

Smoothness for Simultaneous Composition of Mechanisms with Admission^{*}

Martin Hoefer, Thomas Kesselheim, and Bojana Kodric

MPI Informatik and Saarland University, Saarbrücken, Germany
`{martin.hoefer,thomas.kesselheim,bojana.kodric}@mpi-inf.mpg.de`

Abstract. We study social welfare of learning outcomes in mechanisms with admission. In our repeated game there are n bidders and m mechanisms, and in each round each mechanism is available for each bidder only with a certain probability. Our scenario is an elementary case of simple mechanism design with incomplete information, where availabilities are bidder types. It captures natural applications in online markets with limited supply and can be used to model access of unreliable channels in wireless networks. If mechanisms satisfy a smoothness guarantee, existing results show that learning outcomes recover a significant fraction of the optimal social welfare. These approaches, however, have serious drawbacks in terms of plausibility and computational complexity. Also, the guarantees apply only when availabilities are stochastically independent among bidders. In contrast, we propose an alternative approach where each bidder uses a single no-regret learning algorithm and applies it in all rounds. This results in what we call *availability-oblivious* coarse correlated equilibria. It exponentially decreases the learning burden, simplifies implementation (e.g., as a method for channel access in wireless devices), and thereby addresses some of the concerns about Bayes-Nash equilibria and learning outcomes in Bayesian settings. Our main results are general composition theorems for smooth mechanisms when valuation functions of bidders are lattice-submodular. They rely on an interesting connection to the notion of correlation gap of submodular functions over product lattices.

1 Introduction

Truthful mechanism design is a central challenge at the intersection of economics and computer science, but many fundamental techniques are only very rarely used in practice. For example, sponsored search auctions are used on a daily basis and generate billions of dollars in revenue, but they are based on simple and non-truthful procedures to allocate ads on search result pages. In contrast, truthful mechanisms often involve heavy algorithmic machinery, complicated allocation techniques, or other hurdles to easy and transparent implementation.

A recent trend is to study non-truthful and conceptually “simple” mechanisms for allocation in markets and their inherent loss in system performance.

^{*} This work is supported by DFG through Cluster of Excellence MMCI.

The idea is to analyze the induced game among the bidders and bound the quality of (possibly manipulated) outcomes in equilibrium. In a seminal paper, Syrgkanis and Tardos [25] propose a general technique for bounding social welfare of these equilibria, based on a so-called “smoothness” technique. These guarantees apply even to mixed Bayes-Nash equilibria in environments with composition of mechanisms. For example, in a combinatorial auction we might not sell all items via a complicated truthful mechanism, but instead sell each item simultaneously via simple individual single-item auctions. Such a mechanism is obviously not truthful, since bidders are not even able to express their valuations for all subsets of items. However, if bidders have complement-free XOS valuations, the (expected) social welfare of allocations in a mixed Bayes-Nash equilibrium turns out to be a constant-factor approximation of the optimal social welfare.

While this is a fundamental insight into non-truthful mechanisms, it is not well-understood how this result extends under more realistic conditions. In particular, there has been recent concern about the plausibility and computational complexity of exact and approximate Bayes-Nash equilibria [7]. For more general Bayesian concepts based on no-regret learning strategies in repeated games, there are two natural approaches – either bidder types are drawn newly with bids, or types are drawn only once initially. While the latter is not really in line with the idea of incomplete information (bidders could communicate their type in the course of learning, see [7]), the former is in general hard to obtain. Also, the composition theorem applies only if bidders’ types are drawn independently.

In this paper, we study a variant of simultaneous composition of mechanisms and show how to avoid the drawbacks of the Bayesian approach. Our scenario is motivated by limited availability or admission: Suppose bidders try to acquire items in a repeated online market, in which m items are sold simultaneously via, say, first-price auctions. However, in each round only some of the items are actually available for purchase. This scenario can be phrased in the Bayesian framework when bidder i ’s type is given by the set of items available to him. To obtain an equilibrium in the Bayesian sense, each bidder would have to consider a complicated bid vector and satisfy an equilibrium condition for each of the possible 2^m subsets of items.

In contrast, here we assume that bidders do not even get to know (or are not able to account for) *their own availabilities* before making bids in each round. We assume they learn with no-regret strategies in a way that is *oblivious* to their own and all other bidders’ availabilities. Thereby, bidders arrive at what one might term an *availability-oblivious* coarse-correlated equilibrium – a bid distribution not tailored to the specific availabilities of bidders, which can be computed (approximately) in polynomial time. Our main result is that for a large class of valuation functions, we can apply smoothness ideas in this framework and prove bounds that mirror the guarantees above. The guarantees apply even if some bidders learn obliviously and others follow a Bayes-Nash bidding strategy. In particular, we cover a broad domain with simultaneous composition of weakly smooth mechanisms in the sense of [25] when bidders have lattice-submodular valuations. Our study covers cases where availabilities are correlated among bid-

ders and provides lower bounds for combinatorial auctions with item-bidding and XOS valuations. As a part of our analysis, we use the concept of correlation gap from [1] for submodular functions over product lattices.

1.1 Our Contribution

We assume that every mechanism satisfies a weak smoothness bound (for more details see Section 2 below) with parameters $\lambda, \mu_1, \mu_2 \geq 0$. It is known that for each individual mechanism, this implies an upper bound of $(\max(1, \mu_1) + \mu_2)/\lambda$ on the price of anarchy for no-regret learning outcomes and Bayes-Nash equilibria. Furthermore, the same bound also applies for outcomes of multiple simultaneous mechanisms that are tailored to availabilities, i.e., not oblivious.

In Section 3 we consider smoothness for oblivious learning and composition with independent availabilities, where in each round t , each mechanism j is available to each bidder i independently with probability $q_{i,j}$. Our smoothness bound involves the above parameters and the correlation gap of the class of valuation functions. In particular, if valuations v_i come from a class \mathcal{V} with a correlation gap of $\gamma(\mathcal{V})$, the price of anarchy becomes $\gamma(\mathcal{V}) \cdot (\max(1, \mu_1) + \mu_2)/\lambda$.

Our construction uses smoothness of simultaneous composition from [25]. However, since learning is oblivious, the deviations establishing smoothness must be independent of availability. Here we use correlation gap to relate the value for independent deviations to that of type-dependent Bayesian deviations. Correlation gap is a notion originally defined for submodular set functions in [2]. It captures the worst-case ratio between the expected value of independent and correlated distributions over elements with the same marginals. We use an extension of this notion from [1] to Cartesian products of outcome spaces such as product lattices. For the class \mathcal{V} of monotone lattice-submodular valuations, we prove a correlation gap of $\gamma(\mathcal{V}) = e/(e-1)$, which simplifies and slightly extends previous results.

In Section 4, we analyze oblivious learning for composition with correlated availabilities in the form of “everybody-or-nobody” – each mechanism is either available to all bidders or to no bidder. The probability for availability of mechanism j is q_j , and availabilities are independent among mechanisms. In this case, we simulate independence by assuming that each bidder draws random types and outcomes for himself. We also consider distributions where outcomes are drawn independently according to the marginals from the optimal correlated distribution over outcomes. While these two distributions are directly related via correlation gap, the technical challenge is to show that there is a connection to the value obtained by the bidder. For lattice-submodular functions, we show a smoothness bound that implies a price of anarchy of $4e/(e-1) \cdot (\max(1, \mu_1) + \mu_2)/\lambda^2$.

For neither of the results is it necessary that all bidders follow our oblivious-learning approach. We only require that bidders have no regret compared to this strategy. This is also fulfilled if some or all bidders determine their bids based on the actually available items rather than in the oblivious way.

Finally, in Section 5 we show a lower bound for simultaneous composition of single-item first-price auctions with general XOS valuation functions. The corre-

lation gap for such functions is known to be large [2], but this does not directly imply a lower bound on the price of anarchy for oblivious learning. We provide a class of instances where the price of anarchy for oblivious learning becomes $\Omega((\log m)/(\log m \log m))$. This shows that for XOS functions it is impossible to generalize the constant price of anarchy for single-item first-price auctions.

Our results have additional implications beyond auctions for the analysis of regret learning in wireless networks. We discuss these in Appendix A.

1.2 Related Work

Closely related to our work are combinatorial auctions with item bidding, where multiple items are being sold in separate auctions. Bidders are generally interested in multiple items. However, depending on the bidder, some items may be substitutes for others. As the auctions work independently, bidders have to strategize in order to buy not too many items simultaneously. In a number of papers [9,5,18,16] the efficiency of Nash and Bayes-Nash equilibria has been studied. It has been shown that, if the single items are sold in first or second price auctions and if the valuation functions are XOS or subadditive, the price of anarchy is constant. Limitations of this approach are shown in [10,24].

