The Limits of Posted Prices in Random-Order Combinatorial Auctions

Frederick V. Qiu Aviad Rubinstein

November 3, 2026

Abstract

Combinatorial Auctions are a cornerstone problem in Algorithmic Game Theory: we want to allocate m items to n bidders with combinatorial (more specifically, XOS) valuations. The goal is to (approximately) maximize the social welfare, i.e. the sum of bidders' values for their allocations. In this paper we study the case where bidders arrive online in random order, and the seller needs to find item prices that balance between affordability for the current bidder and saving enough value for future bidders. A long line of works has given improving approximation guarantees: $O(\log^2 m)$ [DNS06], $O(\log m \log\log m)$ [Dob07], $O(\log m)$ [KV12], $O(\sqrt{\log m})$ [Dob16], $O((\log\log m)^3)$ [AS19], $O((\log\log m)^2)$ [AKS21], and most recently $O(\log\log m)$ [RS27]. Can this last $\log\log m$ factor be shaved to finally obtain a constant factor approximation?

We prove a strong negative answer: no truthful posted price mechanism that allocates in online random order can obtain a constant factor approximation. Our lower bound holds even for sellers that are computationally unbounded and have a full view of previous bidders' valuations.

Disclosure of helpful artificial and human intelligence. The proof of Theorem 2.1 was primarily human. The authors proved an $\Omega(\log\log\log m)$ lower bound, then supplied GPT 6 Astra with the idea of a role vector to achieve the final $\Omega(\log\log m)$ bound. GPT gave the idea to leverage utility-maximization in a more structured manner. The proof of Theorem 1.1 is based on the authors' ideas, but GPT was helpful for efficiently working out the details of the construction. In particular, the idea to independently sample the auxiliary pieces of the construction, embed the core in the intersection, and then leverage the independence of the auxiliary pieces for the probabilistic argument was a key implementation detail supplied by GPT.

More importantly, this paper would not have been possible without the human intelligence of Shahar Dobzinski, Sahil Singla, Jan Vondrak, and Mingwei Yang; in particular, it was Shahar who conjectured the lower bound against sellers that have full view of previous bidders' valuations.

1 Introduction

In a combinatorial auction, there is a set of items $M := [m]$ that we wish to allocate among n bidders. Each bidder $i \in [n]$ holds a private combinatorial valuation function $v_i : 2^M \to \mathbb{R}_+$; specifically, in this paper we focus on XOS valuations,[1] i.e. functions of the form $v_i(S) = \max_{\boldsymbol{u} \in \mathcal{U}} \sum_{x \in S} u_x$. An *allocation* is an n-tuple of pairwise disjoint subsets of M. The objective is to find an allocation $(S_1, \ldots, S_n)$ that approximately maximizes the welfare $\sum_{i \in [n]} v_i(S_i)$.

Since each bidder i only knows their own valuation v_i, two challenges arise:

1. The bidders must communicate in order to find a high welfare allocation.
2. A bidder may attempt to strategically manipulate the protocol by lying about their valuation, reducing overall welfare in order to achieve higher personal utility.

On one hand, ignoring incentives, [KRTV13, DKL+24] give $\Theta(1)$-approximation algorithms that process bidders in online random order and can be implemented with only polynomial communication. However, these algorithms do not incentivize truthful behavior. On the other hand, the VCG mechanism [Vic61, Cla71, Gro73] shows how to design a truthful mechanism given a welfare-maximizing algorithm. However, welfare-maximization (and hence the VCG auction) requires interactive, exponential communication. A central problem at the intersection of Economics and Computing is understanding whether it is possible to achieve communication-efficient, truthful mechanisms with good approximation guarantees.

One important class of mechanisms is *posted price* mechanisms. The mechanism visits bidders one at a time, and for each bidder i, it simply posts an item price vector $\boldsymbol{p}$ on the unallocated items. The bidder then declares their utility-maximizing set $S^* = \operatorname{argmax}_{S \subseteq M}\{v_i(S) - \boldsymbol{p}(S)\}$, and is allocated S^*. The mechanism may use information from previous visits to set prices for the next bidder, but the mechanism is truthful because no bidder can influence the prices they face. Ever since their introduction to the combinatorial auctions literature, posted price mechanisms have been the state-of-the-art polynomial-communication truthful mechanisms. There is a long line of work on posted price mechanisms [DNS06, Dob07, KV12, Dob16, AS19, AKS21], with the recent work of [RS27] giving an $O(\log \log m)$-approximation for XOS valuations.

One can also consider posted prices in an algorithmic setting. [DKL+24] give an $O(1)$-approximation for XOS valuations using a posted-price-like algorithm: every bidder faces a pricing over *all* items (including previously allocated ones), and chooses their favorite bundle. They are then allocated the available subset of their bundle. Hence, this algorithm is "almost" truthful, but not quite.

It is therefore more than natural to hope that we can shave the last $\log \log m$ factor and obtain a truthful posted price mechanism achieving a constant factor approximation. Our main result proves unconditionally that this is unfortunately not the case:

Theorem 1.1. *Posted price mechanisms cannot guarantee an $O(1)$-approximation of the welfare in Combinatorial Auctions with XOS valuations.*

[1] To avoid trivial pathological examples, the literature has focused on obtaining results for classes of valuation functions. We focus on the well-studied class of *XOS valuations* (sometimes also "fractionally subadditive"), which sit between submodular and subadditive. XOS has emerged as arguably the cleanest abstraction of "complement free" in this line of work.

1.1 Related Work

Truthful Combinatorial Auctions. As we have previously mentioned, there is a long line of work on truthful mechanisms using posted price and price learning techniques for combinatorial auctions for complement-free valuations [DNS06, Dob07, KV12, Dob16, AS19, AKS21, RS27]. Our work is the first to give an impossibility result for truthful posted price mechanisms for complement-free valuations.

Besides posted prices in the online setting, there is a rich body of literature on truthful mechanisms for combinatorial auctions in the polynomial communication [HKMT04, DNS10, QW24, AKSW20, RTWZ24] and polynomial computation [Dob11, DV11, DV12b, DV12a, DV16] settings. Our impossibility result stands in contrast to these settings, in that the mechanism is allowed unbounded computation and is given full access to the valuations of previously visited bidders. Nevertheless our result informs this line of research by ruling out an entire leading approach towards constant approximation guarantees: posted price mechanisms.

Online Combinatorial Auctions. A related line of work studies online combinatorial auctions, where bidders must be served irrevocably upon arrival. For complement-free valuations, approximation algorithms have been developed in the secretary setting [AKS21, RS27], the prophet inequality setting [FGL15, DFKL17, DKL20], and the prophet secretary setting [EHKS18]. More recent work investigates how little distributional information suffices, obtaining prophet inequalities using only a small number of samples [DKL$^+$24, CD26, RS27]. While these works mostly focus on the posted price and price learning paradigms, there are also works considering online algorithms that do not fall in the category of item pricing [KRTV13, CC23, MNP$^+$26].

2 Technical Overview

Single-Query Lower Bound. Two common communication models in combinatorial auctions are the *value query* and *demand query* models. For a valuation $v : 2^M \to \mathbb{R}_+$, a value oracle takes as input a set $S \subseteq M$ and outputs $v(S)$. A demand oracle takes as input a price vector $\boldsymbol{p} \in \mathbb{R}_+^m$ and outputs $\operatorname{argmax}_{S \subseteq M}\{v_i(S) - \boldsymbol{p}(S)\}$. Note that a demand oracle is precisely what is needed to implement a posted price mechanism.

An XOS valuation is one which can be written as a maximum of additive *clauses*. Therefore, a natural query model is the *clause query* model. Let $\mathcal{U} \subseteq \mathbb{R}_+^m$, and let $v : 2^M \to \mathbb{R}_+$ be an XOS valuation defined $v(S) = \max_{\boldsymbol{u} \in \mathcal{U}} \boldsymbol{u}(S)$. Then a clause oracle for $v, \mathcal{U}$ takes as input a price vector $\boldsymbol{p} \in \mathbb{R}_+^m$, and outputs both the utility-maximizing set and its utility-maximizing clause.

We observe that both the state-of-the-art mechanisms of [AKS21, RS27], and the non-truthful algorithm of [DKL$^+$24] can be implemented using a single clause query per bidder. For this class of mechanisms, we obtain a simpler and stronger lower bound, matching [RS27] up to constants:

Theorem 2.1. *Single-query posted price mechanisms cannot guarantee an $o(\log\log m)$-approximation of the welfare in Combinatorial Auctions with XOS valuations.*

We begin with an overview of this simpler lower bound, which is both interesting on its own and serves as a technical stepping stone towards our general result.

Let the number of bidders $n = \Theta(\log\log m)$. Our construction partitions the items into n blocks $Z_1, \ldots, Z_n$. Bidder i will be the *owner* of block i, and have high value for items in Z_i, and low

value for items in Z_j for $j \neq i$. In order to obtain high welfare, the mechanism must allocate many items in Z_i to bidder i. The key idea is that the mechanism cannot distinguish which bidders own which items, and therefore, if the mechanism posts prices which allow bidder i to purchase Z_i, these prices also allow bidder i to *steal* items owned by other bidders, *destroying* owner value while gaining negligible value for themselves. The question then becomes how to make a bidder with "low" value steal "high" value items; after all, the mechanism merely needs to set a "medium" price to ensure each item goes to the correct bidder.

First, the construction muddies the meanings of low and high. In particular, suppose each block is partitioned into $\kappa = 2^{\Theta(n)}$ *ranks*. For some sufficiently large base β, let there be $\beta^{\kappa-k}$ items of rank k, and let the base value of each rank k item be β^k. The total value of each rank is therefore constant. If bidder i owns an item x, then their value for x is $\beta^{1/2}$ times the base value. Now, even though bidder j has vastly more value for Z_j than bidder i, there is no single price which allows bidder j to capture this value without also allowing bidder i to destroy it.

However, the above intuition only applies if the mechanism cannot determine the rank of the items; once the mechanism knows that an item x is rank k, it can simply price x at $\beta^k + 1$.

Information Hiding. The key is to equip bidder i with a large set of clauses, which encode bidder i with enough information to steal the items from Z_j, but not enough information for the mechanism to learn the exact rank of the items. For example, suppose the mechanism knows nothing about the items, and uniformly posts a price of $\beta^k + 1$. At these prices, the owner can afford the items in ranks $[k, \kappa]$, and a non-owner can afford the items in ranks $[k+1, \kappa]$. We would like for bidder i to steal every item in ranks $[k+1, \kappa]$, but we don't want bidder i's clause to reveal too much information. Then the ideal clause for bidder i to use is one that gives the items of ranks $[1, k]$ value β^1, and the items of ranks $[k+1, \kappa]$ value β^{k+1}. Under these values, bidder i is able to steal the most they could possibly steal, while revealing the least amount of information possible: whether the rank of each item was above or below k.

Since $\kappa = 2^{\Theta(n)}$, and the mechanism only gets a single query per bidder, the mechanism can never learn the item ranks with enough precision to post prices which allow owners to purchase their items but prevent non-owners from stealing those same items.

The requirement $n = \Theta(\log \log m)$ is because the construction uses $\approx \beta^\kappa = 2^{2^{\Theta(n)}}$ items.

Utility Maximization. The last piece of the puzzle is figuring out how to actually make it so the clause which steals and hides information is actually the utility-maximizing clause. To do this, suppose that we can give each clause a "bias," meaning a fixed bonus awarded whenever the clause is chosen. We can essentially do this WLOG, see Proposition 3.8.

Then in some clauses, the bias will be large but the item values will be lower, while in other clauses, the bias will be small but the item values will be higher. The effect is that a clause with higher item value yields more utility when those items are affordable; otherwise, having a higher bias is preferred. In particular, by having clauses increase the item values linearly while decreasing the bias quadratically, we can make utility-maximization choose a clause which semantically optimizes an affordability statistic.

For example, suppose we had a set S of items $\{1, \dots, 1000\}$ receiving prices from the same distribution, and we wanted to estimate the fraction of items in S whose price is below 100. Then define 1001 clauses parameterized by an integer $q \in \{0, 1, \dots, 1000\}$. For each clause q, the value of item $x \in [1000]$ is $99 + \mathbb{1}(x \leq q)$, and the bias is $-q^2/2000$. For simplicity, suppose no price

lies in $[99, 100]$. Then a quick calculation shows that the utility-maximizing clause in expectation is $q = 1000p$, where p is the probability an item has price below 100. A construction in its entirety might embed a small set of these "detector" items into a larger population of items, then tie the measured statistic q to side effects on the normal items to create sophisticated bidder responses to utility-maximization.

Full-Revelation Lower Bound. We now explain how to generalize the ideas from the single-query lower bound to the following fully general result.

Theorem 2.2. *Posted price mechanisms cannot guarantee an $o(\log\log\log m)$-approximation of the welfare in Combinatorial Auctions with XOS valuations, even with unbounded computation and full access to the valuations of previously visited bidders.*

The main new idea is to imagine a huge pool of $\eta = m^{2n}$ *names*. Draw a random injection $\tau : [n] \to [\eta]$ which assigns each bidder identity $1, \ldots, n$ a name $\tau(i)$. Independently for each name-set $L \subseteq [\eta]$ of size at most n, sample the nested sets $Z_{L,1} \supseteq \cdots \supseteq Z_{L,\kappa}$. The interpretation is that $Z_{L,k}$ is where the name-set L "believes" the items of rank $\geq k$ lie. Each name-set L corresponds to a clause, which values items according to which set $Z_{L,k} \setminus Z_{L,k+1}$ the item belongs.

The true rank k items Z_k lie in a set where all the real names agree on, i.e.,

$$Z_k \subseteq \bigcap_{L \subseteq \tau([n])} Z_{L,k} \ .$$

Finally, give every bidder i a *name list* $N_i \subseteq [\eta]$. N_i contains every real name $\tau([n])$, and contains each other name independently w.p. only $1/m$. Bidder i has exactly the clauses corresponding to name-sets $L \subseteq N_i$.

