are i.i.d uniform variables on $[0, 1]^n$
- (3) With probability at least $1 - T(1 - \sigma)^z$, $\{v^1, \dots, v^T\} \subseteq \{x_j^t \mid t \in [T], j \in [z]\}$

The third property of the coupling in Lemma 5.2 allows us to upper bound any set-monotone function in presence of an adaptive smooth adversary by switching in the larger set of i.i.d. random variables. In particular, using Lemma 5.2 in (8) we can change the expectation over the adaptive adversaries to the expectation over independent uniform distributions, by increasing the number of rounds from T to $T' \approx T \log(T)/\sigma$, and the expression (8) will only increase. Hence, it suffices to find a set $\overline{\mathcal{M}}$ such that the expression of (8) is small when the expectation is over independent uniform distributions. Now we can use a trick to make sure that the summands in (8) are smooth and apply our uniform concentration Theorem 4.2 to get that it suffices to upper bound the following

$$\sup_{M \in \mathcal{M}_\Gamma^D} \min_{M' \in \overline{\mathcal{M}}} \mathbb{E}_{v \sim U} \left[\left| \mathbb{1} \left[\text{REV}(M, v^t) \geq \frac{i}{N} \right] - \mathbb{1} \left[\text{REV}(M', v^t) \geq \frac{i}{N} \right] \right| \right]. \quad (9)$$

This suggests that $\overline{\mathcal{M}}$ has to be constructed so that it covers all the possible revenue functions of deterministic, smooth, DSIC, and IR mechanisms over the uniform distribution. This is a much stronger requirement compared to the cover that we constructed in the oblivious case. We find such a cover for which (9) is small for the subclass $\mathcal{M}_\Gamma^D \subseteq \mathcal{M}_\Gamma$ of deterministic auctions mechanisms. We give a brief description of its construction.

Whereas the cover given in Corollary 4.7 did not make use of any specific auction properties (we constructed it by finding a collection of approximately-optimal auctions given an initial net of distributions), the cover we use in this argument is constructed directly over the space of allocation functions. Further, we utilize the DSIC property to show that (9) is small.

Similar to our other arguments, we begin by considering lattices of the form \mathcal{L}_k^n . We consider the finite set of all DSIC auctions that have constant allocation function on any cubelet $B \in \mathcal{L}_k^n$. As these allocation functions themselves are constructed from a lattice, they are effectively perfectly 0-smooth. This, coupled with the Γ -smoothness of the auction, allows us to ensure that the summands in (8) are smooth. Next, for any auction $M \in \mathcal{M}_\Gamma^D$, we consider the corresponding mechanism $M^* \in \overline{\mathcal{M}}$ that allocates the item to bidder i if the bid profile v lies in the inner approximation of the set of bid profiles of the original set of bids for which M allocated the item to bidder i . By the Γ -smoothness of the auction, the boundary of this set is close to that of its inner approximation. As the price function of a DSIC auction depends only on the location of the allocation boundary, we may conclude that the revenue (and its level sets) of M^* are indeed close to M as desired. Finding such a cover completes our proof of Theorem 5.1.

The above is only a high level description of the proof and the actual proof is more involved in many ways. We refer to Appendix C for a full description of the proof of Theorem 5.1.

6 LOWER BOUNDS

In this section we show two lower bounds that justify our assumptions about the smoothness of both the distribution and the set of auctions. In particular, we show that if one of the smoothness assumptions is violated then for every online algorithm there exists an oblivious adversary for which the algorithm will have linear regret. Both of our lower bounds holds even in the case where we have $n = 2$ number of bidders and illustrate the difficulty of online learning optimal auctions even with a constant number of agents.

Our lower bound constructions share a common idea which is that if we have correlation among bidders then we can use the bid of the first bidder to infer the bid of the second bidder without violating incentive compatibility. This way the optimal auction can almost extract the social welfare as revenue. Any algorithm that compares with this optimal auction needs to infer almost exactly the correlation among the bidders which as we will show is not possible.

In Section 6.1 we show that for any algorithm ALG, there exists an instance with non-smooth distributions for which ALG has linear regret even compared to the optimal smooth auction in hindsight. In fact, we show that ALG will have linear regret even compared to the optimal competitive monotone and deterministic auction in hindsight. In Section 6.2 we show that for any algorithm ALG, there exists an instance with smooth distributions for which ALG has linear regret compared to the optimal auction in hindsight from the set of all auctions. We defer the full proofs of the results in this section to Appendices D.1 and D.2.