Many of these proofs follow a similar pattern, namely showing smoothness. This concept has been introduced by [22,23] to analyze correlated and Bayes-Nash equilibria of general games. In [25] it was adjusted to mechanisms, and it was shown that simultaneous or sequential composition of smooth mechanisms is again smooth. Combinatorial auctions with item bidding are an example of a simultaneous composition. To show smoothness of the combined mechanism, it is thus enough to show smoothness of each single auction. Other examples of smooth mechanisms are position auctions with generalized second price [8,21] and greedy auctions [19]. The smoothness approach for mixed Bayes-Nash equilibria shown in [25] is, in fact, slightly more general and continues to hold for variants of Bayesian correlated equilibrium [17].

The complexity of finding such equilibria has been studied only very recently. It has been shown in [7,15] that equilibria are hard to find in some settings. In contrast, in [14] a different auction format is studied that yields good bounds on social welfare for equilibria that can be found more easily. Although similar in spirit, our approach is different – it shows that in some scenarios agents can reduce the computational effort and still obtain reasonably good states with existing mechanisms.

As such, our approach is closer to recent work [13] that shows hardness results for learning full-information coarse-correlated equilibria in simultaneous single-item second-price auctions with unit-demand bidders. As a remedy, a form of so-called no-envy learning is proposed, in which bidders use a different form of bidding that enables convergence in polynomial time. While achieving a general no-regret guarantee against all possible bid vectors is hard, we note here that our approach based on smoothness requires only a guarantee with respect to bids that are derived directly from the XOS representation of the bidder valuation. As such, bidders can obtain the guarantees required for our results in

polynomial time. Conceptually, we here treat a different problem – the impact of availabilities, and more generally, different bidder types on learning outcomes in repeated mechanism design.

A model with dynamic populations in games has recently been considered in [20]. Each round a small portion of players are replaced by others with different utility functions. When players use algorithms that minimize a notion called adaptive regret, smoothness conditions and the resulting bounds on the price of anarchy continue to hold if there are solutions which remain near-optimal over time with a small number of structural changes. Using tools from differential privacy, these conditions are shown for some special classes of games, including first-price auctions with unit-demand or gross-substitutes valuations. In contrast, our scenario is orthogonal, since we consider much more general classes of mechanisms and allow changes in each round for possibly all players. However, our model of change captures the notion of availability and therefore is much more specific than the adversarial approach of [20].

The notion of correlation gap was defined and analyzed for stochastic optimization in [2,1]. The notion was used in [26] for analyzing revenue maximization with sequential auctions, which is very different from our approach.

2 Model and Preliminaries

There are n bidders that participate in m simultaneous mechanisms. Each mechanism $j \in [m]$ is a pair $M_j = (f_j, p_j)$, consisting of an outcome function and payment functions. More formally, function $f_j: B_j \rightarrow \mathcal{X}_j$ maps every bid vector $b_{\cdot,j}$ on mechanism j into an outcome space \mathcal{X}_j . The function $p_j = (p_{1,j}, \dots, p_{n,j})$ defines a payment for each bidder. That is, depending on the bid vector, $p_{i,j}: B_j \rightarrow \mathbb{R}_{\geq 0}$ defines the non-negative payment for bidder i in mechanism j .

We consider a repeated framework with oblivious learning in a simultaneous composition of mechanisms with availabilities. There are T rounds and in each round the bidders participate in m simultaneous mechanisms. In round $t = 1, \dots, T$, each bidder places a bid $b_{i,j}^t$ for each mechanism, the mechanism determines the outcome and the payments, and bidder i has a utility function $u_i(b^t) = v_i(f(b^t)) - p_i(b^t)$, where v_i is a valuation function over vectors of outcomes and $p_i = \sum_j p_{i,j}(b^t)$. In addition, in each round we assume that each mechanism is available to each bidder with a certain probability. We let the Bernoulli random variable $A_{i,j} = 1$ if mechanism j is available to bidder i . Due to availability, the mechanisms must also be applicable when only subsets of bidders are placing bids. For this reason, it will be convenient to assume that the outcome space for mechanism $j \in [m]$ is $\mathcal{X}_j = \mathcal{X}_{1,j} \times \dots \times \mathcal{X}_{n,j}$ and $x_j \in \mathcal{X}_j$ is $x_j = (x_{i,j})_{i \in [n]}$. We assume that each bidder, for whom the mechanism is not available, must place a bid of “0”. If bidder i bids 0 for mechanism j , we assume $f_j(0, b_{-i,j}) = \perp_{i,j}$, where $\perp_{i,j}$ is a “losing” outcome, and payment $p_{i,j}(0, b_{-i,j}) = 0$. For convenience, we will denote by $f = (f_j)_{j \in [m]}$ the composed mechanism and by $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_m$ its outcome space.

Oblivious Learning We assume *oblivious learning* – each bidder runs a single no-regret learning algorithm and uses the utility of every round as feedback, no matter how the availability in each round turns out. In hindsight, the average history of play for oblivious learning becomes an availability-oblivious variant of coarse-correlated equilibrium [6]. Hence, the outcomes of oblivious learning are captured by the coarse-correlated equilibria in the following one-shot game: First, all bidders simultaneously place a bid for every mechanism. They know only the probability distribution of the availabilities. Only after they placed their bids, the availability of each mechanism for each bidder is determined at random.

Definition 1. *An availability-oblivious coarse-correlated equilibrium is a distribution over bid vectors b (independent of A) such that, in expectation over all availabilities, it is not beneficial for any bidder i to switch to another bid b'_i . For each i and each b'_i , we have $\mathbf{E}[u_i(b'_i, b_{-i})] \leq \mathbf{E}[u_i(b)]$.*

Indeed, our results also hold for a larger class of equilibria, in which a subset of bidders might not be oblivious to availabilities. For our guarantees, it is enough to consider distributions over bidding strategies b which might depend on A such that, in expectation over all availabilities, it is not beneficial for any bidder i to switch to another bid b'_i . For each i and each b'_i , we have $\mathbf{E}[u_i(b'_i, b_{-i})] \leq \mathbf{E}[u_i(b)]$. Note that both ordinary coarse-correlated equilibria and availability-oblivious ones fulfill this property.

We bound the performance of these equilibria by deriving suitable smoothness bounds.

Smoothness We assume that each mechanism j satisfies *weak smoothness* as defined in [25]. For any valuations $v_{i,j}: \mathcal{X}_j \rightarrow \mathbb{R}^{\geq 0}$ there are (possibly randomized) deviations¹ $b'_{i,j}$ for each $i \in [n]$ such that for all bid vectors $b_{\cdot,j}$

$$\begin{aligned} & \mathbf{E} \left[\sum_{i \in [n]} v_{i,j}(f_j(b'_{i,j}, b_{-i,j})) - p_{i,j}(b'_{i,j}, b_{-i,j}) \right] \\ & \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - \mu_1 \cdot \sum_{i \in [n]} p_{i,j}(b_{\cdot,j}) - \mu_2 \sum_{i \in [n]} h_{i,j}(b_{i,j}, f_j(b_{\cdot,j})) \quad , \quad (1) \end{aligned}$$

where $h_{i,j}(b_{i,j}, x_j) = \max_{b_{-i,j}: f_j(b_{\cdot,j}) = x_j} p_{i,j}(b_{\cdot,j})$. For intuition, assume that (1) holds with $\mu_2 = 0$. Consider a learning outcome with a no-regret guarantee where every bidder i can gain at most ϵ in any fixed deviation, i.e., $\mathbf{E}[v_{i,j}(f_j(b_{\cdot,j})) - p_{i,j}(b_{\cdot,j})] \geq \mathbf{E}[v_{i,j}(f_j(b'_{i,j}, b_{-i,j})) - p_{i,j}(b'_{i,j}, b_{-i,j})] - \epsilon$. Applying (1) pointwise

$$\sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j})) - p_{i,j}(b_{\cdot,j})] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - \mu_1 \cdot \sum_{i \in [n]} \mathbf{E}[p_{i,j}(b_{\cdot,j})] - n\epsilon,$$

¹ In slight contrast to [25], we here assume that the smoothness deviations of a bidder do not depend on his own current bid. This serves to simplify our exposition and can be incorporated into our analysis.

which implies for social welfare

$$\sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) + (1 - \mu_1) \cdot \sum_{i \in [n]} \mathbf{E}[p_{i,j}(b_{\cdot,j})] - n\epsilon.$$

Every bidder i can stay away from the market and payments are non-negative, so $0 \leq \mathbf{E}[p_{i,j}(b_{\cdot,j})] \leq \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] + \epsilon$ and

$$\max(1, \mu_1) \sum_{i \in [n]} \mathbf{E}[v_{i,j}(f_j(b_{\cdot,j}))] \geq \lambda \cdot \max_{x_j \in \mathcal{X}_j} \sum_{i \in [n]} v_{i,j}(x_j) - (n + \mu_1)\epsilon.$$

Thus, for $\epsilon \rightarrow 0$, the price of anarchy tends to $\max(1, \mu_1)/\lambda$. More generally, (1) implies a bound on the price of anarchy of $(\mu_2 + \max(1, \mu_1))/\lambda$ for many equilibrium concepts. If $\mu_2 > 0$, then the bound relies on an additional no-overbidding assumption, which directly transfers to our results. For details see [25].