Now, grant the mechanism full knowledge of the nested sets $Z_{L,1} \supseteq \cdots \supseteq Z_{L,\kappa}$ for every name-set L. Then there are two things that a mechanism can try to do: it can learn where the valuable items are so it can protect them, or it can manipulate utility-maximization to make the next bidder i steal some random low value items instead of the valuable items. The first would require the mechanism to guess the real names of the unvisited bidders, which is unlikely because the full name pool is massive. The second is also unlikely: suppose that the mechanism tries to manipulate the bidder i into using the clause corresponding to L. But bidder i only has clauses which have name-sets contained in N_i, and the probability that $L \subseteq N_i$ is vanishing.

At a high level, the name pool is what allows us to get around full-revelation: even though the mechanism can see everything, it can't distinguish which pieces of information are useful. It also can't just choose a price that makes use of lots of information, because the information is independently sampled and pulls in random directions. The same concepts introduced in the single-query exposition of stealing and information hiding via utility-maximization are still used in the full-revelation lower bound, but we defer their details to the proof.

Note that this exposition should not be taken as a literal description of the construction, but a simplification of main ideas needed to obtain a lower bound in the full-revelation model.

3 Definitions, Notation, and Simplifications

We first formally state the model.

Definition 3.1 (XOS)**.** A valuation $v : 2^M \to \mathbb{R}_+$ is *XOS* if there exist *clauses* $\mathcal{U} \subseteq \mathbb{R}_+^m$ such that $v(S) = \sup_{\boldsymbol{u} \in \mathcal{U}} \boldsymbol{u}(S)$ for all $S \subseteq M$. We call any such $\mathcal{U}$ a *representation* of v. We assume $|\mathcal{U}| < \infty$.

Definition 3.2 (Demand Oracle)**.** A *demand oracle* for a valuation $v : 2^M \to \mathbb{R}_+$ takes a price vector $\boldsymbol{p} \in [0, \infty]^m$ as input, and outputs the lexicographically smallest (according to the string given by the indicator vector) utility-maximizing set S.

Definition 3.3 (XOS Clause Oracle)**.** A *clause oracle* for a fixed representation $\mathcal{U} \subseteq \mathbb{R}_+^m$ of an XOS valuation $v : 2^M \to \mathbb{R}_+$ takes a price vector $\boldsymbol{p} \in [0, \infty]^m$ as input. The oracle outputs the lexicographically smallest utility-maximizing set S and the lexicographically smallest clause $\boldsymbol{u} \in \mathcal{U}$ such that $v(S) = \boldsymbol{u}(S)$.

Definition 3.4 (Single-Query Model)**.** An *online single-query* mechanism visits each bidder one at a time (in a possibly adaptive order). When the mechanism visits a bidder, it declares whether the bidder is a *sample* bidder or not, then makes a query to that bidder's XOS clause oracle, where previously allocated items are required to have price ∞. If the bidder is a sample bidder, they are allocated nothing. Otherwise, they are allocated their reported set.

Definition 3.5 (Full-Revelation Model)**.** An *online full-revelation* (OFR) mechanism visits each bidder one at a time (in a possibly adaptive order). When the mechanism visits a bidder, it makes a query to that bidder's demand oracle, where previously allocated items are required to have price ∞. The bidder is allocated their reported set, and the mechanism learns the bidder's full valuation.

Definition 3.6 (Approximation Ratio)**.** A mechanism attains an *α-approximation* for a valuation class $\mathcal{V}$ if for all $v_1, \ldots, v_n \in \mathcal{V}$, the allocation $S_1, \ldots, S_n$ output by the mechanism satisfies

$$\mathbb{E}\left[\sum_{i \in [n]} v_i(S_i)\right] \quad \geq \quad \frac{1}{\alpha} \max_{(T_1, \ldots, T_n)} \sum_{i \in [n]} v_i(T_i) \ ,$$

where the expectation is taken over the randomness of the mechanism.

For every lower bound, we will define a distribution over instances such that there is a uniform bound on the expected welfare over all deterministic mechanisms. Therefore, it suffices to consider only deterministic mechanisms by an averaging argument.

Throughout the paper, we will use the following notational conventions. For a vector $\boldsymbol{u} \in \mathbb{R}^M$ and set $S \subseteq M$, we denote $\boldsymbol{u}(S) = \sum_{i \in S} u_i$. When convenient, we may denote a singleton set $\{x\}$ simply as x. All logarithms are base 2. $\mathbb{Z}_+$ and $\mathbb{R}_+$ include 0. We assume that all parameters are chosen such that all relevant quantities are integers (e.g., if we state that a set $S \subseteq M$ has size $m/10$, we implicitly assume m is divisible by 10).

3.1 Reduction to Biased XOS Valuations

An XOS valuation is a maximum over linear functions. In our lower bound constructions, it will be convenient to allow each of these linear functions to have a nonzero *bias*.

Definition 3.7 (Biased XOS Representation)**.** A *biased XOS representation* of a valuation $v : 2^M \to \mathbb{R}_+$ is a set of clauses $\mathcal{U} \subseteq \mathbb{R}_+^{m+1}$ such that $v(S) = \max_{\boldsymbol{u} \in \mathcal{U}} \{\boldsymbol{u}(S) + u_0 \, \mathbb{1}(S \neq \emptyset)\}$ for all $S \subseteq M$. Given a valuation and a biased XOS representation, a *biased clause oracle* takes a price vector $\boldsymbol{p} \in [0, \infty]^m$ as input, and outputs the lexicographically smallest utility-maximizing set S and a utility-maximizing clause $\boldsymbol{u} \in \mathcal{U}$ such that $v(S) = \boldsymbol{u}(S) + u_0 \, \mathbb{1}(S \neq \emptyset)$. The clause may be chosen using an arbitrary deterministic tie-breaking rule. If the utility-maximizing set is $\emptyset$, the oracle outputs the clause $\mathbf{0}$.

Proposition 3.8. *Let $\mathcal{U} \subseteq \mathbb{R}_+^{m+1}$ be a biased XOS representation of $v : 2^M \to \mathbb{R}_+$. Then:*

1. *v is XOS.*
2. *v has an XOS representation $\mathcal{U}' \subseteq \mathbb{R}_+^m$ such that the responses of the clause oracle given by $\mathcal{U}'$ are no more informative than the responses of some biased XOS clause oracle given by $\mathcal{U}$.*

Proof. For every $\boldsymbol{u} \in \mathcal{U}$ and every $z \in M$, define $\boldsymbol{w}^{(\boldsymbol{u},z)} \in \mathbb{R}_+^m$ by

$$w_x^{(\boldsymbol{u},z)} \quad = \quad \begin{cases} u_x & x < z \\ u_x + u_0 & x = z \\ 0 & x > z \end{cases} .$$

Define $\mathcal{U}' = \{\boldsymbol{w}^{(\boldsymbol{u},z)} : \boldsymbol{u} \in \mathcal{U},\, z \in M\} \cup \{\mathbf{0}\}$. For $\boldsymbol{u} \in \mathcal{U}$, we see that for non-empty $S \subseteq M$,

$$\boldsymbol{u}(S) + u_0 \quad = \quad \boldsymbol{w}^{(\boldsymbol{u},\max S)}(S) \quad = \quad \max_{z \in M} \boldsymbol{w}^{(\boldsymbol{u},z)}(S) \ ,$$

Therefore, taking maximums over $\mathcal{U}$ shows that $\mathcal{U}'$ is an XOS representation of v, which proves (1).

We now show that some biased clause oracle given by $\mathcal{U}$ can simulate the responses of the clause oracle given by $\mathcal{U}'$. Consider a query $\boldsymbol{p} \in [0,\infty]^m$, and let S be the output of the clause oracle given by $\mathcal{U}'$. If $S = \emptyset$, then $\mathbf{0}$ is the lexicographically smallest utility-maximizing clause.

Otherwise, the clause oracle outputs some clause $\boldsymbol{w}^{(\boldsymbol{u},z)}$. If $z < \max S$, then the item $\max S$ yields zero utility and S is not the lexicographically smallest utility-maximizer. If $z > \max S$, then the clause $\boldsymbol{w}^{(\boldsymbol{u},\max S)}$ obtains the same utility while being lexicographically smaller.

Hence, after S and $\boldsymbol{u}$ are determined, the utility-maximizing clause is $\mathbf{0}$ if $S = \emptyset$ and $\boldsymbol{w}^{(\boldsymbol{u},\max S)}$ otherwise. Since S and $\boldsymbol{u}$ are determined in a fixed manner, there exists a biased clause oracle given by $\mathcal{U}$ which simulates the clause oracle given by $\mathcal{U}'$, which proves (2). □

Therefore, we may WLOG consider biased XOS representations and biased clause oracles.

3.2 Concentration Inequalities

We will colloquially use Chernoff and Hoeffding bounds without much formal bookkeeping. We will also use the following two concentration bounds, which we state explicitly. For a finite set S and a function f whose domain contains S, define

$$\mathrm{Avg}[f \mid S] \quad = \quad \frac{1}{|S|} \sum_{x \in S} f(x) \ ,$$

where $\mathrm{Avg}[f \mid \emptyset] = 0$. If the domain of f is S, then we simply write $\mathrm{Avg}[f]$.

Theorem 3.9 (Hoeffding for Sampling Without Replacement [BM15])**.** *Let Ω be a finite set, $s \in [|\Omega|]$, and $S \subseteq \Omega$ be a uniformly random set of size s. Then for any $f : \Omega \to [0,1]$ and $t > 0$,*

$$\Pr\left[\left|\mathrm{Avg}[f \mid S] - \mathrm{Avg}[f \mid \Omega]\right| \geq t\right] \quad \leq \quad 2e^{-2st^2} \ .$$

Corollary 3.10. *Let $S \subseteq \Omega$ be sampled by independently including each item $x \in \Omega$ w.p. q. Then for any $f : \Omega \to [0,1]$,*

$$\Pr\left[|S| \geq s,\ \left|\mathrm{Avg}[f \mid S] - \mathrm{Avg}[f \mid \Omega]\right| \geq t\right] \quad \leq \quad 2e^{-2st^2} \ .$$

Proof. Condition on $|S| = s'$, which allows us to invoke Theorem 3.9. Averaging over $s' \geq s$ completes the proof. □

3.3 Partitions

Definition 3.11 (Partitions)**.** Let $\kappa \in \mathbb{Z}_+$. For a partition $\mathcal{K}$ of $[\kappa]$, let $\mathcal{K}(k)$ denote the set $K \in \mathcal{K}$ such that $k \in K$.

Definition 3.12 (Geometric Partition)**.** For an interval $K \subseteq [\kappa]$, define $\mathcal{L}(K)$ to be the partition of K into intervals of doubling size (i.e., $1, 2, 4, 8, \dots$), starting from the right and possibly truncating the leftmost interval. For a partition $\mathcal{K}$ of $[\kappa]$ into intervals, let $\mathcal{L}(\mathcal{K}) = \bigcup_{K \in \mathcal{K}} \mathcal{L}(K)$. For example, $\mathcal{L}([30, 80]) = \{[30, 49], (49, 65], (65, 73], (73, 77], (77, 79], (79, 80]\}$.

Definition 3.13. For $s \in \mathbb{Z}_+$, define $\mathcal{K}_s = \mathcal{L}^s(\{[\kappa]\})$.

In our lower bounds, we will use $\mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_n$ as the rank information partition schedule. That is, after visiting the i^{th} bidder, the mechanism will not be able to distinguish the ranks of items to a finer granularity than $\mathcal{K}_i$.

The reason we choose the refinement given by $\mathcal{L}$ is as follows. If the mechanism cannot distinguish the ranks in the interval $K \subseteq [\kappa]$, and some fraction f of the value from K is affordable to the owner, we want a non-owner to be able to destroy a comparable fraction g of value. The obstacle is that the owner's value may be much higher than the non-owner's value, because the non-owner value gets rounded down to the bottom of its rank interval. However, notice that in the partition $\mathcal{L}(K)$, every pair of consecutive intervals $C, D \in \mathcal{L}(K)$ satisfies $|C| \le 2|D|$. Therefore, if the price distribution is uniform over the interval K, and the owner can obtain an f fraction of value from C, then the non-owner must be able to obtain a g fraction of value from D. Summing over the consecutive intervals in $\mathcal{L}(K)$ gives the desired relation. It also turns out the refinement $\mathcal{L}$ does not expose too much information. We prove these two facts below.