6.1 Smooth Auctions Non-smooth Distributions

In our lower bound we have two bidders, the first bidder has value either $1/2$ or 1 for the item whereas the second bidder has value that is always less than $1/4$ for the item. This means that the revenue optimal auction will definitely allocate the item to the first bidder because we can extract revenue at least $1/2$ from them whereas from the second bidder we can extract revenue at most $1/4$. So the first bidder always gets the item and the question is at what price. In our construction the bid of the second bidder is completely determined from the bid of the first bidder. Hence looking at the bid of the second bidder we know whether the first bidder has value $1/2$ or 1 for the item and we can charge the first bidder his whole value as a price without violating incentive compatibility or individual rationality. To make sure that this is doable with a deterministic competitive monotone auction, we use a simple threshold rule to associates the first bid with the second bid. Now to make sure that any algorithm has linear regret we pick the second bid to be always closer to the threshold in a way that all the previous information is not helpful to decide the optimal reserve price for the next time step. This resembles the impossibility result of online learning of threshold functions when the distribution of data is non-smooth.

THEOREM 6.1. *For any online learning algorithm ALG and every T there is a distribution $\mathcal{D}_T = \mathcal{D}_{1T} \times \dots \times \mathcal{D}_{TT}$ where \mathcal{D}_{iT} is a product distribution over $[0, 1]^2$, such that $\text{AVGREGRET}(\mathcal{D}_T, (M_t)_{t \in [T]}, \mathcal{M}) \geq T/4$, where $(M_t)_{t \in [T]}$ is the sequence of auctions generated by ALG and \mathcal{M} is the set of all auctions that are deterministic, competitive monotone, DSIC, and IR.*

6.2 Non-smooth Auctions Smooth Distributions

The starting point of our lower bound in this section is the same as in Section 6.1. We have two bidders, the first bidder always has higher valuation than the second bidder and from the second bidder we can infer the bid of the first bidder. Of course, since we want the distributions to be smooth we cannot use point masses and hence the second bid can only approximately imply the distribution of the first bid. The idea to construct such a smooth distribution is to define a continuous function $y = f(x)$ between the first bid y and the second bid x . Now for every x we define y to be distributed uniformly in an interval $[f(x) - q, f(x) + q]$. Now the optimal reserved price given x when y is distributed this way is some value $r(x) \in [f(x) - q, f(x) + q]$ which means that our only chance of achieving good revenue is to be able to estimate $f(x)$ with an error up to q . But since f can be any continuous function, we can show that it is impossible to find a good estimation of f with any finite number of samples and hence it is impossible to find good reserve prices. This implies in fact a lower bound even for the offline learning problem as we see in the following theorem.

THEOREM 6.2. *For every $T > 0$ and for any learning algorithm ALG that receives T i.i.d. samples from an unknown distribution and outputs a mechanism, there exists a distribution \mathcal{D} over $[0, 1]^2$ with maximum density 8 such that*

$$\mathbb{E}_{v \sim \mathcal{D}} [\text{REV}(M^*, v)] - \mathbb{E} [\mathbb{E}_{v \sim \mathcal{D}} [\text{REV}(\text{ALG}(z_1, \dots, z_T), v)]] \geq \frac{1}{192},$$

where the outside expectation of the second term is over the T samples z_1, \dots, z_T that are drawn from \mathcal{D} which are input to the learning algorithm ALG and M^* is the revenue optimal auction for \mathcal{D} .

Theorem 6.2 implies that even the offline problem of estimating the optimal auction from any finite number of samples is impossible if we only assume that the underline distribution is smooth. It is a straightforward corollary that this lower bound transfers to the online learning setting and we get the following corollary.