Valuation Functions Our main results apply for the class of monotone lattice-submodular valuations. Suppose for every mechanism j the set $\mathcal{X}_{i,j}$ of possible outcomes for bidder i forms a lattice $(\mathcal{X}_{i,j}, \succeq_{i,j})$ with a partial order $\succeq_{i,j}$. Bidder i has a *lattice-submodular valuation* v_i if and only if it is submodular on the product lattice $(\mathcal{X}_i, \succeq_i)$ of outcomes for bidder i : $\forall x_i, \tilde{x}_i \in \mathcal{X}_i : v_i(x_i \vee \tilde{x}_i) + v_i(x_i \wedge \tilde{x}_i) \leq v_i(x_i) + v_i(\tilde{x}_i)$. In the paper, we concentrate on distributive lattices, for which this definition is equivalent to the diminishing marginal returns property:

$$\forall z_i \succeq_i y_i \in \mathcal{X}_i \implies \forall t \in \mathcal{X}_i : v_i(t \vee y_i) - v(y_i) \geq v_i(t \vee z_i) - v(z_i).$$

Lattice-submodular functions generalize submodular set functions but are a strict subclass of XOS functions. Bidder i has an *XOS valuation* v_i if and only if there are additive functions v_i^1, v_i^2, \dots with $v_i^{k_i}(x_i) = \sum_j v_{i,j}^{k_i}(x_{i,j})$ for every $x_{i,j} \in \mathcal{X}_{i,j}$ and $v_i(x_i) = \max_{k_i} v_i^{k_i}(x_i)$.

3 Composition with Independent Admission

We first consider simultaneous composition of smooth mechanisms with independent availabilities. Here, all random variables $A_{i,j}$ are independent, and we let $q_{i,j} = \Pr[A_{i,j} = 1]$.

Definition 2. Let v be a valuation function on a product lattice, coming from a class of valuation functions \mathcal{V} . Given vectors x^1, \dots, x^k and numbers $\alpha_1, \dots, \alpha_k \in [0, 1]$ such that $\sum_{j=1}^k \alpha_j = 1$, determine another vector y at random by setting component y_i to x_i^j independently with probability α_j . Then, the smallest γ s.t. $\sum_{j=1}^k \alpha_j v(x^j) \leq \gamma \cdot \mathbf{E}[v(y)]$ is the correlation gap of class \mathcal{V} .

Theorem 1. Suppose bidder valuations are monotone and come from a class \mathcal{V} with a correlation gap of $\gamma(\mathcal{V})$. The price of anarchy for oblivious learning for simultaneous composition of weakly (λ, μ_1, μ_2) -smooth mechanisms with valuations from \mathcal{V} and fully independent availability is at most $\gamma(\mathcal{V}) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$.

Before the proof of the main theorem of this section, we note that in Appendix C.1 we also prove an upper bound of $e/(e-1)$ on the correlation gap of lattice-submodular valuations with diminishing marginal returns. This result slightly generalizes the result of [1] from composition of totally ordered sets to arbitrary product lattices.

Lemma 1 (Correlation Gap on a Product Lattice). *Let v be a function with diminishing marginal returns on a product lattice. Given vectors x^1, \dots, x^k and numbers $\alpha_1, \dots, \alpha_k \in [0, 1]$ such that $\sum_{j=1}^k \alpha_j = 1$, determine another vector y at random by setting component y_i to x_i^j independently with probability α_j . Then $\mathbf{E}[v(y)] \geq (1 - \frac{1}{e}) \sum_{j=1}^k \alpha_j v(x^j)$.*

From here, we arrive at the following corollary of the main theorem.

Corollary 1. *The price of anarchy for oblivious learning for simultaneous composition of weakly (λ, μ_1, μ_2) -smooth mechanisms with monotone lattice-submodular valuations and fully independent availability is at most $e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$.*

Proof of Theorem 1. We will prove the theorem by defining an availability-oblivious (randomized) deviation b'_i for each player i such that the following inequality will hold for any (not necessarily availability-oblivious) bidding strategy b :

$$\begin{aligned} & \sum_i \mathbf{E}[u_i(b'_i, b_{-i})] \\ & \geq \frac{1}{\gamma(\mathcal{V})} \cdot \lambda \cdot \sum_i \mathbf{E}[v_i(x^*)] - \mu_1 \sum_i \mathbf{E}[p_i(b)] - \mu_2 \sum_i \mathbf{E}[h_i(b_i, f(b))] \quad , \quad (2) \end{aligned}$$

where x^* denotes the (random) optimal outcome. From this inequality, whose form is in fact exactly that of the smoothness condition (1), the claim of the theorem follows as described in Section 2.

In more detail, to attain the aforementioned inequality, we will relate each player's utility for deviating to b'_i to the utility he could achieve if he was allowed to see and react upon the availabilities. In that case, he could simply use the smoothness deviation tailored to the specific availability profile $A_i = (A_{i,1}, \dots, A_{i,m})$ that he is encountering. We denote this non-oblivious smoothness deviation by $b_i^{A_i}$. Because the global mechanism is a simultaneous composition of (λ, μ_1, μ_2) -smooth mechanisms, it is again (λ, μ_1, μ_2) -smooth. Therefore we know that the non-oblivious deviations $b_i^{A_i}$ do exist, and they satisfy the smoothness inequality (1) by definition.

We proceed to define, for each player i , the availability-oblivious deviation b'_i . First, bidder i assumes for himself a reduced valuation function $\bar{v}_i = \alpha \cdot v_i$, for some appropriate α to be chosen later. The deviation b'_i is a composition of component-wise independent deviations $b'_{i,j}$, i.e. $b'_i = (b'_{i,1}, \dots, b'_{i,m})$ where each $b'_{i,j}$ is chosen independently. To arrive at $b'_{i,j}$, bidder i assumes that mechanism j is available to him and draws all other availabilities independently according

to probabilities $q_{i',j'}$. This means that he draws availabilities for all other players on all mechanisms and also his own availabilities on all mechanisms other than j . Now he has a full availability profile, and therefore he can consider the non-oblivious smoothness deviation. He observes the j -th component of this smoothness deviation and sets $b'_{i,j}$ to be equal to it. Note that $b'_{i,j}$ will be applied only with the probability that mechanism j is in fact available to bidder i , i.e. with probability $q_{i,j}$.

Next, we want to compare $u_i(b'_i, b_{-i})$ and $u_i(b_i^{A_i}, b_{-i})$. Let us focus on the valuation $v_i(f(b'_i, b_{-i}))$ first. The non-oblivious smoothness deviation $b_i^{A_i}$ is a vector whose components are correlated. More precisely, to form this bid we observe A_i , sample the availabilities A_{-i} and bids b_{-i} of other players, and take the optimal allocation x^* for the resulting availability profile A . Then, we determine the ℓ for which $\bar{v}_i(x_i^*) = \sum_j \bar{v}_{i,j}^\ell(x_{i,j}^*)$ and use $\bar{v}_{i,j}^\ell$ for determining $b_{i,j}^{A_i}$ (note that A_i can be regarded as bidder i 's type in a Bayesian sense, for more details see [25]). Therefore, the components of $b_i^{A_i}$ are correlated through the common choice of ℓ . Our deviation b'_i is assembled by setting $b'_{i,j} = (b_{i,j}^{A_i})_{k_j}$ independently for each j .