Lemma 3.14. *$|\mathcal{K}_s| \le 5^s \sqrt{\kappa}$ for $s \in \mathbb{Z}_+$.*

Proof. For a partition $\mathcal{K}$ of $[\kappa]$ into intervals, define $\Phi(\mathcal{K}) = \sum_{K \in \mathcal{K}} \sqrt{|K|}$. Then

$$\Phi(\mathcal{L}(\mathcal{K})) \quad \le \quad \sum_{K \in \mathcal{K}} \sum_{\ell=0}^{\lceil \log |K| \rceil} 2^{\ell/2} \quad \le \quad \sum_{K \in \mathcal{K}} \frac{2^{\lceil \log |K| \rceil / 2}}{1 - 1/\sqrt{2}} \quad \le \quad \sum_{K \in \mathcal{K}} 5\sqrt{|K|} \quad = \quad 5 \cdot \Phi(\mathcal{K}) \ .$$

Therefore,

$$|\mathcal{K}_s| \quad \le \quad \Phi(\mathcal{K}_s) \quad \le \quad 5^s \Phi(\mathcal{K}_0) \quad = \quad 5^s \sqrt{\kappa} \ . \qquad \square$$

Lemma 3.15. *Let $s \in \mathbb{Z}_+$, $f, g : [\kappa] \to [0, 1]$, and $a, b \ge 0$. Suppose that for every $K \in \mathcal{K}_s$ and every pair of consecutive intervals $C, D \in \mathcal{L}(K)$, $\operatorname{Avg}[f \mid C] \le a \operatorname{Avg}[g \mid D] + b$. Then*

$$\sum_{k \in [\kappa]} f(k) \quad \le \quad 2a \sum_{k \in [\kappa]} g(k) + b\kappa + 5^s \sqrt{\kappa} \ .$$

Proof. Fix $K \in \mathcal{K}_s$. For consecutive $C, D \in \mathcal{L}(K)$, we have $|C| \le 2|D|$. Therefore,

$$\sum_{k \in C} f(k) \quad = \quad |C| \operatorname{Avg}[f \mid C] \quad \le \quad a|C| \operatorname{Avg}[g \mid D] + b|C| \quad \le \quad 2a \sum_{k \in D} g(k) + b|C| \ .$$

Summing over consecutive pairs of $\mathcal{L}(K)$ yields every $k \in K$ on LHS except the rightmost element, and every $k \in K$ on RHS except the leftmost interval. But since $f(k) \le 1$, we have

$$\sum_{k \in K} f(k) \quad \le \quad 2a \sum_{k \in K} g(k) + b|K| + 1 \ .$$

Summing over $K \in \mathcal{K}_s$ and applying Lemma 3.14 completes the proof. $\square$

4 Single-Query Lower Bound

For sufficiently large n, let

- $\kappa = 2^{20n}$ be the number of *ranks*.
- $\beta = \kappa^\kappa$ and $\lambda = \beta^\kappa$.
- $\mathrm{val}(k) = \beta^k$ for $k \in [\kappa]$ be the *value* of items of rank k.
- $\mathrm{num}(k) = \lambda/\beta^k$ for $k \in [\kappa]$ be the (relative) *number* of items of rank k.

Let $R = [-\kappa, \kappa] \cap \mathbb{Z}$. We give each item a private label

$$(i, k, \boldsymbol{r}, \ell) \ , \qquad \text{where } \ i \in [n], \ \ k \in [\kappa], \ \ \boldsymbol{r} \in R^n, \ \ \ell \in [\kappa \,\mathrm{num}(k)] \ .$$

For each item, we call i the *owner*, k the *rank*, $\boldsymbol{r}$ the *role*, and ℓ the *layer*. Each label is assigned to λ^{10} items. We call the λ^{10} items associated with a single label a *fiber*. The number of items is

$$m \quad \coloneqq \quad n\kappa\lambda^{10}(2\kappa+1)^n \sum_{k\in[\kappa]} \mathrm{num}(k) \quad \le \quad \lambda^{12} \ ,$$

and $n = \Theta(\log\log\lambda) = \Theta(\log\log m)$. The private labels are randomly permuted onto the items.

Fix some $i \in [n]$. For $j \neq i$ and $\boldsymbol{r}_{-i} \in R^{n-1}$, call the set of labels $\{(j, k, (r_i, \boldsymbol{r}_{-i}), \ell) : k \in [\kappa], r_i \in R, \ell \in [\kappa\,\mathrm{num}(k)]\}$ an *i-slice*. Hence, the labels in an i-slice vary over their rank, i-role, and layer. Let Σ_i denote the set of i-slices, and note that $|\Sigma_i| = (n-1)(2\kappa+1)^{n-1} \le \beta^{1/8}$.

Define $\mathcal{A} = \{0, 1, \ldots, \kappa\}^{\kappa\times\kappa}$, and for a matrix $A \in \mathcal{A}$, define

$$s(A) \quad = \quad \min\{s \in \mathbb{Z}_+ : \mathcal{K}_s(k) = \mathcal{K}_s(r) \implies A_k = A_r \ \ \forall k, r \in [\kappa]\} \ .$$

In other words, $\mathcal{K}_{s(A)}$ is the coarsest partition for which differing rows in A belong to different intervals. Additionally define $\mathrm{rank}(k, A) = \min \mathcal{K}_{s(A)+1}(k)$ and $\mathrm{val}(k, A) = \mathrm{val}(\mathrm{rank}(k, A))$.

4.1 The Full Instance and Road Map

Let $\boldsymbol{A} \in \mathcal{A}^{\Sigma_i}$ denote a list of matrices. Define $\mathrm{Avg}[A]$ or $\mathrm{Avg}[A_\sigma]$ as the average over the entries of the list of matrices or matrix. Let $\boldsymbol{u}^{(i,\boldsymbol{A})} \in \mathbb{R}_+^{m+1}$, where for an item x with label $(i, k, \boldsymbol{r}, \ell)$,

$$u_x^{(i,\boldsymbol{A})} \quad = \quad \begin{cases} \beta^{1/2}\,\mathrm{val}(k) & \boldsymbol{r} = \boldsymbol{0} \\ 0 & \text{o.w.} \end{cases} \ ,$$

and for an item x with label $(j, k, \boldsymbol{r}, \ell)$ for $j \neq i$ and $\sigma = (j, \boldsymbol{r}_{-i})$,

$$u_x^{(i,\boldsymbol{A})} \quad = \quad \begin{cases} \mathrm{val}(k, A_\sigma) & r_i = 0 \\ \mathrm{val}(-r_i) & -r_i \in [k] \setminus \mathrm{rank}(k, A_\sigma) \\ \mathrm{val}(r_i) - 1 + \mathbb{1}(\ell \le A_{\sigma,k,r_i}) & r_i, \ell \in [\kappa] \\ 0 & \text{o.w.} \end{cases} \ .$$

Independently draw uniformly random offsets $\theta_1, \ldots, \theta_n \in [0, 1]$, and let

$$u_0^{(i,\boldsymbol{A})} \quad = \quad \lambda^{10}\Big(\lambda - \kappa|\Sigma_i|\,\mathrm{Avg}[\theta_i \boldsymbol{A} + \boldsymbol{A}^2/2]\Big) \ .$$

This value is positive because $\lambda \geq \kappa^3|\Sigma_i|$. The valuation $v_i : 2^M \to \mathbb{R}_+$ for bidder i is defined

$$v_i(S) \quad = \quad \max_{\boldsymbol{A} \in \mathcal{A}^{\Sigma_i}} \left\{ \boldsymbol{u}^{(i,\boldsymbol{A})}(S) + u_0^{(i,\boldsymbol{A})} \, \mathbb{1}(S \neq \emptyset) \right\} .$$

Observation 4.1. *WLOG, we may assume the mechanism visits bidders in the fixed order* $1, \ldots, n$.

Proof. Suppose the visited bidders are $I \subseteq [n]$. Choose a permutation $\pi : [n] \to [n]$ where $\pi(i) = i$ for all $i \in I$, and consider the transformed instance where we change the label of every item x from $(j, k, \boldsymbol{r}, \ell)$ to $(\pi(j), k, \pi(\boldsymbol{r}), \ell)$, and set the offset of bidder $\pi(i)$ to θ_i. This preserves the prior distribution. Further, because the mechanism does not observe the private item labels, the observed transcript up to this point is preserved, since the clauses of bidders $i \in I$ do not depend on owner labels $j \notin I$ or j-roles for $j \notin I$. Thus, conditioned on the observed history, the transformation does not change the joint distribution of the unvisited bidders' information, so no unvisited identity is a more advantageous next choice than another. □

Therefore, fix a deterministic mechanism which visits bidders in the fixed order $1, \ldots, n$.

Definition 4.2 (Indistinguishable)**.** Two items are *indistinguishable* at visit i when for each bidder $j \in [i-1]$, the reported value for those items is identical (under the biased clause oracle).

Definition 4.3 (Relevant)**.** An i-slice $(j, \boldsymbol{r}_{-i})$ is *i-relevant* if $j > i$ and $r_1 = \cdots = r_{i-1} = 0$. An item is i-relevant if its i-slice is i-relevant, or its owner is i and role is $\mathbf{0}$. Denote by $X_{i,k}$ the set of i-relevant items of rank k, and denote $X_{i,K} = \bigcup_{k \in K} X_{i,k}$.

Definition 4.4 (Core Items)**.** For $i \in [n]$ and $k \in [\kappa]$, let $Z_{i,k}$ contain all items with owner i, rank k, and role $\mathbf{0}$. Denote $Z_{i,K} = \bigcup_{k \in K} Z_{i,k}$. Note that $Z_{i,k} \subseteq X_{i,k}$.

We now describe the main idea of the construction.

1. Most of what we have discussed up to this point maps onto the intuition from the introduction. An item x has an owner i, which states who has high value for x. An item has a rank k, which controls the exponential value scale. And each rank k has layers $\ell \in [\kappa \, \mathrm{num}(k)]$, which controls the exponentially decreasing number of items at each rank. The main new construct is the role vector $\boldsymbol{r}$.
2. The core items are the true high value items. Recall that our goal is to hide information about the true ranks, by taking an interval of ranks $[a, b]$ and rounding down all the values to $\mathrm{val}(a)$. If we look at the non-owner definition for $u_x^{(i,\boldsymbol{A})}$, this is precisely the point of setting the value of i-role 0 items (which are a superset of the core items) to $\mathrm{val}(k, A_\sigma)$.
3. But just rounding down the value of items makes it so that hiding less information produces more utility. The role vector is designed to counter this: we can think of the items of rank k and i-role $-r_i \in [k]$ as having a default value of $\mathrm{val}(-r_i)$. Then when we want to round down the value of the i-role 0 items to $\mathrm{val}(k, A_\sigma)$, we can compensate by zeroing out the value of the items with i-role $\mathrm{rank}(k, A_\sigma)$.
4. Finally, observe that no valuation v_j with $j < i$ has coefficients depending on i-role. Therefore, the prices posted at visit i are equally distributed across i-role, and we can round down the value of i-role 0 items with no impact on expected utility.
5. Of course, after visiting bidder i, the mechanism will be able to distinguish items by i-role. But since only the items with role $\mathbf{0}$ end up ultimately mattering, the mechanism is free to discard

all items with nonzero i-role, as they are no longer relevant in later visits. The important thing is that by rounding down the value of the i-role 0 items, the full rank information of all the $(i+1)$-relevant items is still hidden.

The other question is how we make utility-maximization find the desired clause. The idea here is that the list of matrices $\boldsymbol{A}$ estimates, for each i-slice σ and rank k, the CDF $A_{\sigma,k} \in \{0, 1, \ldots, \kappa\}^\kappa$ of the price distribution on the rank k items in σ. In particular, $A_{\sigma,k,r}/\kappa$ is roughly the fraction of items with price less than $\mathrm{val}(r)$.

Ideally, when the mechanism cannot distinguish the items within a rank interval K, we want the clause to refine the interval into $\mathcal{L}(K)$, revealing just enough information to steal half the owner affordable items. As described above, the role vector allows us to arbitrarily round down the value of items without impacting utility. Therefore, we just need utility-maximization to actually learn what rank intervals the mechanism is unable to distinguish.

But now observe that if the mechanism cannot distinguish the ranks in $k \in K$, then the distribution of prices for ranks $k \in K$ must all be roughly equal. Therefore, there are three steps: first, we show that utility-maximization selects $A_{\sigma,k}$ which best reflects the true price distribution on rank k items in σ (this will be measured by the i-role $r_i \in [\kappa]$ items with value $\mathrm{val}(r_i) - 1 + \mathbb{1}(\ell \le A_{\sigma,k,r_i})$ together with the subtracted $\mathrm{Avg}[\boldsymbol{A}^2]$ term in the bias). Second, we show that by induction, the same $A_{\sigma,k}$ is chosen for all i-relevant σ. Third, $\mathcal{K}_{s(A_\sigma)}$ is precisely the lowest information partition among $\mathcal{K}_0, \mathcal{K}_1, \mathcal{K}_2, \ldots$ which captures the information reflected by the posted prices. Therefore, rounding the values of rank k items to $\mathrm{val}(k, A_\sigma)$ exactly rounds down the value of rank k items to the bottom of their interval in $\mathcal{L}(\mathcal{K}_{s(A_\sigma)})$.

4.2 Indistinguishable Items Have Similar Realized Prices

In this section, we formally show that the realized distribution of prices on indistinguishable items is close in a meaningful sense.

Definition 4.5 (Empirical CDF)**.** For a non-empty $S \subseteq M$ and prices $\boldsymbol{p} \in [0, \infty]^m$, define

$$F(q, S, \boldsymbol{p}) \quad = \quad \frac{|\{x \in S : p_x < q\}|}{|S|} \ .$$

In other words, $F(q, S, \boldsymbol{p})$ is the fraction of items of S whose price is below q.

Definition 4.6. Let $\mathcal{E}$ be the event that for all non-empty unions of fibers $S, T \subseteq M$ where the items in $S \cup T$ are indistinguishable, any prices $\boldsymbol{p} \in [0, \infty]^m$ posted by the mechanism satisfy

$$\max_{q \in \mathbb{R}_+} |F(q, S, \boldsymbol{p}) - F(q, T, \boldsymbol{p})| \quad \le \quad \lambda^{-4} \ .$$

Proposition 4.7. $\Pr[\mathcal{E}] \ge 1 - \lambda^{-2}$.

Proof. Fix any non-empty unions of fibers S, T. Condition on the assignment of private labels to the items outside $S \cup T$ and the offsets $\theta_1, \ldots, \theta_n$. This fixes $F(q, S \cup T, \boldsymbol{p})$ for all $q, \boldsymbol{p}$, and leaves the remaining private labels randomly permuted onto $S \cup T$.

Let $\mathcal{H}$ be a hypothetical history of reported clauses, and observe that the prices $\boldsymbol{p}_i(\mathcal{H})$ the mechanism posts at visit i is completely fixed by $\mathcal{H}$. Therefore, there are only $n(1 + (1+\kappa)^{\kappa^2 |\Sigma_i|}))^n \le 2^\lambda$ different possible prices the mechanism could post (conditioned on the labels outside $S \cup T$). Note that we do *not* condition on $\mathcal{H}$ being utility-maximizing, which could affect the distribution of the

labels in $S \cup T$. Considering hypothetical $\mathcal{H}$ is only to bound the number of possible responses by the mechanism that we need to union bound over.