COROLLARY 6.3. *For any online learning algorithm ALG and every T there is a distribution \mathcal{D} that is an 8-smooth distribution over $[0, 1]^2$, such that $AVGREGET(\mathcal{D}^{\otimes T}, (M_t)_{t \in [T]}, \mathcal{M}) \geq \frac{T}{192}$, where $(M_t)_{t \in [T]}$ is the sequence of auctions generated by ALG and \mathcal{M} is the set of DSIC and IR auctions.*

7 CONCLUSIONS

In this paper, we give the first positive results on the online learnability of large and non-parametric classes of auctions. To achieve this, we introduce notions of smooth sets and smooth auctions and we show novel learning theoretic results for learning smooth sets and smooth distributions. Our positive result is that learning revenue-optimal smooth auctions over smooth distributions is possible in the non-adaptive and the adaptive cases. We complete our positive results with lower bounds showing that even if one of the smoothness properties is missing, namely the auction class is not smooth or if the bid distribution is not smooth, then online learning of revenue-optimal auctions is impossible.

Our work initiates the study of smooth auctions in the online learning setting and brings forth several interesting open directions:

- *Understand the set of distributions for which the revenue-optimal auction is smooth (or even monotone or competition monotone).* Ignoring optimality, we have already highlighted classes of commonly used auctions that are smooth, revenue monotone, or competition monotone. We have also showed Myerson's optimal auction is smooth. An interesting direction for future work is to examine when these classes of auctions include the optimal auction beyond the independent value distribution.
- *Explore the dependence of the rates on the parameters Γ , σ , n .* The main message of our work is the possibility of online learnability of smooth auctions and we have not dedicated significant effort to optimizing our online learnability rates. Our current bounds are of the form $(\Gamma/\sigma)^n$ and it would be interesting to see whether a rate that is polynomial in n is possible, i.e., $n^{\Gamma/\sigma}$.
- *Explore computationally efficient methods.* This paper is focused on the statistical aspects of the problem since this has been the bottleneck in the existing literature, but exploring computationally efficient methods is a very interesting direction as well.
- *Approximation power of non-parametric auctions.* More generally, our work motivates the study of smooth, monotone, and competition monotone auctions further in both offline correlated and online models. Promising directions for future work include understanding the revenue gap between the optimal correlated auction and that of the optimal smooth, monotone, or competition monotone auctions.

8 ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under grant CCF-2145898, a C3.AI Digital Transformation Institute grant, and Berkeley AI Research Commons grants. This work was partially done while Haghtalab and Zampetakis were visitors at the Simons Institute for the Theory of Computing. Zampetakis was also supported by the Army Research Office under contract W911NF-17-1-0304 as part of the collaboration between US DOD, UK MOD and UK Engineering and Physical Research Council under the Multidisciplinary University Research Initiative.