Formally, let $r_{i,j}^\ell$ denote the conditional probability that the optimum yields an outcome vector x^* that attains its maximum value for bidder i in \bar{v}_i^ℓ , given that $A_{i,j} = 1$. Then, the marginal probability of observing $b_{i,j}^{A_i} = (b_{i,j}^{A_i})_\ell$ is $r_{i,j}^\ell q_{i,j}$. In b'_i we pick ℓ independently for each mechanism with probability $r_{i,j}^\ell$, which yields a combined probability of $r_{i,j}^\ell q_{i,j}$ for availability and deviation. Thus, b'_i simulates the marginal probabilities of outcomes in $b_i^{A_i}$, i.e., $\Pr[f_j(b'_i, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}] = \Pr[f_j(b_i^{A_i}, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}]$ for all $y_{i,j} \in \mathcal{X}_{i,j}$, for each $j \in [m]$. Hence, for fixed A_{-i}, b_{-i} , the two expected valuations $\mathbf{E}[v_i(f(b'_i, b_{-i})) \mid A_{-i}, b_{-i}]$ and $\mathbf{E}[v_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}]$ are related via correlation gap.

Thus, setting $\alpha = 1/\gamma(\mathcal{V})$ and $\bar{v}_i(x) = 1/\gamma(\mathcal{V}) \cdot v_i(x)$ we get

$$\begin{aligned} \mathbf{E}[v_i(f(b'_i, b_{-i})) \mid A_{-i}, b_{-i}] &= \sum_{y \in \mathcal{X}} v_i(y) \cdot \Pr[f(b'_i, b_{-i}) = y \mid A_{-i}, b_{-i}] \\ &= \sum_{y \in \mathcal{X}} v_i(y) \cdot \prod_j \Pr[f_j(b'_i, b_{-i}) = y_{i,j} \mid A_{-i}, b_{-i}] \\ &\geq \frac{1}{\gamma(\mathcal{V})} \cdot \sum_{y \in \mathcal{X}} v_i(y) \cdot \Pr[f(b_i^{A_i}, b_{-i}) = y \mid A_{-i}, b_{-i}] \\ &= \frac{1}{\gamma(\mathcal{V})} \cdot \mathbf{E}[v_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}] \\ &= \mathbf{E}[\bar{v}_i(f(b_i^{A_i}, b_{-i})) \mid A_{-i}, b_{-i}]. \end{aligned}$$

In addition, because payments are simply additive across mechanisms, it is straightforward to see that for every bidder i

$$\mathbf{E}[p_i(b'_i, b_{-i}) \mid A_{-i}, b_{-i}] = \mathbf{E}[p_i(b_i^{A_i}, b_{-i}) \mid A_{-i}, b_{-i}].$$

This allows to apply the smoothness bound for Bayesian mechanisms with independent types from [25] to derive

$$\begin{aligned}
& \sum_i \mathbf{E}[u_i(b'_i, b_{-i})] \\
&= \sum_i \mathbf{E}[v_i(f(b'_i, b_{-i}))] - \mathbf{E}[p_i(b'_i, b_{-i})] \\
&\geq \sum_i \mathbf{E}[\bar{v}_i(f(b_i^{A_i}, b_{-i}))] - \mathbf{E}[p_i(b_i^{A_i}, b_{-i})] \\
&\geq \lambda \cdot \sum_i \mathbf{E}[\bar{v}_i(x^*)] - \mu_1 \sum_i \mathbf{E}[p_i(b)] - \mu_2 \sum_i \mathbf{E}[h_i(b_i, f(b))] \\
&= \frac{\lambda}{\gamma(\mathcal{V})} \cdot \sum_i \mathbf{E}[v_i(x^*)] - \mu_1 \sum_i \mathbf{E}[p_i(b)] - \mu_2 \sum_i \mathbf{E}[h_i(b_i, f(b))]
\end{aligned}$$

This proves the desired smoothness guarantee and implies the theorem. \square

4 Composition with Everybody-or-Nobody Admission

We consider the case in which at each point in time each mechanism is either available to all bidders or to none. We let $A_j = A_{i,j}$ for all $i \in [n]$ and $q_j = \Pr[A_j = 1]$. Note that all A_j are assumed to be independent.

Let the social optimum be denoted by x^* . We assume that $x_j^* = \perp_j$ if $A_j = 0$. Otherwise, x^* might have different values, depending on the availabilities of other mechanisms. Let us denote the possible outcomes by x_j^1, x_j^2, \dots and let $r_j^\ell := \Pr[x_j^* = x_j^\ell \mid A_j = 1]$. That is, r_j^ℓ is the marginal probability of x_j^ℓ conditioned on j being available. Theorem 2 formulates our main result in this section.

Theorem 2. *The price of anarchy for oblivious learning for simultaneous composition of weakly (λ, μ_1, μ_2) -smooth mechanisms with monotone lattice-submodular valuations and everybody-or-nobody admission is at most $4e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda^2$.*

Proof. We will prove that, for each bidder i and each mechanism j there are randomized deviation strategies $b'_{i,j}$ that are independent of the availabilities such that the following smoothness guarantee holds against any (potentially non-oblivious) bidding strategy b :

$$\begin{aligned}
& \sum_i \mathbf{E}[u_i(b'_i, b_{-i})] \\
&\geq \left(1 - \frac{1}{e}\right) \frac{\lambda^2}{4} \sum_i \mathbf{E}[v_i(x^*)] - \mu_1 \sum_i \mathbf{E}[p_i(b)] - \mu_2 \sum_i \mathbf{E}[h_i(b_i, f(b))] .
\end{aligned}$$

From this guarantee the claim of the theorem again follows as described in Section 2.

To define $b'_{i,j}$, every bidder i draws two vectors z^i and \tilde{t}^i at random as follows. He sets z_j^i to x_j^ℓ with probability r_j^ℓ/α , where $\alpha = 2/\lambda$, and to \perp_j with the remaining probability. Furthermore, he sets \tilde{t}_j^i to x_j^ℓ with probability $q_j r_j^\ell$ and to \perp_j with the remaining probability. These draws are performed independent of any availabilities. Observe that for each i , we have $\mathbf{E} [\sum_{i'} v_{i'}(\tilde{t}^i)] \geq (1 - \frac{1}{e}) \mathbf{E} [\sum_{i'} v_{i'}(x^*)]$ by Lemma 1.

Due to the random draws, each bidder i' defines functions $w_{i',j}^{i'}: \Omega_j \rightarrow \mathbb{R}$ for each bidder i and each mechanism j . Function $w_{i',j}^{i'}$ maps an outcome of mechanism j , denoted by y_j , to a real number as follows

$$w_{i',j}^{i'}(y_j) = v_i(\tilde{t}_1^{i'}, \dots, \tilde{t}_{j-1}^{i'}, y_j \wedge z_j^{i'}, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^{i'}, \dots, \tilde{t}_{j-1}^{i'}, \perp_j, \dots, \perp_m) .$$

Note that these functions do not necessarily reflect the actual value any outcome might have. They are only used to define the deviation strategy: bidder i' pretends all bidders i , including himself, would have valuations $w_{i',j}^{i'}$ for the outcome of mechanism j . This gives him a deviation strategy $b'_{i',j}$ by setting $b'_{i',j} = b_{i',j}^*(w_{1,j}^{i'}, \dots, w_{n,j}^{i'})$ as defined by the smoothness of mechanism j .

The proofs for the following three lemmas are presented in Appendix C.2, C.3, C.4.

Lemma 2. *For every bidder i and deviating bids $b'_{i,j} = b_{i,j}^*(w_{1,j}^i, \dots, w_{n,j}^i)$,*

$$\mathbf{E} [v_i(f(b'_i, b_{-i}))] \geq \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i}))] - \frac{1}{\alpha(\alpha+1)} \mathbf{E} [v_i(\tilde{t}^i)] .$$

Lemma 3. *For the adjusted functions w we can apply smoothness to obtain*

$$\begin{aligned} & \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] . \end{aligned}$$

Lemma 4. *For function w^1 , random vectors z_j^1 and \tilde{t}^1 , and every mechanism j*

$$\sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] = \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^1)] .$$

The bound from Lemma 3 has striking similarities to the smoothness bound (1). However, it is expressed in terms of the functions $w_{i,j}^{i'}$ rather than the actual valuation functions v_i . The other two Lemmas show that, in expectation, these functions are close enough to the functions v_i so that this bound actually suffices to prove the main result:

$$\sum_i \mathbf{E} [u_i(b'_i, b_{-i})] = \sum_i \mathbf{E} \left[v_i(f(b'_i, b_{-i})) - \sum_j p_{i,j}(b'_{i,j}, b_{-i}) \right]$$

$$\begin{aligned}
&\geq \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] - \frac{1}{\alpha(\alpha+1)} \sum_i \mathbf{E} [v_i(\tilde{t}^i)] \\
&\hspace{25em} \text{(by Lemma 2)} \\
&\geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \\
&\quad - \frac{1}{\alpha(\alpha+1)} \sum_i \mathbf{E} [v_i(\tilde{t}^1)] \hspace{10em} \text{(by Lemma 3)} \\
&= \sum_i \left(\frac{\lambda}{\alpha} - \frac{1}{\alpha(\alpha+1)} \right) \mathbf{E} [v_i(\tilde{t}^i)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \ . \\
&\hspace{25em} \text{(by Lemma 4)}
\end{aligned}$$

By setting $\alpha = \frac{2}{\lambda}$

$$\begin{aligned}
\sum_i \mathbf{E} [u_i(b'_i, b_{-i})] &\geq \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(\tilde{t}^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \\
&\geq \left(1 - \frac{1}{e}\right) \frac{\lambda^2}{4} \sum_i \mathbf{E} [v_i(x^*)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \ .
\end{aligned}$$

The last step follows from Lemma 1.