Fix any such possible price vector $\boldsymbol{p}$ and $q \in \mathbb{R}_+$. Since $F(q, S, \boldsymbol{p})$ and $F(q, T, \boldsymbol{p})$ are the average of indicators with mean $F(q, S \cup T, \boldsymbol{p})$, Hoeffding for sampling without replacement and the fact that $|S|, |T| \geq \lambda^{10}$ yield

$$\begin{aligned}
\Pr[|F(q,S,\boldsymbol{p}) - F(q,T,\boldsymbol{p})| \geq \lambda^{-4}] &\leq \Pr[|F(q,S,\boldsymbol{p}) - F(q,S\cup T,\boldsymbol{p})| \geq \lambda^{-4}/2] \\
&\quad + \Pr[|F(q,S\cup T,\boldsymbol{p}) - F(q,T,\boldsymbol{p})| \geq \lambda^{-4}/2] \\
&\leq 4\exp(-2(\lambda^{-4}/2)^2\lambda^{10}) \\
&\leq \exp(-\lambda^2/10) \ .
\end{aligned}$$

Since $\boldsymbol{p}$ is fixed, there are at most $m+1$ distinct values taken on by each of $F(\cdot, S, \boldsymbol{p})$ and $F(\cdot, T, \boldsymbol{p})$. Union bounding over these at most $2m+1$ threshold values, and all 2^λ possible prices $\boldsymbol{p}_i(\mathcal{H})$ yields

$$\Pr\left[\max_{q\in\mathbb{R}_+} |F(q,S,\boldsymbol{p}_i(\mathcal{H})) - F(q,T,\boldsymbol{p}_i(\mathcal{H}))| \geq \lambda^{-4}\right] \leq (2m+1)2^\lambda \exp(-\lambda^2/10) \ .$$

Deconditioning on the labels outside $S \cup T$ and union bounding over the at most $4^{n\kappa|R|^n\lambda} \leq \exp(\lambda^{3/2})$ possible choices of labels for $S \cup T$ completes the proof. □

4.3 The Clause Selected by Utility-Maximization

In this section, we describe the utility-maximizing clause. In particular, we first show that at visit i, the utility-maximizing clause chooses the same matrix for all i-relevant $\sigma \in \Sigma_i$. Therefore, on the items which matter for future visits, the mechanism only learns rank information. We then show that the mechanism never learns more than the partition $\mathcal{K}_i$ after each visit i.

Definition 4.8. For $i \in [n]$, $k, r \in [\kappa]$, and $\boldsymbol{p} \in [0,\infty]^m$, define

$$B_{i,k,r}(\boldsymbol{p}) = \int_{\mathrm{val}(r)-1}^{\mathrm{val}(r)} F(q, X_{i,\mathcal{K}_{i-1}(k)}, \boldsymbol{p})\, dq \ ,$$

where we may omit $\boldsymbol{p}$ when obvious from context. We let $B_i \in [0,1]^{\kappa\times\kappa}$ denote the matrix formed from the entries $B_{i,k,r}$. Notice that the rows $B_{i,k} = B_{i,k'}$ when $\mathcal{K}_{i-1}(k) = \mathcal{K}_{i-1}(k')$.

Observe that $B_{i,k,r}$ is precisely the expected utility gain when the value of an item in $X_{i,\mathcal{K}_{i-1}(k)}$ goes from value $\mathrm{val}(r)-1$ to value $\mathrm{val}(r)$. This occurs when the matrix entry $A_{\sigma,k,r}$ increases for i-relevant σ and $k \in K$.

Proposition 4.9. *W.p. at least $1-\lambda^{-1}$, $\mathcal{E}$ and the following hold simultaneously for every $i \in [n]$:*

1. *Before visit i, the items in $X_{i,K}$ are indistinguishable for all $K \in \mathcal{K}_{i-1}$.*
2. *When the utility-maximizing set is non-empty, there exists a matrix $A_i^* \in \mathcal{A}$ such that $s(A_i^*) \leq i-1$, and the utility-maximizing clause $\boldsymbol{A}$ at visit i satisfies $A_\sigma = A_i^*$ for all i-relevant $\sigma \in \Sigma_i$*

Proof. Condition on $\mathcal{E}$ for the rest of the proof.

We prove by induction on i. Suppose that (1) holds for visit i. This is true for $i = 1$ because all items start off indistinguishable. We will show that (1) for visit i implies (2) holds w.h.p., which in turn implies (1) for visit $i+1$. Union bounding over the n visits then completes the proof.

Let $\boldsymbol{p} \in [0,\infty]^m$ be the prices posted by the mechanism at visit i. If the utility-maximizing set is empty, then the reported clause is $\mathbf{0}$ and the mechanism learns nothing. Therefore, assume the utility-maximizing set is non-empty for the rest of the proof.

Let $U_i(\boldsymbol{A}, \boldsymbol{p})$ be the utility of bidder i when they use clause $\boldsymbol{u}^{(i,\boldsymbol{A})}$ under prices $\boldsymbol{p}$. For an i-slice $\sigma \in \Sigma_i$, let $X_{i,\sigma}$ denote the set of items in σ, and define

$$U_i(\sigma, A_\sigma, \boldsymbol{p}) \quad = \quad \sum_{x \in X_{i,\sigma}} (\boldsymbol{u}_x^{(i,\boldsymbol{A})} - p_x)_+ - \kappa\lambda^{10} \operatorname{Avg}[\theta_i A_\sigma + A_\sigma^2/2]$$

to be the utility of the items in $X_{i,\sigma}$ under prices $\boldsymbol{p}$, minus an additional term. Then because the bias is always earned on a non-empty set,

$$U_i(\boldsymbol{A}, \boldsymbol{p}) \quad = \quad W_i(\boldsymbol{p}) + c_i(\boldsymbol{A}, \boldsymbol{p}) + \lambda^{11} + \sum_{\sigma \in \Sigma_i} U_i(\sigma, A_\sigma, \boldsymbol{p}) \ ,$$

where $W_i(\boldsymbol{p})$ is the utility from the items owned by i, and

$$c_i(\boldsymbol{A}, \boldsymbol{p}) \quad = \quad \min\left\{0, \max_{x \in M}\{u_x^{(i,\boldsymbol{A})} - p_x\}\right\}$$

is the potential negative correction for purchasing a "negative utility" item to earn the bias. In particular, observe that $|c_i(\boldsymbol{A}, \boldsymbol{p}) - c_i(\boldsymbol{A'}, \boldsymbol{p})| \le \lambda$ for all $\boldsymbol{A}, \boldsymbol{A'}$ because only non-owner items can change value across clauses, and their value is bounded in $[0, \lambda]$.

Lemma 4.10. *For all i-relevant σ,*

$$\operatorname{Avg}[(A_\sigma^* - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] \quad \le \quad \min_{A \in \mathcal{A}} \operatorname{Avg}[(A - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] + \lambda^{-2} \ .$$

Proof. Fix some i-relevant σ. For a fiber S, every item has the same value w, so the utility contributed by S is

$$\lambda^{10} \int_0^w F(q, S, \boldsymbol{p})\, dq \ .$$

For $k \in [\kappa]$, consider the fibers of the labels $\{(\sigma, k, r_i, \ell) : r_i \in \{0, -1, \ldots, -\kappa\}, \ell \in [\kappa \operatorname{num}(k)]\}$. No matter the matrix A chosen for σ, the number of times each value shows up on the fibers is identical. Therefore, by indistinguishability of $X_{i,\mathcal{K}_{i-1}(k)}$, the utility contributed by these fibers is

$$\delta_{i,k}^-(\sigma, A, \boldsymbol{p}) + \kappa\lambda^{10} \operatorname{num}(k) \sum_{r \in [k]} \int_0^{\operatorname{val}(r)} F(q, X_{i,\mathcal{K}_{i-1}(k)}, \boldsymbol{p})\, dq \ ,$$

where $\delta_{i,k}^-(\sigma, A, \boldsymbol{p})$ is absolutely bounded by $\lambda^{-4} \cdot \kappa\lambda^{10} \operatorname{num}(k) \sum_{r \in [k]} \operatorname{val}(r) \le 2\kappa\lambda^7$ on the event $\mathcal{E}$. Now, for $k \in [\kappa]$, consider the fibers of the labels $\{(\sigma, k, r_i, \ell) : r_i, \ell \in [\kappa]\}$. By indistinguishability of $X_{i,\mathcal{K}_{i-1}(k)}$, the utility contributed by these fibers is

$$\begin{aligned}
&\delta_{i,k}^+(\sigma, A, \boldsymbol{p}) + \lambda^{10} \sum_{r,\ell \in [\kappa]} \int_0^{\operatorname{val}(r)-1+\mathbb{1}(\ell \le A_{k,r})} F(q, X_{i,\mathcal{K}_{i-1}(k)}, \boldsymbol{p})\, dq \\
&\quad = \quad \delta_{i,k}^+(\sigma, A, \boldsymbol{p}) + \kappa\lambda^{10} \sum_{r \in [\kappa]} \int_0^{\operatorname{val}(r)-1} F(q, X_{i,\mathcal{K}_{i-1}(k)}, \boldsymbol{p})\, dq + \lambda^{10} \sum_{r \in [\kappa]} A_{k,r} B_{i,k,r}(\boldsymbol{p}) \ ,
\end{aligned}$$

where $\delta^+_{i,k}(\sigma, A, \boldsymbol{p})$ is absolutely bounded by $\lambda^{-4} \cdot \kappa\lambda^{10} \sum_{r\in[\kappa]} \mathrm{val}(r) \le 2\kappa\lambda^7$ on the event $\mathcal{E}$.

In other words, for $A \in \mathcal{A}$, we have

$$\begin{aligned} U_i(\sigma, A, \boldsymbol{p}) &= U_i(\sigma, \boldsymbol{p}) + \lambda^{10} \sum_{k\in[\kappa]} \sum_{r\in[\kappa]} A_{k,r} B_{i,k,r}(\boldsymbol{p}) - \kappa\lambda^{10} \operatorname{Avg}[\theta_i A + A^2/2] + \delta_i(\sigma, A, \boldsymbol{p}) \\ &= U_i(\sigma, \boldsymbol{p}) + \frac{\kappa\lambda^{10}}{2} \operatorname{Avg}[2\kappa A B_i(\boldsymbol{p}) - 2\theta_i A - A^2] + \delta_i(\sigma, A, \boldsymbol{p}) \\ &= U_i(\sigma, \boldsymbol{p}) + \frac{\kappa\lambda^{10}}{2} \operatorname{Avg}[(\kappa B_i(\boldsymbol{p}) - \theta_i)^2] \\ &\qquad - \frac{\kappa\lambda^{10}}{2} \operatorname{Avg}[(A - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] + \delta_i(\sigma, A, \boldsymbol{p}) \ , \end{aligned}$$

where $U_i(\sigma, \boldsymbol{p})$ is a term not depending on A, $\operatorname{Avg}[(\kappa B_i(\boldsymbol{p}) - \theta_i)^2]$ does not depend on A, and $\delta_i(\sigma, A, \boldsymbol{p})$ is absolutely bounded by $4\kappa^2\lambda^7$ on the event $\mathcal{E}$. Therefore,

$$\begin{aligned} &\operatorname{Avg}[(A^*_\sigma - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] - \frac{2}{\kappa\lambda^{10}} \Big(\delta_i(\sigma, A^*_\sigma, \boldsymbol{p}) + c_i(\boldsymbol{A}^*, \boldsymbol{p}) \Big) \\ &\quad\le \min_{A\in\mathcal{A}} \left\{ \operatorname{Avg}[(A - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] - \frac{2}{\kappa\lambda^{10}} \Big(\delta_i(\sigma, A, \boldsymbol{p}) + c_i((A, \boldsymbol{A}^*_{-\sigma}), \boldsymbol{p}) \Big) \right\} \\ &\quad\le \min_{A\in\mathcal{A}} \operatorname{Avg}[(A - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] - \frac{2}{\kappa\lambda^{10}} \min_{A\in\mathcal{A}} \Big\{ \delta_i(\sigma, A, \boldsymbol{p}) + c_i((A, \boldsymbol{A}^*_{-\sigma}), \boldsymbol{p}) \Big\} \ . \end{aligned}$$

Bounding the contribution of the error terms by $4(4\kappa^2\lambda^7+\lambda)/(\kappa\lambda^{10}) \le \lambda^{-2}$ completes the proof. □

Corollary 4.11. *W.p. at least $1 - (2n\lambda)^{-1}$, it holds for all i-relevant σ that*

$$A^*_\sigma = \operatorname*{argmin}_{A\in\mathcal{A}} \operatorname{Avg}[(A - (\kappa B_i(\boldsymbol{p}) - \theta_i))^2] \ ,$$

and that the minimizer is unique.

Proof. First, observe that the minimizer is unique w.p. 1 because $\theta_i \sim \mathrm{Unif}([0,1])$. Next, observe that by Lemma 4.10, A^*_σ is not the minimizer only if the second-best-minimizer is within λ^{-2} of A^*_σ. This can only happen if there exists $k, r \in [\kappa]$ so that the fractional part of $\kappa B_i(\boldsymbol{p}) - \theta_i$ lies within $\kappa^2\lambda^{-2}$ of $1/2$. The probability this occurs, not conditioned on $\mathcal{E}$, is at most $2\kappa^4/\lambda^2$. By Proposition 4.7, conditioning on $\mathcal{E}$ only distorts this probability by at most λ^{-2}. □

Thus, (2) holds. Finally, observe that the minimizer A^*_i being unique implies that whenever two rows $k, k' \in [\kappa]$ satisfy $B_{i,k} = B_{i,k'}$, we also have $A^*_{i,k} = A^*_{i,k'}$. By construction, this occurs when $\mathcal{K}_{i-1}(k) = \mathcal{K}_{i-1}(k')$, and thus, $s(A^*_i) \le i-1$. This implies that for two i-relevant items with i-role 0 and ranks k, k' satisfying $\mathcal{K}_i(k) = \mathcal{K}_i(k')$, they receive the same coefficient. Since the $(i+1)$-relevant items were precisely the i-relevant items with i-role 0 and owner $j > i$, it holds that before visit $i+1$, the items in the sets $X_{i+1,K}$ for $K \in \mathcal{K}_i$ are indistinguishable.