REFERENCES

- Noga Alon, Shai Ben-David, Nicolo Cesa-Bianchi, and David Haussler. 1997. Scale-sensitive dimensions, uniform convergence, and learnability. *J. ACM* 44, 4 (1997), 615–631.
- Martin Anthony, Peter L Bartlett, Peter L Bartlett, et al. 1999. *Neural network learning: Theoretical foundations*. Vol. 9. Cambridge University Press, Cambridge.
- Maria-Florina Balcan and Avrim Blum. 2006. Approximation algorithms and online mechanisms for item pricing. In *Proceedings of the 7th ACM Conference on Economics and Computation (EC)*. 29–35.
- Maria-Florina Balcan, Travis Dick, and Ellen Vitercik. 2018. Dispersion for data-driven algorithm design, online learning, and private optimization. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 603–614.
- Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. 2009. Agnostic Online Learning. In *Proceedings of the Conference on Learning Theory (COLT)*.
- Dirk Bergemann, Benjamin Brooks, and Stephen Morris. 2020. Countering the winner’s curse: Optimal auction design in a common value model. *Theoretical Economics* 15, 4 (2020), 1399–1434.
- Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin. 2022. Smoothed online learning is as easy as statistical learning. In *Proceedings of the Conference on Learning Theory (COLT)*. PMLR, 1716–1786.
- Avrim Blum and Jason D. Hartline. 2005. Near-Optimal Online Auctions. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, USA, 1156–1163.
- Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K Warmuth. 1989. Learnability and the Vapnik-Chervonenkis dimension. *J. ACM* 36, 4 (1989), 929–965.
- Sebastien Bubeck, Nikhil R Devanur, Zhiyi Huang, and Rad Niazadeh. 2017. Online auctions and multi-scale online learning. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*. 497–514.
- Nicolo Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. 2014. Regret minimization for reserve prices in second-price auctions. *IEEE Transactions on Information Theory* 61, 1 (2014), 549–564.
- Nicolo Cesa-Bianchi and Gábor Lugosi. 1999. On prediction of individual sequences. *The Annals of Statistics* 27, 6 (1999), 1865–1895.
- Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *Proceedings of the 47th annual ACM Symposium on Theory of computing (STOC)*. 243–252.
- Jacques Cr  mer and Richard P McLean. 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society* (1988), 1247–1257.
- Constantinos Daskalakis and Vasilis Syrgkanis. 2022. Learning in auctions: Regret is hard, envy is easy. *Games and Economic Behavior* 134 (2022), 308–343.
- Nikhil R Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. 2016. The sample complexity of auctions with side information. In *Proceedings of the 48th annual ACM Symposium on Theory of Computing (STOC)*. 426–439.
- Miroslav Dudik, Nika Haghtalab, Haipeng Luo, Robert E Schapire, Vasilis Syrgkanis, and Jennifer Wortman Vaughan. 2020. Oracle-efficient online learning and auction design. *J. ACM* 67, 5 (2020), 1–57.
- Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. 1956. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics* (1956), 642–669.
- Yoav Freund and Robert E Schapire. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences* 55, 1 (1997), 119–139.
- Hu Fu, Nima Haghtalab, Jason Hartline, and Robert Kleinberg. 2014. Optimal auctions for correlated buyers with sampling. In *Proceedings of the 2014 ACM conference on Economics and computation (EC)*. 23–36.
- Edgar N Gilbert. 1952. A comparison of signalling alphabets. *The Bell system technical journal* 31, 3 (1952), 504–522.
- Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. 2019. Settling the sample complexity of single-parameter revenue maximization. In *Proceedings of the 51st annual ACM Symposium on Theory of Computing (STOC)*. 662–673.
- Rishi Gupta and Tim Roughgarden. 2016. A PAC approach to application-specific algorithm selection. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science (ITCS)*. 123–134.
- Nika Haghtalab. 2018. *Foundation of Machine Learning, by the People, for the People*. Ph. D. Dissertation. Carnegie Mellon University.
- Nika Haghtalab, Yanjun Han, Abhishek Shetty, and Kunhe Yang. 2022. Oracle-efficient online learning for smoothed adversaries. *Advances in Neural Information Processing Systems* 35 (2022), 4072–4084.
- Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. 2020. Smoothed Analysis of Online and Differentially Private Learning. *Advances in Neural Information Processing Systems* 33 (2020), 9203–9215.
- Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. 2021. Smoothed analysis with adaptive adversaries. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 942–953.
- Nick Littlestone. 1988. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning* 2, 4 (1988), 285–318.

- Bodo Manthey. 2020. Smoothed Analysis of Local Search. In *Beyond the Worst-Case Analysis of Algorithms*, Tim Roughgarden (Ed.). Cambridge University Press, Chapter 13, 285–308.
- Paul R Milgrom and Robert J Weber. 1982. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* (1982), 1089–1122.
- Jamie H Morgenstern and Tim Roughgarden. 2015. On the pseudo-dimension of nearly optimal auctions. *Advances in Neural Information Processing Systems* 28 (2015).
- Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
- Alexandros Psomas, Ariel Schwartzman, and S Matthew Weinberg. 2019. Smoothed analysis of multi-item auctions with correlated values. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*. 417–418.
- Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. 2011. Online learning: Stochastic, constrained, and smoothed adversaries. *Advances in neural information processing systems* 24 (2011).
- Tim Roughgarden and Joshua R Wang. 2019. Minimizing regret with multiple reserves. *ACM Transactions on Economics and Computation* 7, 3 (2019), 1–18.
- Daniel A Spielman and Shang-Hua Teng. 2004. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM* 51, 3 (2004), 385–463.
- Vladimir Vapnik and A Ya Chervonenkis. 1968. On the uniform convergence of relative frequencies of events to their probabilities. In *Doklady Akademii Nauk USSR*, Vol. 181. 781–787.
- Rom Rubenovich Varshamov. 1957. Estimate of the number of signals in error correcting codes. *Doklady Akad. Nauk, SSSR* 117 (1957), 739–741.
- Chunxue Yang and Xiaohui Bei. 2021. Learning Optimal Auctions with Correlated Valuations from Samples. In *International Conference on Machine Learning (ICML)*. PMLR, 11716–11726.
- Bin Yu. 1997. Assouad, fano, and le cam. In *Festschrift for Lucien Le Cam*. Springer, 423–435.