Note that technically the mechanism could be randomized itself. Our results extend to this case in a straightforward way.

5 A Lower Bound for General XOS Functions

In this section we consider combinatorial auctions with item bidding and first-price auctions. We can apply the previous analysis, since for each bidder the outcomes form a trivial 2-element lattice – winning an item is the supremum outcome, not winning is the infimum outcome. In the analysis, observe that each bidder determines a random allocation of items according to the probabilities in the optimum. Based on these allocations, bidders determine the valuations $w_{i,j}^i$, which in turn form the basis for the deviation. The first-price auction with general bidding space is $(1 - 1/e, 1, 0)$ -smooth [25]. If valuation functions are submodular, the composition theorems can be applied to yield the following corollary.

Corollary 2. *The price of anarchy for oblivious learning for simultaneous composition of single-item first-price auctions with monotone submodular valuations and fully independent availability is at most $1/(1-1/e)^2$; for everybody-or-nobody admission it is at most $4/(1-1/e)^3$.*

For more general XOS valuations, we prove a lower bound that with oblivious bidding we will not be able to show a guarantee based on the smoothness parameters – even for a single bidder, so the bound applies without assumptions on correlation among bidders. The proof can be found in Appendix C.5.

Theorem 3. *In a simultaneous composition of discrete first-price single-item auctions with m items and XOS valuations, the price of anarchy for pure Nash equilibria with oblivious bidding can be as large as $\Omega((\log m)/(\log \log m))$, while each single mechanism is weakly $(1/2, 1, 0)$ -smooth.*

6 Conclusion

In this paper, we have studied an oblivious variant for no-regret learning in repeated games with incomplete information and proved a composition theorem for smooth mechanisms. The bounds show that even if bidders apply learning algorithms independently of their types, they can still obtain outcomes that approximate the optimal social welfare within a small ratio.

Our primary motivation are changes over time on the supply side. That is, bidders value items the same at all times but are constrained when they can buy them. A different interpretation that leads to the same model is when bidders value items differently from time to time. Here the valuation for a bundle has the special structure that it is given by the value of a fixed submodular function evaluated on the intersection of this bundle with a random set.

There is potential to generalize this approach to other interesting settings. For example, one could consider general independent types, where the complete availability-vector of a single bidder is drawn from a bidder-specific distribution, and for each bidder this is done independently. In Appendix B, we give a partial answer and show how our techniques can be extended to the following case. Consider simultaneous single-item auctions with unit-demand valuations, i.e., $v_i(S) = \max_{j \in S} v_{i,j}$. The distribution over valuations is such that for each item the value $v_{i,j}$ is independently drawn from a distribution of small support. Independent availabilities can be captured in this setting by setting $v_{i,j}$ to a fixed value or to 0 with the respective probabilities.

References

1. S. Agrawal. *Optimization under Uncertainty: Bounding the Correlation Gap*. PhD thesis, Stanford University, 2011.
2. S. Agrawal, Y. Ding, A. Saberi, and Y. Ye. Correlation robust stochastic optimization. In *Proc. 21st Symp. Discr. Algorithms (SODA)*, pages 1087–1096, 2010.
3. M. Andrews and M. Dinitz. Maximizing capacity in arbitrary wireless networks in the SINR model: Complexity and game theory. In *Proc. 28th IEEE Conf. Computer Communications (INFOCOM)*, pages 1332–1340, 2009.
4. E. I. Asgeirsson and P. Mitra. On a game theoretic approach to capacity maximization in wireless networks. In *Proc. 30th IEEE Conf. Computer Communications (INFOCOM)*, pages 3029–3037, 2011.
5. K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proc. 22nd Symp. Discr. Algorithms (SODA)*, pages 700–709, 2011.
6. A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In N. Nisan, É. Tardos, T. Roughgarden, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 4. Cambridge University Press, 2007.

7. Y. Cai and C. Papadimitriou. Simultaneous bayesian auctions and computational complexity. In *Proc. 15th Conf. Econom. Comput. (EC)*, pages 895–910, 2014.
8. I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. P. Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. *J. Econom. Theory*, 156:343–388, 2015.
9. G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proc. 35th Intl. Coll. Autom. Lang. Program. (ICALP)*, pages 820–832, 2008.
10. G. Christodoulou, A. Kovács, A. Sgouritsa, and B. Tang. Tight bounds for the price of anarchy of simultaneous first price auctions. *ACM Trans. Econom. Comput.*, 4(2):9, 2016.
11. J. Dams, M. Hoefer, and T. Kesselheim. Sleeping experts in wireless networks. In *Proc. 27th Intl. Symp. Distrib. Comput. (DISC)*, pages 344–357, 2013.
12. J. Dams, M. Hoefer, and T. Kesselheim. Jamming-resistant learning in wireless networks. In *Proc. 41st Intl. Coll. Autom. Lang. Program. (ICALP)*, volume 2, pages 447–458, 2014.
13. C. Daskalakis and V. Syrgkanis. Learning in auctions: Regret is hard, envy is easy. In *Proc. 57th Symp. Foundations of Computer Science (FOCS)*, 2016. To appear.
14. N. Devanur, J. Morgenstern, V. Syrgkanis, and M. Weinberg. Simple auctions with simple strategies. In *Proc. 16th Conf. Econom. Comput. (EC)*, pages 305–322, 2015.
15. S. Dobzinski, H. Fu, and R. D. Kleinberg. On the complexity of computing an equilibrium in combinatorial auctions. In *Proc. 26th Symp. Discr. Algorithms (SODA)*, pages 110–122, 2015.
16. M. Feldman, H. Fu, N. Gravin, and B. Lucier. Simultaneous auctions are (almost) efficient. In *Proc. 45th Symp. Theory of Computing (STOC)*, pages 201–210, 2013.
17. J. D. Hartline, V. Syrgkanis, and É. Tardos. No-regret learning in repeated Bayesian games. In *Proc. 28th Conf. Adv. Neural Information Processing Systems (NIPS)*, pages 3043–3051, 2015.
18. A. Hassidim, H. Kaplan, Y. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proc. 12th Conf. Electr. Commerce (EC)*, pages 295–296, 2011.
19. B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proc. 21st Symp. Discr. Algorithms (SODA)*, pages 537–553, 2010.
20. T. Lykouris, V. Syrgkanis, and É. Tardos. Learning and efficiency in games with dynamic population. In *Proc. 27th Symp. Discr. Algorithms (SODA)*, pages 120–129, 2016.
21. R. Paes Leme and É. Tardos. Pure and bayes-nash price of anarchy for generalized second price auction. In *Proc. 51st Symp. Foundations of Computer Science (FOCS)*, pages 735–744, 2010.
22. T. Roughgarden. Intrinsic robustness of the price of anarchy. *Comm. ACM*, 55(7):116–123, 2012.
23. T. Roughgarden. The price of anarchy in games of incomplete information. In *Proc. 13th Conf. Electr. Commerce (EC)*, pages 862–879, 2012.
24. T. Roughgarden. Barriers to near-optimal equilibria. In *Proc. 55th Symp. Foundations of Computer Science (FOCS)*, pages 71–80, 2014.
25. V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proc. 45th Symp. Theory of Computing (STOC)*, pages 211–220, 2013.
26. Q. Yan. Mechanism design via correlation gap. In *Proc. 22nd Symp. Discr. Algorithms (SODA)*, pages 710–719, 2011.