This completes the induction. A union bound over $\mathcal{E}$ and all n visits completes the proof. □

4.4 Owner Affordability Allows Stealing

At visit i, let $S_i(\boldsymbol{p})$ be the output utility-maximizing set for v_i, and set $S_i(\boldsymbol{p}) = \emptyset$ if bidder i is a sample bidder. Define $f_i, g_{i,j} : [\kappa] \to [0,1]$ to be the fractions of items from each owner-rank pair affordable and allocated to bidder i, i.e.,

$$f_i(k) = F(\beta^{1/2}\,\mathrm{val}(k), Z_{i,k}, \boldsymbol{p})\, \mathbb{1}(i \text{ is not sample}) \qquad \text{and} \qquad g_{i,j}(k) = \frac{|S_i(\boldsymbol{p}) \cap Z_{j,k}|}{|Z_{i,k}|} \ .$$

Lemma 4.12. *Suppose the event of Proposition 4.9 holds. Then for every visit $i \in [n]$ and $j > i$,*

$$\sum_{k\in[\kappa]} f_i(k) \quad \le \quad 2\sum_{k\in[\kappa]} g_{i,j}(k) + \kappa^{2/3} \ .$$

Proof. Fix a visit i. If bidder i is a sample bidder or their utility-maximizing set is $\emptyset$, we are done because $f_i = 0$.

Otherwise, by Proposition 4.9, the same matrix A_i^* is used for every i-relevant i-slice, which includes the slices containing the sets $Z_{j,k}$ for $j > i$ and $k \in [\kappa]$.

Fix $j > i$ and $K \in \mathcal{K}_{s(A_i^*)}$, and let $C, D \in \mathcal{L}(K)$ be consecutive intervals. Let $k \in C$ and $k' \in D$, and let $r = \operatorname{rank}(k', A_i^*) = \min D$. Then

$$\begin{aligned}
f_i(k) \quad &\le \quad \int_{\mathrm{val}(r)-1}^{\mathrm{val}(r)} F(q, Z_{i,k}, \boldsymbol{p})\, dq \\
&\le \quad B_{i,k,r}(\boldsymbol{p}) + \kappa^{-1} \\
&\le \quad \kappa^{-1} A^*_{i,k,r} + 3\kappa^{-1} \\
&= \quad \kappa^{-1} A^*_{i,k',r} + 3\kappa^{-1} \\
&\le \quad B_{i,k',r}(\boldsymbol{p}) + 5\kappa^{-1} \\
&\le \quad \int_{\mathrm{val}(r)-1}^{\mathrm{val}(r)} F(q, Z_{j,k'}, \boldsymbol{p})\, dq + 6\kappa^{-1} \\
&\le \quad g_{i,j}(k') + 6\kappa^{-1} \ .
\end{aligned}$$

The first line follows because $\mathrm{val}(r) - 1 \ge \beta^{1/2}\,\mathrm{val}(k)$. The second and sixth lines follow because $Z_{i,k}$ lies in the indistinguishable set $X_{i,\mathcal{K}_{i-1}(k)}$, and similarly for k'. The third and fifth lines follow because $A^*_{i,k,r}/\kappa \in \{0, \kappa^{-1}, 2\kappa^{-1}, \dots, 1\}$ minimizes the distance to $B_{i,k,r} - \theta_i/\kappa$, and similarly for k'. The fourth line follows because $\mathcal{K}_{s(A_i^*)}(k) = \mathcal{K}_{s(A_i^*)}(k')$. The last line follows because $\mathrm{val}(k', A_i^*) = \mathrm{val}(r)$. Applying Lemma 3.15 completes the proof. □

4.5 Wrapping Up

Proof of Theorem 2.1. The non-owner welfare (including the bias) is bounded by

$$\begin{aligned}
\sum_{i\in[n]} v_i(M \setminus Z_{i,[\kappa]}) \quad &\le \quad n|\Sigma_i|\lambda^{10}\Bigg(\lambda + \sum_{k\in[\kappa]}\sum_{r\in[k]}\Big(\kappa\,\mathrm{num}(k)\,\mathrm{val}(r)\Big) + \sum_{k\in[\kappa]}\sum_{r\in[\kappa]}\Big(\kappa\,\mathrm{val}(r)\Big)\Bigg) \\
&\le \quad 5n|\Sigma_i|\kappa^2\lambda^{11} \\
&\le \quad \beta^{1/4}\kappa^2\lambda^{11} \ .
\end{aligned}$$

By Lemma 4.12, the owner welfare (no bias, and on the event of Proposition 4.9) is bounded by

$$\begin{aligned}
\beta^{1/2}\kappa\lambda^{11}\sum_{i\in[n]}\sum_{k\in[\kappa]} f_i(k) \quad &\le \quad \beta^{1/2}\kappa\lambda^{11}\Bigg(\sum_{k\in[\kappa]} f_n(k) + 2\sum_{i\in[n-1]}\sum_{k\in[\kappa]} g_{i,n}(k) + n\kappa^{2/3}\Bigg) \\
&\le \quad 2\beta^{1/2}\kappa\lambda^{11}\sum_{k\in[\kappa]}\Bigg(f_n(k) + \sum_{i\in[n-1]} g_{i,n}(k)\Bigg) + \beta^{1/4}\kappa^2\lambda^{11} \\
&\le \quad (2 + \beta^{-1/4})\beta^{1/2}\kappa^2\lambda^{11} \ ,
\end{aligned}$$

where the last line comes from feasibility of the allocation (the total allocated rank fraction is 1). On the other hand, the optimal welfare OPT is at least

$$\sum_{i\in[n]} v_i(Z_{i,[\kappa]}) \quad \geq \quad n\beta^{1/2}\kappa^2\lambda^{11} \ .$$

Thus, by Proposition 4.9, the approximation ratio is lower bounded by

$$\frac{n}{2+2\beta^{-1/4}+n\lambda^{-1}} \quad \geq \quad \frac{n}{3} \quad = \quad \Theta(\log\log m) \ . \qquad \square$$

5 Full-Revelation Lower Bound

For sufficiently large n, let

- $\kappa = 2^{20n}$ be the number of *ranks*.
- $\mathrm{val}(k) = \kappa^{\kappa^k}$ for $k \in [\kappa]$ be the *value* of items of rank k. Denote $\mathrm{val}(0) = 0$.
- $\lambda = \mathrm{val}(\kappa)$.
- $m = n\lambda^{100}$ be the number of items.
- $\eta = m^{2n}$ be the size of the *name pool*.
- $\varepsilon = \lambda^{-10}$ be a discretization/error parameter.

Observe that $n = \Theta(\log\log\log m)$.

Partition the items into n blocks $Z_1, \ldots, Z_n$ of size λ^{100}. Sample a uniformly random permutation $\pi : [n] \to [n]$. Bidder i will be the *owner* of $Z_{\pi(i)}$, and a *non-owner* of Z_t for $t \neq \pi(i)$. Each bidder's valuation will be the direct sum of an XOS valuation for each block. Let us restrict attention to a block Z_t, and describe the owner and non-owner valuations for Z_t. Note that the following random objects we sample are per-block, and the blocks are sampled independent of π and each other.

Sample a uniformly random injection $\tau_t : [n] \to [\eta]$ assigning each bidder a name. Call the set $\tau_t([n])$ the *real names*. Sample $N_{t,1}, \ldots, N_{t,n} \subseteq [\eta]$ conditionally independently on $\tau_t([n])$ such that each set contains every real name w.p. 1 and every other name independently w.p. $1/m$.

A *context* consists of a non-empty set $L \subseteq [\eta]$ of size at most n and an element $\ell \in L$. For a context $\gamma = (\ell, L)$, denote $|\gamma| = |L|$, and say that $\gamma \subseteq L'$ if $L \subseteq L'$. Let Γ denote the set of all contexts.

For each $\gamma \in \Gamma$, independently draw a function $\psi_{t,\gamma} : Z_t \to \{0, 1, \ldots, \kappa\}$ where independently for each $x \in Z_t$,

$$\Pr[\psi_{t,\gamma}(x) = k] \quad = \quad \begin{cases} \dfrac{2}{\kappa\,\mathrm{val}(k)} & k \in [\kappa] \\ 1 - \displaystyle\sum_{k\in[\kappa]} \dfrac{2}{\kappa\,\mathrm{val}(k)} & k = 0 \end{cases} \ .$$

For each $\gamma \in \Gamma$, we also independently draw a function $\xi_{t,\gamma} : Z_t \to [0, 2]$ uniformly randomly. Let $\psi'_{t,\gamma}$ be the function defined by $\psi'_{t,\gamma}(x) = \psi_{t,\gamma}(x)\,\mathbb{1}(\xi_{t,\gamma}(x) \leq 1)$.

For a set $L \subseteq [\eta]$, let $\mathcal{P}_{t,L}$ be the partition of Z_t induced by $\{\psi'_{t,\gamma} : \gamma \subseteq L\}$, i.e., items $x, y \in Z_t$ are in the same part of $\mathcal{P}_{t,L}$ if $(\psi'_{t,\gamma}(x) : \gamma \subseteq L) = (\psi'_{t,\gamma}(y) : \gamma \subseteq L)$. For a context $\gamma = (\ell, L)$, let $\mathcal{P}_{t,\gamma} = \mathcal{P}_{t,L\setminus\ell}$. For a partition $\mathcal{P}$ of Z_t, we say that a function f on Z_t is $\mathcal{P}$-*measurable* if $f|_P$ is constant for all $P \in \mathcal{P}$. For a partition $\mathcal{P}$ of Z_t, let $\mathcal{P}(x)$ denote the set $P \in \mathcal{P}$ such that $x \in P$.

For a function f on Z_t, let $\Pi_{t,L}f$ be the function defined by $(\Pi_{t,L}f)(x) = \text{Avg}[f \mid \mathcal{P}_{t,L}(x)]$. In other words, $\Pi_{t,L}f$ is the $\mathcal{P}_{t,L}$-measurable function formed by averaging over each $P \in \mathcal{P}_{t,L}$. For a function f on $[\kappa] \times Z_t$, let $\Pi_{t,L}f$ be the function defined by $(\Pi_{t,L}f)(k,x) = \text{Avg}[f(k,\cdot) \mid \mathcal{P}_{t,L}(x)]$. Define $\Pi_{t,(\ell,L)}$ identically to $\Pi_{t,L\setminus\ell}$.

5.1 The Non-Owner Valuation

For a context γ and function $A : [\kappa] \times Z_t \to \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$, define $\boldsymbol{u}^{(t,\gamma,A)} \in \mathbb{R}_+^{|Z_t|+1}$ by

$$u_0^{(t,\gamma,A)} \quad = \quad |Z_t|(1 - \text{Avg}[A^2]) + 8\varepsilon|Z_t|(n - |\gamma|)$$

and

$$u_x^{(t,\gamma,A)} \quad = \quad \begin{cases} 2\,\text{val}(\psi_{t,\gamma}(x)) & \xi_{t,\gamma}(x) \le 1 + A(\psi_{t,\gamma}(x), x) \\ 0 & \text{o.w.} \end{cases} ,$$

where we consider $A(0,\cdot) = 0$. In other words, $\psi_{t,\gamma}$ determines the candidate value of every item, and each item actually has its candidate value w.p. in $[1/2, 1]$, depending on the function A. The bias decreases when the function A increases or when $|\gamma|$ increases.

For $i \in [n]$, let $\Gamma_{t,i}$ be the set of all contexts $\gamma = (\ell, L)$ such that $\ell = \tau_t(i)$ and $L \subseteq N_{t,i}$. For $L \subseteq [\eta]$, let $\mathcal{A}_{t,L}$ be the set of functions $A : [\kappa] \times Z_t \to \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$ such that $A(k,\cdot)$ is $\mathcal{P}_{t,L}$-measurable for all $k \in [\kappa]$, and define $\mathcal{A}_{t,(\ell,L)} = \mathcal{A}_{t,L\setminus\ell}$ similarly. We highlight that $\mathcal{A}_{t,L}$ is itself a random object. Then the bidder i non-owner valuation $v_{t,i} : 2^{Z_t} \to \mathbb{R}_+$ is given by

$$v_{t,i}(S) \quad = \quad \begin{cases} \max\limits_{\gamma \in \Gamma_{t,i}} \max\limits_{A \in \mathcal{A}_{t,\gamma}} \{\boldsymbol{u}^{(t,\gamma,A)}(S) + u_0^{(t,\gamma,A)}\} & S \neq \emptyset \\ 0 & S = \emptyset \end{cases} .$$

Note that $v_{t,i}$ is also defined for the owner of Z_t.