A Applications beyond Auctions

Our results have interesting implications beyond mechanisms that incorporate standard auction formats. A very intriguing one is channel allocation in wireless networks. The overall problem is to maximize the utilization of a wireless channel while avoiding interference. To this end, the following game was defined in [3]: Each player i corresponds to a pair of a sender s_i and a receiver r_i . The transmission from s_i to r_i is successful if the signal-to-interference-plus-noise ratio (SINR) is high enough. This means that the incoming interference from senders transmitting simultaneously plus ambient noise is by a factor smaller than the intended signal. Formally, transmission i is successful if

$$\frac{\frac{p}{d(s_i, r_i)^\alpha}}{\sum_{j \in S \setminus \{i\}} \frac{p}{d(s_j, r_i)^\alpha} + \nu} \geq \beta .$$

Here $p > 0$ is the (fixed) power level, $S \subseteq [n]$ is the set of simultaneous transmissions; $\alpha > 0$, $\beta > 0$, and $\nu \geq 0$ are constants.

To derive a game, each player has two strategies b_i : either he decides to transmit or not to. The best possible outcome is a successful transmission. An unsuccessful transmission is the worst possible outcome. Due to the energy consumption, it is considered to be even worse than not transmitting at all. This is reflected in the following utility function.

$$u_i(b) = \begin{cases} 1 & \text{if } b_i = 1 \text{ and } i \text{ is successful against } b_{-i} \\ -1 & \text{if } b_i = 1 \text{ and } i \text{ is not successful against } b_{-i} \\ 0 & \text{if } b_i = 0 \end{cases}$$

The robust price of anarchy of this game is constant [4]. In every coarse correlated equilibrium, the expected number of successful transmissions is only a constant smaller than the maximum possible number of simultaneous successful transmissions.

Quite surprisingly, this game corresponds to a smooth mechanism as follows. Each player decides whether to transmit; a player always has valuation 2 for making a successful transmission. However, whenever making a transmission (successful or not), the bidder has to pay 1. This is comparable to an all-pay auction, where each bidder has to pay his bid, regardless of whether he wins the respective item.

Theorem 4. *The mechanism representing the channel-allocation game is weakly $(1, \mu_1, \mu_2)$ -smooth for $\mu_1 = O(1)$ and $\mu_2 = 0$.*

Proof. Let $S \subseteq N$ be a maximum set of players that can transmit simultaneously. Define b' by setting $b'_i = 1$ for $i \in S$ and $b'_i = 0$ for $i \notin S$. That is, $u_i(b'_i, b_{-i}) = 0$ for all $i \notin S$. Consider some bid vector b , let T be the set of players making a transmission attempt. Note that by our definition $\sum_i p_i(b) = |T|$.

Furthermore, $i \in S$ is successful under (b'_i, b_{-i}) if and only if

$$\frac{\frac{p}{d(s_i, r_i)^\alpha}}{\sum_{j \in T \setminus \{i\}} \frac{p}{d(s_j, r_i)^\alpha} + \nu} \geq \beta ,$$

for which it is sufficient to have

$$\sum_{j \in T \setminus \{i\}} \frac{d(s_i, r_i)^\alpha}{d(s_j, r_i)^\alpha} + \frac{d(s_i, r_i)^\alpha}{p} \nu < \frac{1}{\beta} ,$$

which is equivalent to

$$\sum_{j \in T \setminus \{i\}} a_{j,i} < 1 \quad \text{where } a_{j,i} = \min \left\{ 1, \frac{1}{\frac{1}{\beta} - \frac{d(s_i, r_i)^\alpha}{p} \nu} \frac{d(s_i, r_i)^\alpha}{d(s_j, r_i)^\alpha} \right\} .$$

This implies $u_i(b'_i, b_{-i}) \geq 1 - 2 \sum_{j \in T \setminus \{i\}} a_{j,i}$. Taking the sum over all $i \in S$, we get

$$\sum_{i \in S} u_i(b'_i, b_{-i}) \geq |S| - 2 \sum_{j \in T} \sum_{i \in S \setminus \{j\}} a_{j,i} .$$

Lemma 11 in [4] shows that $\sum_{i \in S \setminus \{j\}} a_{j,i} \leq C$ for some constant C because $S \setminus \{j\}$ is a feasible set. This gives us

$$\sum_{i \in N} u_i(b'_i, b_{-i}) = \sum_{i \in S} u_i(b'_i, b_{-i}) \geq |S| - 2 \sum_{j \in T} C = |S| - 2C \sum_{i \in N} p_i(b) .$$

□

By applying our composition theorems, we obtain a constant price of anarchy for oblivious learning in this game even when we have multiple channels with fully independent or everybody-or-nobody availability. This simplifies and generalizes an approach based on sleeping expert learning in [11]. Furthermore, our analysis can also be conducted similarly for other interference models with a bounded independence condition, see [11,12] for a discussion.

B Extension to Changing Unit-Demand Functions

We now consider a case in which valuations change over time rather than the supply. In particular, we consider a unit-demand setting, i.e., there are values $v_{i,j}$ such that $v_i(S) = \max_{j \in S} v_{i,j}$. We assume that each of the $v_{i,j}$ is an independent random variable in which constantly many outcomes have a positive probability. So, for a fixed player i , the valuation is defined such that for $k = 1, \dots, K$ we let $v_{i,j} = v_{i,j}^{(k)}$ with probability $q_{i,j}^{(k)}$, $\sum_{k=1}^K q_{i,j}^{(k)} = 1$. Without loss of generality, let $v_{i,j}^{(1)} \geq v_{i,j}^{(2)} \geq \dots \geq v_{i,j}^{(K)}$.

To apply availability-oblivious learning, player i now makes K copies of each item j . The k th copy of item j has value $v_{i,j}^{(k)}$, and it is available whenever

$v_{i,j} \geq v_{i,j}^{(k)}$. Note that, when restricting the consideration to only the most valuable item, we can equivalently assume that availabilities of items are drawn independently with probability $q_{i,j}^{(k)} / \sum_{k'=k}^K q_{i,j}^{(k')}$ for the k th copy of item j .

By the same argument as in Section 3, we then have

$$\begin{aligned} \sum_i \mathbf{E}[u_i(b'_i, b_{-i})] &= \sum_i \mathbf{E}[v_i(f(b'_i, b_{-i}))] - \mathbf{E}[p_i(b'_i, b_{-i})] \\ &\geq \left(1 - \frac{1}{e}\right) \sum_i \mathbf{E}[v_i(f(b_i^{A_i}, b_{-i}))] - \mathbf{E}[p_i(b_i^{A_i}, b_{-i})] , \end{aligned}$$

when comparing the availability-oblivious deviation b'_i with the availability-aware ones $b_i^{A_i}$.

Therefore, if each single mechanism is weakly (λ, μ_1, μ_2) -smooth, the price of anarchy for oblivious learning is at most $e/(e-1) \cdot (\mu_2 + \max(1, \mu_1))/\lambda$.

C Missing Proofs

C.1 Proof of Lemma 1

Without loss of generality, let $v(x^1) \geq v(x^2) \geq \dots \geq v(x^k)$. For each component $i \in [m]$, let $J_i \in [k]$ be the random variable of the index of the vector from which y_i was taken.

Let z be defined by

$$z_i = \begin{cases} \perp_i & \text{if } J_i = 1 \\ y_i & \text{otherwise} \end{cases}$$

If $J_i \neq 1$, we have

$$v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m) = 0 .$$

Otherwise, if $J_i = 1$, we have

$$\begin{aligned} &v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m) \\ &\geq v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, y_i, z_{i+1}, \dots, z_m) - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m) \\ &= v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m). \end{aligned}$$

That is, in combination, we get

$$\begin{aligned} &\mathbf{E}[v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m)] \\ &\geq \alpha_1 \mathbf{E}[v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) \\ &\quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m) \mid J_i = 1] \\ &= \alpha_1 \mathbf{E}[v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) \\ &\quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m)] , \end{aligned}$$

where the last step uses the independence of the components.