5.2 The Owner Valuation

Let $\mathcal{K}_s = \mathcal{L}^s(\{[\kappa]\})$, and define $\varphi_s(k) = \min \mathcal{K}_s(k) - 1$. For $I \subseteq [n]$ and $k \in [\kappa]$, define

$$Z_{t,k,I} \quad = \quad \bigcap_{\gamma \subseteq \tau_t(I)} \{x \in Z_t : \psi'_{t,\gamma}(x) \ge \varphi_{|\gamma|}(k)\} ,$$

where $Z_{t,k,\emptyset} = Z_t$. Also fix some $x \in Z_t$ and define

$$\vartheta_{t,k,I} \quad = \quad \Pr[x \in Z_{t,k,I}] .$$

Lemma 5.1. *$\vartheta_{t,k,I} \ge 1/\sqrt{\text{val}(k)}$ for all $I \subseteq [n]$ and $k \in [\kappa]$.*

Proof. The statement is immediate for $k = 1$ or $I = \emptyset$ because $\vartheta_{t,1,I} = \vartheta_{t,k,\emptyset} = 1$. Otherwise, since

every $\psi_{t,\gamma}$ and $\xi_{t,\gamma}$ is sampled independently, we have

$$\begin{aligned}
\Pr[x \in Z_{t,k,I}] &\geq \prod_{\gamma \subseteq \tau_t(I)} \Big(\Pr[\psi_{t,\gamma}(x) \geq \varphi_{|\gamma|}(k)] \Pr[\xi_{t,\gamma}(x) \leq 1] \Big) \\
&\geq \prod_{\gamma \subseteq \tau_t(I)} \frac{\Pr[\psi_{t,\gamma}(x) = k-1]}{2} \\
&\geq \prod_{\gamma \subseteq \tau_t(I)} \frac{1}{\kappa^{1+\kappa^{k-1}}} \\
&\geq \frac{1}{(\kappa^{1+\kappa^{k-1}})^{2^n n}} \\
&\geq \frac{1}{\sqrt{\operatorname{val}(k)}} \ .
\end{aligned}$$

□

Items $x \in Z_t$ are independently assigned to one or none of $Z_{t,1}, \ldots, Z_{t,\kappa}$, where

$$\Pr[x \in Z_{t,k} \mid \tau_t, \{\psi'_{t,\gamma} : \gamma \in \Gamma\}] = \frac{\mathbb{1}(x \in Z_{t,k,[n]})}{\vartheta_{t,k,[n]} \operatorname{val}(k)}$$

for $k \in [\kappa]$. This distribution is well-defined because by Lemma 5.1,

$$\sum_{k \in [\kappa]} \frac{\mathbb{1}(x \in Z_{t,k,[n]})}{\vartheta_{t,k,[n]} \operatorname{val}(k)} \leq \sum_{k \in [\kappa]} \frac{1}{\sqrt{\operatorname{val}(k)}} \leq \frac{\kappa}{\sqrt{\operatorname{val}(1)}} \leq 1 \ .$$

The owner valuation $w_t : 2^{Z_t} \to \mathbb{R}_+$ is additive, where

$$w_t(x) = \begin{cases} \operatorname{val}(k) & x \in Z_{t,k} \\ 0 & \text{o.w.} \end{cases} \ .$$

Observe that $\mathbb{E}[|Z_{t,k}|] = |Z_t|/\operatorname{val}(k)$. Therefore, $\mathbb{E}[w_t(Z_{t,k})] = |Z_t| = \lambda^{100}$, and the expected optimal welfare of the full instance is at least $n\kappa\lambda^{100}$.

5.3 The Full Instance, Randomness Recap, and Road Map

The valuation $v_i : 2^M \to \mathbb{R}_+$ for bidder i is defined

$$v_i(S) = w_{\pi(i)}(S \cap Z_{\pi(i)}) + \sum_{t \neq \pi(i)} v_{t,i}(S \cap Z_t) \ ,$$

which is XOS because direct sums of XOS functions are XOS. Further, the lexicographically smallest utility-maximizing set of v_i is simply the union of the lexicographically smallest utility-maximizing sets of $w_{\pi(i)}$ and $v_{t,i}$ for $t \neq \pi(i)$. Hence, we may analyze the utility of v_i by analyzing the block-wise utility for $w_{\pi(i)}$ and $v_{t,i}$ for $t \neq \pi(i)$.

Observe that the randomness in the instance comes from the owner permutation π, the real name functions $\{\tau_t : t \in [n]\}$, the name lists $\{N_{t,i} : t, i \in [n]\}$, the functions $\{\psi_{t,\gamma}, \xi_{t,\gamma} : t \in [n], \gamma \in \Gamma\}$, and the conditional draws of $\{Z_{t,k} : t \in [n], k \in [\kappa]\}$. For the analysis, we will assume that

$\{\psi_{t,\gamma}, \xi_{t,\gamma} : t \in [n], \gamma \in \Gamma\}$ are known to the mechanism in advance. Additionally, after the mechanism visits bidder i, it will be granted $\pi(i)$, $\{\tau_t(i) : t \in [n]\}$, and $\{N_{t,i} : t \in [n]\}$. The only other information revealed to the mechanism after visiting bidder i is the sets $\{Z_{\pi(i),k} : k \in [\kappa]\}$ which have high value for the owner. Hence, the only information hidden to the mechanism are which blocks are owned by the unvisited bidders, and for each such block, the unvisited bidders' names and name lists, and the high value owner sets.

Observation 5.2. *WLOG, we may assume the mechanism visits bidders in the fixed order* $1, \ldots, n$.

Proof. Suppose the unvisited bidders are $I \subseteq [n]$. Choose a permutation $\sigma : I \to I$, and consider the transformed instance where for all $i \in I$ and $t \in [n]$, we set $\pi'(i) = \pi(\sigma(i))$, $\tau'_t(i) = \tau_t(\sigma(i))$, and $N'_{t,i} = N_{t,\sigma(i)}$. Since the owner ranks $Z_{t,k}$ depend on $\tau'_t([n]) = \tau_t([n])$, this transformation preserves the prior distribution. Further, the observed transcript up to this point is preserved, because σ only applies to unvisited bidders and information revealed after a visit. Thus, conditioned on the observed history, the transformation does not change the joint distribution of the unvisited bidders' information, so no unvisited identity is a more advantageous next choice than another. □

Therefore, fix a deterministic mechanism which visits bidders in the fixed order $1, \ldots, n$.

Definition 5.3 (Relevant)**.** A block-visit pair $(t, i) \in [n] \times [n]$ is *relevant* if $t \notin \pi([i-1])$, i.e., the owner of Z_t has not yet been visited.

The following is an idealization of how the construction works:

1. After visiting bidders $[i-1]$, the mechanism no longer needs to consider blocks $\pi([i-1])$, as their owners have already been visited and the non-owner value is negligible. Therefore, the mechanism wants to set prices on the remaining blocks such that the high value items are *affordable* to the owner but not to the non-owners.
2. Focus on such a block Z_t. Using the functions $\{\psi_{t,\gamma}, \xi_{t,\gamma} : \gamma \in \Gamma\}$ and the names $\tau_t([i-1])$, the mechanism is able to deduce the sets $Z_{t,k,[i-1]} \supseteq Z_{t,k}$.
3. Let r, r' be such that $r < r'$, $\mathcal{K}_{i-1}(r) = \mathcal{K}_{i-1}(r')$, and $\mathcal{K}_i(r) \neq \mathcal{K}_i(r')$. In other words, r, r' belong to the same interval in $\mathcal{K}_{i-1}$, but r' belongs to a higher interval in the refinement $\mathcal{K}_i$.
4. Because $Z_{t,r,[i-1]}, Z_{t,r',[i-1]}$ are the intersection of sets defined using $\varphi_s(r)$ for $s < i$, we have $Z_{t,r,[i-1]} = Z_{t,r',[i-1]}$, so the prices set by the mechanism result in roughly the same affordability on the items in $Z_{t,r}$ and $Z_{t,r'}$.
5. Then the mechanism visits bidder i. Suppose that bidder i is a non-owner of Z_t, and that they respond with a clause using context $\gamma \subseteq \tau_t([i])$; for simplicity, assume $\gamma = (\tau_t(i), \tau_t([i]))$. Then recall that every $x \in Z_{t,r'}$ satisfies $\psi'_{t,\gamma}(x) \geq \varphi_i(r') = \min \mathcal{K}_i(r') - 1 \geq r$. Therefore, if $Z_{t,r}$ were affordable for the owner of the block, then $Z_{t,r'}$ is affordable for bidder i.
6. Therefore, if the owner can recover some value W, a non-owner can destroy value $\Omega(W)$.

The idealized execution is essentially the same as in the single-query lower bound; the heavy-lifting is proving that the additional information provided to the mechanism does not help. We first prove that utility-maximization indeed causes bidder i to respond in the way claimed in (5). We then prove that the balance claim in (4) holds. Finally, we argue rigorously why owner affordability implies non-owner stealing.

Throughout the proof, we will make reference to a "good event" which holds w.p. $1 - \lambda^{-1}$ and makes the arguments go through. When such probabilistic claims appear, we will provide intuition for why they should hold, but defer their full proofs to Section A.

Definition 5.4 (The Good Event)**.** The *good event* $\mathcal{E}$ is the event that Proposition A.3 holds (which includes the event of Proposition A.2), and occurs w.p. at least $1-\lambda^{-1}$. We omit explicitly listing all the conditions of the good event, instead stating them piecemeal as they become relevant in the proof. Their proofs are aggregated because they fall into a common framework.

5.4 The Clause Selected by Utility-Maximization

In this section, we will show which clause maximizes the utility for a non-owner valuation for the block Z_t. In particular, we will show that any utility-maximizing clause chooses $A : [\kappa] \times Z_t \to \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$ which best approximates the function $B_{t,\boldsymbol{p}} : [\kappa] \times Z_t \to [0,1]$ given by

$$B_{t,\boldsymbol{p}}(k,x) \quad = \quad \left(1 - \frac{p_x}{2\,\mathrm{val}(k)}\right)_+ .$$

We can think of $B_{t,\boldsymbol{p}}(k,x)$ as the fraction of utility captured from item x in a clause which assigns x value $2\,\mathrm{val}(k)$. Define

$$W_t(\boldsymbol{p}) \quad = \quad \sum_{x \in Z_t} \Big(w_t(x)\, \mathbb{1}(w_t(x) > p_x) \Big)$$

as the value obtained by the owner valuation under prices $\boldsymbol{p} \in [0,\infty]^{Z_t}$, and define

$$U_t(\gamma, A, \boldsymbol{p}) \quad = \quad \frac{1}{|Z_t|} \max_{S \subseteq Z_t} \Big\{ \boldsymbol{u}^{t,\gamma,A}(S) + u_0^{(t,\gamma,A)}\, \mathbb{1}(S \neq \emptyset) - \boldsymbol{p}(S) \Big\}$$

to be the normalized utility obtained by the clause $\boldsymbol{u}^{(t,\gamma,A)}$ under prices $\boldsymbol{p}$.

Proposition 5.5. *Suppose $\mathcal{E}$ holds and $W_t(\boldsymbol{p}) \geq |Z_t|$. Let (t,i) be a relevant block-visit pair where the mechanism posts prices $\boldsymbol{p} \in [0,\infty]^{Z_t}$. Let $\gamma^* = (\tau_t(i), L^*)$ and $A^* : [\kappa] \times Z_t \to \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$ define a utility-maximizing clause for $v_{t,i}$. Then:*

1. *$L^* \subseteq \tau_t([i])$.*
2. *$\mathrm{Avg}[(A^* - \Pi_{t,\tau_t([i-1])} B_{t,\boldsymbol{p}})^2] \leq 8\varepsilon n$.*

Proof. For a context γ, and functions $A, B : [\kappa] \times Z_t \to [0,1]$, define

$$U_t(\gamma, A, B) \quad = \quad \frac{1}{|Z_t|} \sum_{x \in Z_t} \sum_{k \in [\kappa]} \Big(2\,\mathrm{val}(k)\, \mathbb{1}(\psi_{t,\gamma}(x) = k)\, \mathbb{1}(\xi_{t,\gamma}(x) \leq 1 + A(k,x))\, B(k,x) \Big) \ .$$

Note that $U_t(\gamma, A, B_{t,\boldsymbol{p}})$ is precisely the normalized utility of the clause $\boldsymbol{u}^{(t,\gamma,A)}$ when the bias is set to 0. Therefore, we have the implication

$$U_t(\gamma, A, B_{t,\boldsymbol{p}}) \ > \ 0 \qquad \Longrightarrow \qquad U_t(\gamma, A, \boldsymbol{p}) \ = \ U_t(\gamma, A, B_{t,\boldsymbol{p}}) + u_0^{(t,\gamma,A)}/|Z_t| \ . \tag{1}$$

Claim 5.6. *Suppose $\mathcal{E}$ holds. Then for all $\gamma \in \Gamma_{t,i}$ and $A \in \mathcal{A}_{t,\gamma}$,*

$$\Big| U_t(\gamma, A, B_{t,\boldsymbol{p}}) - \Big(2\,\mathrm{Avg}[B_{t,\boldsymbol{p}}] + 2\,\mathrm{Avg}[AB_{t,\boldsymbol{p}}] \Big) \Big| \quad \leq \quad 2\varepsilon \ .$$

Proof Sketch. A direct computation shows that for a fixed choice of $\gamma = (\ell, L)$, A, and B, conditioning on $\{\psi_{t,\gamma'}, \xi_{t,\gamma'} : \gamma' \subseteq L \setminus \ell\}$ gives that $U_t(\gamma, A, B)$ is a sum over independent random variables with $\mathbb{E}[U_t(\gamma, A, B)] = 2\,\mathrm{Avg}[B] + 2\,\mathrm{Avg}[AB]$. Therefore, we expect $U_t(\gamma, A, B)$ to concentrate, and a union bound yields concentration for all γ, A.

The challenge is that the concentration holds over randomness of $\psi_{t,\gamma}, \xi_{t,\gamma}$, but this is public information on which the mechanism can condition its prices and hence $B_{t,\boldsymbol{p}}$. The trick is to discretize the image of $B_{t,\boldsymbol{p}}$ and union bound over all possible discretizations. Naively, this fails because the mechanism can overfit for any fixed γ, A by looking at $\psi_{t,\gamma}, \xi_{t,\gamma}$. Indeed, we can achieve failure probability as low as $\exp(-\lambda^{77})$, but we need to union bound over $\approx \exp(\lambda^{100})$ choices of $B_{t,\boldsymbol{p}}$.