Note that, by diminishing marginal returns, we have

$$\begin{aligned} & v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1, z_{i+1}, \dots, z_m) - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, \perp_i, z_{i+1}, \dots, z_m) \\ & \geq v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1 \vee z_i, z_{i+1}, \dots, z_m) \\ & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, z_i, z_{i+1}, \dots, z_m). \end{aligned}$$

Applying furthermore $x_i^1 \vee z_i = x_i^1 \vee y_i$ and taking the expectation, we get

$$\begin{aligned} & \mathbf{E}[v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m)] \\ & \geq \alpha_1 \mathbf{E}[v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, x_i^1 \vee y_i, z_{i+1}, \dots, z_m) \\ & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, z_i, z_{i+1}, \dots, z_m)] . \end{aligned}$$

Overall, we get

$$\begin{aligned} \mathbf{E}[v(y)] &= \mathbf{E}\left[v(z) + \sum_{i=1}^m v(y_1, \dots, y_i, z_{i+1}, \dots, z_m) - v(y_1, \dots, y_{i-1}, z_i, \dots, z_m)\right] \\ &\geq \mathbf{E}[v(z)] + \sum_{i=1}^m \alpha_1 \mathbf{E}[v(x_1^1 \vee y_1, \dots, x_i^1 \vee y_i, z_{i+1}, \dots, z_m) \\ & \quad - v(x_1^1 \vee y_1, \dots, x_{i-1}^1 \vee y_{i-1}, z_i, \dots, z_m)] \\ &= \mathbf{E}[v(z)] + \alpha_1 (\mathbf{E}[v(x^1 \vee y)] - \mathbf{E}[v(z)]) \\ &\geq (1 - \alpha_1) \mathbf{E}[v(z)] + \alpha_1 v(x^1) . \end{aligned}$$

By applying this argument inductively, we also get

$$\mathbf{E}[v(y)] \geq \sum_{i=1}^k \alpha_i v(x^1) \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) .$$

As we assumed $v(x^1) \geq v(x^2) \geq \dots \geq v(x^k)$, the FKG inequality gives us

$$\sum_{i=1}^k \alpha_i v(x^1) \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) \geq \left(\sum_{i=1}^k \alpha_i \prod_{i'=1}^{i-1} (1 - \alpha_{i'}) \right) \left(\sum_{i=1}^k \alpha_i v(x^1) \right) .$$

This implies

$$\mathbf{E}[v(y)] \geq \left(\frac{1}{k} \sum_{i=1}^k \left(1 - \frac{1}{k}\right)^{i-1} \right) \left(\sum_{i=1}^k \alpha_i v(x^1) \right) \geq \left(1 - \frac{1}{e}\right) \sum_{i=1}^k \alpha_i v(x^1) . \square$$

C.2 Proof of Lemma 2

Let $y_j^i = f_j(b'_{i,j}, b_{-i}) \wedge z_j^i$ and $\hat{z}_j^i = z_j^i$ if j is available and $\hat{z}_j^i = \perp_j$ otherwise. Notice that $f_j(b'_{i,j}, b_{-i}) \wedge \hat{z}_j^i = f_j(b'_{i,j}, b_{-i}) \wedge z_j^i$. This is because $\hat{z}_j^i = z_j^i$ when j is available and $f_j(b'_{i,j}, b_{-i}) = \perp_j$ when j is not available.

By monotonicity, we have $v_i(f(b'_i, b_{-i})) \geq v_i(f(b'_i, b_{-i}) \wedge z^i) = v_i(y^i)$. Furthermore, we can decompose $v_i(y^i)$ into a telescoping sum by

$$v_i(y^i) = \sum_j v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m) .$$

Next, we bound each of these terms independently using diminishing marginal returns multiple times

$$\begin{aligned} & v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m) \\ & \geq v_i(\hat{z}_1^i, \dots, \hat{z}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i, \dots, \hat{z}_{j-1}^i, \perp_j, \dots, \perp_m) \\ & \geq v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\ & = v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\ & \quad - (v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\ & \quad \quad \quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m)) \\ & \geq v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\ & \quad - (v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m)) \\ & = v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) \\ & \quad - (v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)) \\ & \quad + v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m) . \end{aligned}$$

That is, by linearity of expectation, we have

$$\begin{aligned} & \mathbf{E} [v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m)] \geq \\ & \underbrace{\mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]}_{\text{part 1}} \\ & \quad - \underbrace{\mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]}_{\text{part 2}} \\ & \quad + \underbrace{\mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]}_{\text{part 3}} . \end{aligned}$$

To simplify part 2, we use the fact that $\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i$ and \hat{z}_j^i are independent. Therefore, we have

$$\begin{aligned} & \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\ & = \sum_{\ell} \frac{q_j r_j^\ell}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, x_j^\ell, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\ & = \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] . \end{aligned}$$

For the same reason, we can also bound part 1 by using

$$\mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]$$

$$\begin{aligned}
&= \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) \\
&\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m)] \\
&\quad + \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
&\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&\leq \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) \\
&\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&\quad + \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
&\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&= (\alpha + 1) \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) \\
&\quad - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)],
\end{aligned}$$

which implies

$$\begin{aligned}
&\mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \hat{z}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&\geq \frac{1}{\alpha + 1} \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)]
\end{aligned}$$

Finally, part 3 is precisely the definition of $\mathbf{E} [w_{i,j}^i(y_j^i)]$. Therefore, in combination, we get

$$\begin{aligned}
&\mathbf{E} [v_i(y_1^i, \dots, y_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(y_1^i, \dots, y_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&\geq \frac{1}{\alpha + 1} \mathbf{E} [v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_j^i \vee \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\hat{z}_1^i \vee \tilde{t}_1^i, \dots, \hat{z}_{j-1}^i \vee \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&\quad - \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^i, \dots, \tilde{t}_j^i, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^i, \dots, \tilde{t}_{j-1}^i, \perp_j, \dots, \perp_m)] \\
&\quad + \mathbf{E} [w_{i,j}^i(y_j^i)]
\end{aligned}$$

Taking the sum over all j , we get two telescoping sums, which simplify to $\mathbf{E} [v_i(\hat{z}^i \vee \tilde{t}^i)]$ (part 1) and $\mathbf{E} [v_i(\tilde{t}^i)]$ (part 2). This gives us

$$\begin{aligned}
\mathbf{E} [v_i(f(b'_i, b_{-i}))] &\geq \mathbf{E} [v_i(y^i)] \geq \frac{1}{\alpha + 1} \mathbf{E} [v_i(\hat{z}^i \vee \tilde{t}^i)] - \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^i)] + \sum_j \mathbf{E} [w_{i,j}^i(y_j^i)] \\
&\geq \sum_j \mathbf{E} [w_{i,j}^i(y_j^i)] - \frac{1}{\alpha(\alpha + 1)} \mathbf{E} [v_i(\tilde{t}^i)] = \sum_j \mathbf{E} [w_{i,j}^i(f(b'_i, b_{-i}))] - \frac{1}{\alpha(\alpha + 1)} \mathbf{E} [v_i(\tilde{t}^i)] .
\end{aligned}$$

□

C.3 Proof of Lemma 3

Note that functions $w_{i,j}^{i'}$ are identically distributed for different i' and independent of any availabilities. Therefore, we have

$$\mathbf{E} [w_{i,j}^{i'}(f_j(b'_{i,j}, b_{-i}) \mid A_j = 1)] = \mathbf{E} [w_{i,j}^1(f_j(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i})) \mid A_j = 1] \quad (3)$$

and

$$\mathbf{E} [p_{i,j}(b'_{i,j}, b_{-i}) \mid A_j = 1] = \mathbf{E} [p_{i,j}(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i}) \mid A_j = 1] \quad (4)$$

Next, we apply the smoothness of each separate mechanism. Let us first assume mechanism j is available and z^1 and \tilde{t}^1 are fixed arbitrarily. This also fixes the functions $w_{1,j}^1, \dots, w_{n,j}^1$. We pretend these are the actual valuation functions. Then smoothness gives us

$$\begin{aligned} & \sum_i w_{i,j}^1(f_j(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i})) - p_{i,j}(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i}) \\ & \geq \lambda \left(\max_{y \in \Omega_j} \sum_i w_{i,j}^1(y) \right) - \mu_1 \sum_i p_{i,j}(b) - \mu_2 \sum_i h_{i,j}(b_i, f(b)) \\ & \geq \lambda \left(\sum_i w_{i,j}^1(z_j^1) \right) - \mu_1 \sum_i p_{i,j}(b) - \mu_2 \sum_i h_{i,j}(b_i, f(b)) \quad . \end{aligned}$$

Taking the expectation over z^1 and \tilde{t}^1 , we can combine this bound with (3) and (4) to get

$$\begin{aligned} & \sum_i \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i}) \mid A_j = 1] \\ & = \sum_i \mathbf{E} [w_{i,j}^1(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i})) - p_{i,j}(b_{i,j}^*(w_{1,j}^1, \dots, w_{n,j}^1), b_{-i}) \mid A_j = 1] \\ & \geq \mathbf{E} \left[\lambda \sum_i w_{i,j}^1(z_j^1) - \mu_1 \sum_i p_{i,j}(b) - \mu_2 \sum_i h_{i,j}(b_i, f(b)) \mid A_j = 1 \right] \quad . \end{aligned}$$