However, we can instead consider, for each $B_{t,\boldsymbol{p}}$, the probability that λ^{29} "sufficiently independent" choices of γ, A have large deviations; the probability that all choices fail is $\approx \exp(-\lambda^{106})$, which allows us to apply the union bound. We can then prove that since the hidden real names and name lists are still random and sparse conditioned on the mechanism's information, the probability is low that the chosen price vector causes a large deviation on any clause $\boldsymbol{u}^{(t,\gamma,A)}$ which actually appears in $v_{t,i}$. See Proposition A.3-(1) for details. □

Corollary 5.7. *Suppose $\mathcal{E}$ holds and $W_t(\boldsymbol{p}) \geq |Z_t|$. Then for all $\gamma \in \Gamma_{t,i}$ and $A \in \mathcal{A}_{t,\gamma}$,*

$$\left|U_t(\gamma, A, \boldsymbol{p}) - \left(1 + 8\varepsilon n + 2\operatorname{Avg}[B_{t,\boldsymbol{p}}] + \operatorname{Avg}[B_{t,\boldsymbol{p}}^2] - \operatorname{Avg}[(A - B_{t,\boldsymbol{p}})^2] - 8\varepsilon|\gamma|\right)\right| \quad \leq \quad 2\varepsilon \ .$$

Proof. Since the owner values an item $x \in Z_{t,k}$ at $\operatorname{val}(k)$, x is only affordable for the owner if $B(k,x) \geq 1/2$, as $B_{t,\boldsymbol{p}}(k,x)$ gives the normalized utility of x when its value is $2\operatorname{val}(k)$. Therefore,

$$\begin{aligned} W_t(\boldsymbol{p}) \quad &\leq \quad \sum_{k\in[\kappa]} \sum_{x\in Z_t} \Big(\operatorname{val}(k)\, \mathbb{1}(B_{t,\boldsymbol{p}}(k,x) \geq 1/2)\Big) \\ &\leq \quad 2\lambda \sum_{k\in[\kappa]} \sum_{x\in Z_t} B_{t,\boldsymbol{p}}(k,x) \\ &= \quad 2\kappa\lambda |Z_t| \operatorname{Avg}[B_{t,\boldsymbol{p}}] \ . \end{aligned}$$

By Claim 5.6, and the fact that $W_t(\boldsymbol{p}) \geq |Z_t|$, we have

$$U_t(\gamma, A, B_{t,\boldsymbol{p}}) \quad \geq \quad 2\operatorname{Avg}[B_{t,\boldsymbol{p}}] - 2\varepsilon \quad \geq \quad \frac{W_t(\boldsymbol{p})}{\kappa\lambda|Z_t|} - 2\varepsilon \quad \geq \quad \frac{1}{\kappa\lambda} - 2\varepsilon \quad > \quad 0 \ ,$$

which implies $U_t(\gamma, A, \boldsymbol{p}) = U_t(\gamma, A, B_{t,\boldsymbol{p}}) + u_0^{(t,\gamma,A)}/|Z_t|$ by (1). Claim 5.6 then yields

$$\begin{aligned} &\left|U_t(\gamma, A, \boldsymbol{p}) - \left(1 + 8\varepsilon n + 2\operatorname{Avg}[B_{t,\boldsymbol{p}}] + \operatorname{Avg}[B_{t,\boldsymbol{p}}^2] - \operatorname{Avg}[(A - B_{t,\boldsymbol{p}})^2] - 8\varepsilon|\gamma|\right)\right| \\ &\quad = \quad \left|U_t(\gamma, A, B_{t,\boldsymbol{p}}) - \operatorname{Avg}[A^2] - \left(2\operatorname{Avg}[B_{t,\boldsymbol{p}}] + \operatorname{Avg}[B_{t,\boldsymbol{p}}^2] - \operatorname{Avg}[(A - B_{t,\boldsymbol{p}})^2]\right)\right| \\ &\quad = \quad \left|U_t(\gamma, A, B_{t,\boldsymbol{p}}) - \left(2\operatorname{Avg}[B_{t,\boldsymbol{p}}] + 2\operatorname{Avg}[AB_{t,\boldsymbol{p}}]\right)\right| \quad \leq \quad 2\varepsilon \ . \end{aligned}$$

□

Therefore, utility-maximization (approximately) minimizes

$$\operatorname{Avg}[(A - B_{t,\boldsymbol{p}})^2] + 8\varepsilon|\gamma| \ .$$

Recall that $A \in \mathcal{A}_{t,\gamma}$, which imposes that $A(k,\cdot)$ is $\mathcal{P}_{t,\gamma}$-measurable. Adding names to γ incurs a penalty, but allows A to operate over a richer partition. Our objective is to show that the choice of $\gamma^* \subseteq \tau_t([i])$. To this end, for $L \subseteq [\eta]$ with $|L| \leq n-1$ and $B : [\kappa] \times Z_t \to [0,1]$, define

$$\Delta_t(L, B) \quad = \quad \min_{A\in\mathcal{A}_{t,L}} \operatorname{Avg}[(A - B)^2] \ .$$

Claim 5.8. *Suppose $\mathcal{E}$ holds. Then for all $L \subseteq N_{t,i} \setminus \tau_t(i)$ with $|L| \le n-1$,*

$$\Delta_t(L \cap \tau_t([i-1]), B_{t,\boldsymbol{p}}) - \Delta_t(L, B_{t,\boldsymbol{p}}) \quad \le \quad 2\varepsilon$$

Proof Sketch. Consider a fixed B and condition on $\mathcal{P}_{t,L\cap\tau_t([i-1])}$. Then since the partition $\mathcal{P}_{t,L}$ is a permutation invariant refinement of $\mathcal{P}_{t,L\cap\tau_t([i-1])}$, we expect that the best $\mathcal{P}_{t,L}$-measurable approximation of B is not much better than the best $\mathcal{P}_{t,L\cap\tau_t([i-1])}$-measurable approximation.

Again though, the mechanism can condition its prices on the randomness of $\mathcal{P}_{t,L}$, and can therefore make $\Delta_t(L, B_{t,\boldsymbol{p}})$ deviate from $\Delta_t(L\cap\tau_t([i-1]), B_{t,\boldsymbol{p}})$. The trick is similar to the one in Claim 5.6; show that for every $B_{t,\boldsymbol{p}}$, there exists a set $T \subseteq [\eta]$ such that $L \subseteq [\eta]$ has a large deviation only if $L \cap T \neq \emptyset$. Then provided T is small, $\Pr[T \cap N_{t,i} \neq \emptyset] \approx 0$, and therefore the contexts $\gamma \subseteq N_{t,i}$ concentrate. See Proposition A.3-(2) for details. □

Let $\gamma' = (\tau_t(i), L^* \cap \tau_t([i]))$. By optimality of γ^*, A^*, Corollary 5.7, and Claim 5.8,

$$\begin{aligned}
0 \quad &= \quad U_t(\gamma^*, A^*, \boldsymbol{p}) - \max_{\gamma\in\Gamma_{t,i}} \max_{A\in\mathcal{A}_{t,\gamma}} U_t(\gamma, A, \boldsymbol{p}) \\
&\le \quad \min_{\gamma\in\Gamma_{t,i}} \min_{A\in\mathcal{A}_{t,\gamma}} \Big\{ \mathrm{Avg}[(A - B_{t,\boldsymbol{p}})^2] + 8\varepsilon|\gamma| \Big\} - \Big(\mathrm{Avg}[(A^* - B_{t,\boldsymbol{p}})^2] + 8\varepsilon|\gamma^*| \Big) + 4\varepsilon \\
&\le \quad \min_{A\in\mathcal{A}_{t,\gamma'}} \Big\{ \mathrm{Avg}[(A - B_{t,\boldsymbol{p}})^2] \Big\} + 8\varepsilon|\gamma'| - \Big(\mathrm{Avg}[(A^* - B_{t,\boldsymbol{p}})^2] + 8\varepsilon|\gamma^*| \Big) + 4\varepsilon \\
&\le \quad \Delta_t(L^* \cap \tau_t([i-1]), B_{t,\boldsymbol{p}}) - \Delta_t(L^* \setminus \tau_t(i), B_{t,\boldsymbol{p}}) + 8\varepsilon(|\gamma'| - |\gamma^*|) + 4\varepsilon \\
&\le \quad 6\varepsilon - 8\varepsilon|L^* \setminus \tau_t([i])| \ ,
\end{aligned}$$

which implies $|L^* \setminus \tau_t([i])| = 0$, proving Proposition 5.5-(1). We can now prove the following relation.

Lemma 5.9. *Let $A : [\kappa] \times Z_t \to [0,1]$ be $\mathcal{P}_{t,L}$-measurable. Then*

$$\mathrm{Avg}[(A - B_{t,\boldsymbol{p}})^2] \quad = \quad \mathrm{Avg}[(A - \Pi_{t,L} B_{t,\boldsymbol{p}})^2] + \mathrm{Avg}[(\Pi_{t,L} B_{t,\boldsymbol{p}} - B_{t,\boldsymbol{p}})^2] \ .$$

Proof. For any $k \in [\kappa]$ and $P \in \mathcal{P}_{t,L}$,

$$\begin{aligned}
&\mathrm{Avg}[(A - \Pi_{t,L} B_{t,\boldsymbol{p}})(\Pi_{t,L} B_{t,\boldsymbol{p}} - B_{t,\boldsymbol{p}}) \mid k \times P] \\
&\qquad = \quad \mathrm{Avg}[A - \Pi_{t,L} B_{t,\boldsymbol{p}} \mid k \times P] \, \mathrm{Avg}[\Pi_{t,L} B_{t,\boldsymbol{p}} - B_{t,\boldsymbol{p}} \mid k \times P] \quad = \quad 0 \ ,
\end{aligned}$$

because $A - \Pi_{t,L} B_{t,\boldsymbol{p}}$ is a constant over $k \times P$ and can be separated out. Thus, the cross-term in the expansion below disappears:

$$\begin{aligned}
&\mathrm{Avg}[(A - B_{t,\boldsymbol{p}})^2 \mid k \times P] \\
&\qquad = \quad \mathrm{Avg}[(A - \Pi_{t,L} B_{t,\boldsymbol{p}} + \Pi_{t,L} B_{t,\boldsymbol{p}} - B_{t,\boldsymbol{p}})^2 \mid k \times P] \\
&\qquad = \quad \mathrm{Avg}[(A - \Pi_{t,L} B_{t,\boldsymbol{p}})^2 \mid k \times P] + \mathrm{Avg}[(\Pi_{t,L} B_{t,\boldsymbol{p}} - B_{t,\boldsymbol{p}})^2 \mid k \times P] \ .
\end{aligned}$$

Averaging over $k \in [\kappa]$, $P \in \mathcal{P}_{t,L}$ with weights $|P|/(\kappa|Z_t|)$ completes the proof. □

Now, let $\gamma' = (\tau_t(i), \tau_t([i]))$. By optimality of γ^*, A^* and Corollary 5.7,

$$\begin{aligned}
0 \quad &\le \quad U_t(\gamma^*, A^*, \boldsymbol{p}) - \max_{A\in\mathcal{A}_{t,\gamma'}} U_t(\gamma', A, \boldsymbol{p}) \\
&\le \quad \min_{A\in\mathcal{A}_{t,\gamma'}} \mathrm{Avg}[(A - B_{t,\boldsymbol{p}})^2] - \mathrm{Avg}[(A^* - B_{t,\boldsymbol{p}})^2] + 8\varepsilon(|\gamma'| - |\gamma^*|) + 4\varepsilon \\
&\le \quad \min_{A\in\mathcal{A}_{t,\gamma'}} \mathrm{Avg}[(A - B_{t,\boldsymbol{p}})^2] - \mathrm{Avg}[(A^* - B_{t,\boldsymbol{p}})^2] + 8\varepsilon n - 4\varepsilon \ .
\end{aligned}$$

Since $\gamma^* \subseteq \tau_t([i])$, A^* is $\mathcal{P}_{t,\tau_t([i-1])}$-measurable. Also, $\Pi_{t,\gamma'} B_{t,\boldsymbol{p}}$ is the minimizer of $\mathrm{Avg}[(f - B_{t,\boldsymbol{p}})^2]$ among $\mathcal{P}_{t,\gamma'}$-measurable functions, so the best $A \in \mathcal{A}_{t,\gamma'}$ is the discretization of $\Pi_{t,\gamma'} B_{t,\boldsymbol{p}}$ to the grid $\{0, \varepsilon, 2\varepsilon, \dots, 1\}$. Thus, bounding the discretization error by ε^2 and applying Lemma 5.9,

$$\begin{aligned} 0 &\leq \mathrm{Avg}[(\Pi_{t,\gamma'} B_{t,\boldsymbol{p}} - B_{t,\boldsymbol{p}})^2] + \varepsilon^2 - \mathrm{Avg}[(A^* - B_{t,\boldsymbol{p}})^2] + 8\varepsilon n - 4\varepsilon \\ &= 8\varepsilon n - 4\varepsilon + \varepsilon^2 - \mathrm{Avg}[(A^* - \Pi_{t,\gamma'} B_{t,\boldsymbol{p}})^2] \\ \mathrm{Avg}[(A^* - \Pi_{t,\gamma'} B_{t,\boldsymbol{p}})^2] &\leq 8\varepsilon n \ . \end{aligned}$$

□

5.5 Uniform Affordability in the Utility-Maximizing Clause

In this section, we formally show that the average affordability of items in $Z_{t,r}$ and $Z_{t,r'}$ are roughly the same when $\mathcal{K}_{|\gamma^*|-1}(r) = \mathcal{K}_{|\gamma^*|-1}(r')$.