If j is not available, then naturally $f_j(b'_{i,j}, b_{-i}) = \perp_j$ and $p_{i,j}(b'_{i,j}, b_{-i}) = p_{i,j}(b) = 0$. By definition, however, $w_{i,j}^1(z_j^1)$ is independent of the fact whether j is available or not. Therefore, we have

$$\begin{aligned} & \sum_i \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq q_j \mathbf{E} \left[\lambda \sum_i w_{i,j}^1(z_j^1) \mid A_j = 1 \right] - q_j \mathbf{E} \left[\mu_1 \sum_i p_{i,j}(b) + \mu_2 \sum_i h_{i,j}(b_i, f(b)) \mid A_j = 1 \right] \\ & = q_j \mathbf{E} \left[\lambda \sum_i w_{i,j}^1(z_j^1) \right] - \mathbf{E} \left[\mu_1 \sum_i p_{i,j}(b) \right] - \mathbf{E} \left[\mu_2 \sum_i h_{i,j}(b_i, f(b)) \right] \quad . \end{aligned}$$

We can take the sum over all j to get

$$\begin{aligned} & \sum_i \sum_j \mathbf{E} [w_{i,j}^i(f_j(b'_{i,j}, b_{-i})) - p_{i,j}(b'_{i,j}, b_{-i})] \\ & \geq \lambda \sum_i \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] - \mu_1 \sum_i \mathbf{E} [p_i(b)] - \mu_2 \sum_i \mathbf{E} [h_i(b_i, f(b))] \quad . \square \quad (5) \end{aligned}$$

C.4 Proof of Lemma 4

By only plugging in the definitions of $w_{i,j}^1$, z^1 , and \tilde{t}^1 , we get

$$\begin{aligned}
& \sum_j q_j \mathbf{E} [w_{i,j}^1(z_j^1)] \\
&= \sum_j q_j \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, z_j^1, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\
&= \sum_j q_j \sum_{\ell} \frac{r_{\ell}}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, x_j^{\ell}, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\
&= \sum_j \frac{1}{\alpha} \sum_{\ell} q_j r_{\ell} \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, x_j^{\ell}, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\
&= \sum_j \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}_1^1, \dots, \tilde{t}_j^1, \perp_{j+1}, \dots, \perp_m) - v_i(\tilde{t}_1^1, \dots, \tilde{t}_{j-1}^1, \perp_j, \dots, \perp_m)] \\
&= \frac{1}{\alpha} \mathbf{E} [v_i(\tilde{t}^1)] \quad .
\end{aligned}$$

□

C.5 Proof of Theorem 3

Consider the following single-item first-price auction with a discrete bidding space. Each bidder has valuation 0 or 2 for the item, and the set of possible bids is $\{0, 1, 2\}$. The item is sold to one of the bidders with the maximal bid (arbitrary but fixed deterministic tie-breaking), and only this bidder pays his bid. If all bidders bid 0, the item is not given away. It is easy to see that if each bidder i in this auction deviates to half of his valuation, the auction becomes smooth with $\lambda = 1/2$, $\mu_1 = 1$ and $\mu_2 = 0$. Hence, the auction has a price of anarchy of at most 2.

We compose this auction for a set $[m]$ of $m = k^2$ items, for some integer $k > 0$, and every item is sold simultaneously via the first-price auction above. There is a single bidder, and he has an XOS valuation function v as follows. The items are grouped into k groups M_1, \dots, M_k of k items each. For a set of items $S \subseteq [m]$ we have $v(S) = \max_{\ell=1, \dots, k} \sum_{j=1}^m v_j^{\ell}$ with $v_j^{\ell} = 2$ if $j \in M_{\ell}$ and 0 otherwise, for $\ell = 1, \dots, k$. Consequently, $v(S) = 2 \max_{\ell=1, \dots, k} |S \cap M_{\ell}|$. We assume each item $j \in [m]$ is available independently with probability $q_j = 1/k$.

If the bidder can deviate depending on the set of available items, a social optimum b^* is obvious – he considers the group ℓ^* with the maximum number of available items and bids $b_j^* = 1$ for all $j \in M_{\ell^*}$ and 0 otherwise. This way he always obtains a set S of items that maximizes the valuation. Furthermore, this is also the best-response since he obtains the maximum valuation at minimum required payment, and the marginal utility of every obtained item is 1. In contrast, we show that every *oblivious* deterministic best-response bid b allows to recover at most a small fraction of the above described optimum. Thus, even

the price of anarchy for pure equilibria cannot be bounded by the smoothness guarantee.

In the optimum b^* , the bidder gets all available items from the group with the maximum number of available items. The number of available items in a group follows a binomial distribution $B(k, 1/k)$. This scenario is almost identical to throwing k balls uniformly at random into k bins and recording the maximum number of balls in any bin. Now in each bin k balls appear independently at random with probability $1/k$ each, and an almost identical analysis implies

$$\mathbf{E}[v(S(b^*))] = \Theta\left(\frac{\log k}{\log \log k}\right).$$

Now consider an oblivious deterministic best-response b . The valuation function v treats all items of a group in a symmetric way and all groups in a symmetric way. Let us denote by r_ℓ the number of items $j \in M_\ell$ with $b_j = 1$. For a fixed vector r , expected value and payments are the same no matter on which particular items $j \in M_\ell$ a bid $b_j = 1$ is placed. For any two groups M_ℓ and $M_{\ell'}$, the expected valuation and payments remain the same if we change the bid to have $b_j = 1$ for r_ℓ items in $M_{\ell'}$ and $r_{\ell'}$ items in M_ℓ . Moreover, the expected payment depends only on $\sum_\ell r_\ell$. Now suppose there are two groups M_ℓ and $M_{\ell'}$ such that $r_\ell, r_{\ell'} \leq k/2$. This bidding strategy is obviously dominated by any bid that bids 1 on $r_\ell + r_{\ell'}$ items in M_ℓ and none in $M_{\ell'}$. In conclusion, w.l.o.g. we can assume that $r_1 \geq r_2 \geq r_3 \geq \dots \geq r_k$ and there is k' such that $r_\ell \geq k/2$ for $\ell = 1, \dots, k' - 1$, $r_{k'-1} \geq r_{k'} \geq 0$ and $r_\ell = 0$ for $\ell = k' + 1, \dots, k$.

We show that every oblivious best-response b has $\mathbf{E}[v(S(b))] = O(1)$. Let $p(b)$ denote the total payments, X_j denote the event that item j is available, and $Y_\ell = \sum_{j \in M_\ell} X_j$ the number of available items in group M_ℓ , for all $\ell = 1, \dots, k$. Note that

$$\begin{aligned} \mathbf{E}[v(S(b)) - p(b)] &= \mathbf{E}\left[\max_{\ell=1, \dots, k'} \left(\sum_{j \in M_\ell, b_j=1} 2X_j\right) - \sum_{j \in [m], b_j=1} X_j\right] \\ &\leq 2\mathbf{E}\left[\max_{\ell=1, \dots, k'} Y_\ell\right] - \frac{k' - 1}{2}. \end{aligned}$$

Further, for any $d = 1, \dots, k$ we can use Chernoff bounds to see

$$\Pr\left[\max_{\ell=1, \dots, k'} Y_\ell \geq d\right] = 1 - (1 - \Pr[Y_1 \geq d])^{k'} \leq 1 - (1 - e^{-d-1}/d^d)^{k'} \leq \min\{1, k' e^{-d-1}/d^d\}.$$

Hence,

$$\mathbf{E}\left[\max_{\ell=1, \dots, k'} Y_\ell\right] = \sum_{d=1}^k \Pr\left[\max_{\ell=1, \dots, k'} Y_\ell \geq d\right] \leq \sum_{d=1}^k \min\{1, (k'/e) \cdot (e/d)^d\} \leq \frac{3 \log k'}{\log \log k'} + \frac{1}{ek'}.$$

Thus, $\mathbf{E}[v(S(b)) - p(b)] < (6 \log k')/(\log \log k') + 6/(ek') - (k' - 1)/2$, which is positive only for $k' \leq 34$. Every bid b with $k' \geq 35$ is dominated by b' with $b'_j = 0$ for all $j \in [m]$. Hence, for a best-response b we have $\mathbf{E}[v(S(b))] < 17$. This proves the theorem. \square