Proposition 5.10. *Suppose $\mathcal{E}$ holds and $W_t(\boldsymbol{p}) \geq |Z_t|$. Let (t,i) be a relevant block-visit pair where the mechanism posts prices $\boldsymbol{p} \in [0,\infty]^{Z_t}$. Let $\gamma^* = (\tau_t(i), L^*)$, A^* define a utility-maximizing clause. Then for all $k, r, r' \in [\kappa]$ such that $\mathcal{K}_{|\gamma^*|-1}(r) = \mathcal{K}_{|\gamma^*|-1}(r')$,*

$$\left| \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r'}] \right| \leq \lambda^{-2} \ .$$

Proof. Note that $L^* = \tau_t(I \cup i)$ for some $I \subseteq [i-1]$ by Proposition 5.5. By triangle inequality,

$$\left| \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \mathrm{Avg}[A^*(k,\cdot) \mid Z_{t,r,I}] \right|$$

is bounded by

$$\begin{aligned} &\left| \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[i-1]}] \right| \\ &\quad + \left| \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[i-1]}] - \mathrm{Avg}[A^*(k,\cdot) \mid Z_{t,r,[i-1]}] \right| \\ &\quad + \left| \mathrm{Avg}[A^*(k,\cdot) \mid Z_{t,r,[i-1]}] - \mathrm{Avg}[A^*(k,\cdot) \mid Z_{t,r,I}] \right| \ . \end{aligned} \tag{2}$$

$\mathcal{K}_{|I|}$ refines $\mathcal{K}_s$ for all $s \leq |I|$, and $\mathcal{K}_{|I|}(r) = \mathcal{K}_{|I|}(r')$. Therefore, the definitions of $Z_{t,r,I}, Z_{t,r',I}$ are identical, and we have $Z_{t,r,I} = Z_{t,r',I}$. It remains to bound the differences in (2).

Claim 5.11. *Suppose $\mathcal{E}$ holds. Then*

$$\left| \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[i-1]}] \right| \leq 3\varepsilon \ .$$

Proof Sketch. Consider a fixed B and τ_t. Condition on $\{\psi'_{t,\gamma} : \gamma \subseteq \tau_t([i-1])\}$, which fully determines $Z_{t,r,[i-1]}$. Then since each item in $Z_{t,r,[i-1]}$ is independently included in $Z_{t,r}$ w.p. $1/(\vartheta_{t,r,[i-1]}\,\mathrm{val}(r))$, we expect $\mathrm{Avg}[B(k,\cdot) \mid Z_{t,r}]$ to concentrate.

We face the same issue as in Claims 5.6 and 5.8, where the mechanism can see the randomness used to decide $Z_{t,r,[n]}$ (and recall that $Z_{t,r} \subseteq Z_{t,r,[n]} \subseteq Z_{t,r,[i-1]}$). As before, the solution is to observe that $Z_{t,r,[n]}$ depends on $\tau_t([i,n])$, which the mechanism does not know. Therefore, we show that for every $B_{t,\boldsymbol{p}}$, there exists a small set $T \subseteq [\eta]$ such that τ_t has a large deviation only if $\tau_t([i,n]) \cap T \neq \emptyset$. See Proposition A.3-(3) for details. □

Lemma 5.12. *Suppose $\mathcal{E}$ holds and $W_t(\boldsymbol{p}) \geq |Z_t|$. Then*

$$\left| \mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[i-1]}] - \mathrm{Avg}[A^*(k,\cdot) \mid Z_{t,r,[i-1]}] \right| \leq \lambda^{-3} \ .$$

Proof. Since $Z_{t,r,[i-1]}$ is a union of parts from $\mathcal{P}_{t,\tau_t([i-1])}$, we have

$$\operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[i-1]}] \quad = \quad \operatorname{Avg}[(\Pi_{t,\tau_t([i-1])}B_{t,\boldsymbol{p}})(k,\cdot) \mid Z_{t,r,[i-1]}] \ .$$

Also, we have $|Z_{t,r,[i-1]}| \geq |Z_t|/\lambda$ by standard concentration. Formally, this is covered under $\mathcal{E}$; see Proposition A.2-(2). Thus, by Jensen's Inequality and Proposition 5.5,

$$\begin{aligned}
\left|\operatorname{Avg}[(A^* - \Pi_{t,\tau_t([i-1])}B_{t,\boldsymbol{p}})(k,\cdot) \mid Z_{t,r,[i-1]}]\right| &\leq \operatorname{Avg}[|(A^* - \Pi_{t,\tau_t([i-1])}B_{t,\boldsymbol{p}})(k,\cdot)| \mid Z_{t,r,[i-1]}] \\
&\leq (\kappa|Z_t|/|Z_{t,r,[i-1]}|)\operatorname{Avg}[|A^* - \Pi_{t,\tau_t([i-1])}B_{t,\boldsymbol{p}}|] \\
&\leq \kappa\lambda \operatorname{Avg}[|A^* - \Pi_{t,\tau_t([i-1])}B_{t,\boldsymbol{p}}|] \\
&\leq \kappa\lambda\sqrt{\operatorname{Avg}[(A^* - \Pi_{t,\tau_t([i-1])}B_{t,\boldsymbol{p}})^2]} \\
&\leq \kappa\lambda\sqrt{8\varepsilon n} \\
&\leq \lambda^{-3} \ . \qquad \square
\end{aligned}$$

Claim 5.13. *Suppose $\mathcal{E}$ holds. Then*

$$\left|\operatorname{Avg}[A^*(k,\cdot) \mid Z_{t,r,[i-1]}] - \operatorname{Avg}[A^*(k,\cdot) \mid Z_{t,r,I}]\right| \quad \leq \quad \varepsilon \ .$$

Proof. The idea is the same as in Claim 5.11. Condition on $\{\psi'_{t,\gamma} : \gamma \subseteq \tau_t(I)\}$, which fixes $Z_{t,r,I}$ and makes $Z_{t,r,[i-1]}$ a random subset. Luckily, since A^* is $\mathcal{P}_{t,\tau_t(I)}$-measurable, we can directly union bound over all possibilities for A^*. See Proposition A.2-(3) for details. $\square$

Combining the bounds of Claim 5.11, Lemma 5.12, Claim 5.13 with (2) and $Z_{t,r,I} = Z_{t,r',I}$ yields

$$\left|\operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r'}]\right| \quad \leq \quad 2(3\varepsilon + \lambda^{-3} + \varepsilon) \quad \leq \quad \lambda^{-2} \ . \qquad \square$$

5.6 Owner Affordability Allows Stealing

For a relevant block-visit pair (t,i) and prices $\boldsymbol{p} \in [0,\infty]^{Z_t}$, let $S_{t,i}(\boldsymbol{p})$ be the output utility-maximizing set for $v_{t,i}$. Define

$$V_{t,i}(\boldsymbol{p}) \quad = \quad w_t(S_{t,i}(\boldsymbol{p}))$$

to be the amount of owner value destroyed by the non-owner valuation.

Lemma 5.14. *Suppose $\mathcal{E}$ holds. Then for all relevant block-visit pairs (t,i) where the mechanism posts prices $\boldsymbol{p} \in [0,\infty]^{Z_t}$,*

$$W_t(\boldsymbol{p}) \quad \leq \quad 5\,V_{t,i}(\boldsymbol{p}) + \kappa|Z_t|/n^2 \ .$$

Proof. Fix a block-visit pair (t,i). If $W_t(\boldsymbol{p}) \leq |Z_t|$, we are done because $V_{t,i}(\boldsymbol{p}) \geq 0$ and $\kappa/n^2 \geq 1$. Otherwise, let $\gamma^* = (\tau_t(i), \tau_t(I \cup i)), A^*$ be a utility-maximizing clause supporting $S_{t,i}(\boldsymbol{p})$, where $I \subseteq [i-1]$ by Proposition 5.5. Define $f, g : [\kappa] \to [0,1]$ to be the fractions of items from each rank affordable to the owner or destroyed by a non-owner, i.e.,

$$f(k) \quad = \quad \frac{|\{x \in Z_{t,k} : p_x < \operatorname{val}(k)\}|}{|Z_{t,k}|} \qquad \text{and} \qquad g(r) \quad = \quad \frac{|S_{t,i}(\boldsymbol{p}) \cap Z_{t,r}|}{|Z_{t,r}|} \ ,$$

where the denominators are nonzero by $\mathcal{E}$; see Proposition A.2-(1).

Fix $K \in \mathcal{K}_{|I|}$ and a consecutive pair $C, D \in \mathcal{L}(K)$. For $k \in C$ and $r \in D$, because $\gamma^* \subseteq \tau_t([i])$,

$$x \in Z_{t,r} \subseteq Z_{t,r,[i]} \quad \Longrightarrow \quad \psi'_{t,\gamma^*}(x) \geq \varphi_{|I|+1}(r) = \min D - 1 \geq k \ .$$

Thus, the utility-maximizing clause assigns every item in $Z_{t,r}$ value at least $2\,\mathrm{val}(k)$, so an item $x \in Z_{t,r}$ is in $S_{t,i}(\boldsymbol{p})$ when $B_{t,\boldsymbol{p}}(k,x) > 0$. On the other hand, an item $x \in Z_{t,k}$ is affordable to the owner when $B_{t,\boldsymbol{p}}(k,x) > 1/2$. Therefore, since k, r satisfy $\mathcal{K}_{|I|}(k) = \mathcal{K}_{|I|}(r)$, Proposition 5.10 yields

$$f(k) \leq 2\,\mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,k}] \leq 2\,\mathrm{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] + \lambda^{-1} \leq 2\,g(r) + \lambda^{-1} \ .$$

Averaging over $k \in C$ and $r \in D$, we have $\mathrm{Avg}[f \mid C] \leq 2\,\mathrm{Avg}[g \mid D] + \lambda^{-1}$. Applying Lemma 3.15,

$$\sum_{k\in[\kappa]} f(k) \leq 4 \sum_{k\in[\kappa]} g(k) + \frac{\kappa}{\lambda} + \kappa^{2/3} \ . \tag{3}$$

$\mathcal{E}$ guarantees $|\mathrm{val}(k)|Z_{t,k}| - |Z_t|| \leq \varepsilon |Z_t|$; see Proposition A.2-(1). Therefore, we have

$$\begin{aligned} W_t(\boldsymbol{p}) &= \sum_{k\in[\kappa]} \Big(f(k)\,\mathrm{val}(k)|Z_{t,k}| \Big) \leq (1+\varepsilon)|Z_t| \sum_{k\in[\kappa]} f(k) \\ V_{t,i}(\boldsymbol{p}) &= \sum_{k\in[\kappa]} \Big(g(k)\,\mathrm{val}(k)|Z_{t,k}| \Big) \geq (1-\varepsilon)|Z_t| \sum_{k\in[\kappa]} g(k) \ . \end{aligned}$$

Thus, by (3),

$$W_t(\boldsymbol{p}) \leq 4(1+\varepsilon)|Z_t| \sum_{k\in[\kappa]} g(k) + (1+\varepsilon)(\kappa/\lambda + \kappa^{2/3})|Z_t| \leq 5\,V_{t,i}(\boldsymbol{p}) + \frac{\kappa|Z_t|}{n^2} \ . \qquad \square$$

5.7 Wrapping Up

We begin by showing the non-owner welfare is negligible.

Lemma 5.15.

$$\mathbb{E}\left[\sum_{t,i\in[n]} v_{t,i}(Z_t) \right] \leq \frac{\kappa|Z_1|}{n} \ .$$

Proof. Fix any clause (t, γ, A), and define the random variable $Y_x = 2\,\mathrm{val}(\psi_{t,\gamma}(x)) \in [0, 2\lambda]$. Then

$$\boldsymbol{u}^{(t,\gamma,A)}(Z_t) + u_0^{(t,\gamma,A)} \leq \sum_{x\in Z_t} Y_x + 2|Z_t| \ .$$

Since the Y_x are independent with $\mathbb{E}[Y_x] = \sum_{k\in[\kappa]}(2\,\mathrm{val}(k) \cdot 2/(\kappa\,\mathrm{val}(k))) = 4$, Hoeffding yields

$$\Pr\left[\boldsymbol{u}^{(t,\gamma,A)}(Z_t) + u_0^{(t,\gamma,A)} \geq 7|Z_t| \right] \leq \Pr\left[\sum_{x\in Z_t} \Big(Y_x - \mathbb{E}[Y_x] \Big) \geq |Z_t| \right] \leq \exp(-\lambda^2) \ .$$

Taking a union bound over all at most $n\eta^n 2^\lambda \leq \exp(\lambda)$ clauses yields that

$$\Pr\left[\max_{t,i\in[n]} v_{t,i}(Z_t) \geq 7|Z_1| \right] \leq \lambda^{-1} \ .$$

Therefore, the expected non-owner welfare is bounded by

$$\mathbb{E}\left[\sum_{t,i\in[n]} v_{t,i}(Z_t) \right] \leq 7n^2|Z_1| + 3n^2\lambda|Z_1| \Pr\left[\max_{t,i\in[n]} v_{t,i}(Z_t) \geq 7|Z_1| \right] \leq \frac{\kappa|Z_1|}{n} \ . \qquad \square$$

Next, we bound the owner welfare.

Lemma 5.16. *Let $S_1, \dots, S_n$ be the allocation by the mechanism. Then*

$$\mathbb{E}\left[\sum_{i\in[n]} w_{\pi(i)}(S_i \cap Z_{\pi(i)})\right] \quad \le \quad \left(5+\frac{2}{n}\right)\kappa|Z_1| \ .$$

Proof. For $\boldsymbol{p} \in [0,\infty]^m$, we abuse notation and denote $W_t(\boldsymbol{p}) = W_t(\boldsymbol{p}|_{Z_t})$ and $V_{t,i}(\boldsymbol{p}) = V_{t,i}(\boldsymbol{p}|_{Z_t})$. Let $\boldsymbol{p}_1, \dots, \boldsymbol{p}_n \in [0,\infty]^m$ be the prices posted by the mechanism at each visit, and fix a visit $i < n$. Consider the instance where, at the beginning, we exchange the values $\pi(i)$ and $\pi(n)$. This preserves the distribution of the instance and all information revealed before visit i, and therefore the prices $\boldsymbol{p}_i$ posted by the mechanism are identical in the original and the swapped instances. Hence,

$$\mathbb{E}[W_{\pi(i)}(\boldsymbol{p}_i)] \quad = \quad \mathbb{E}[W_{\pi(n)}(\boldsymbol{p}_i)] \ .$$

Since $W_t(\boldsymbol{p}) \le \lambda|Z_1|$ always, Lemma 5.14 yields for all $i \in [n]$ that

$$W_{\pi(n)}(\boldsymbol{p}_i) \quad \le \quad 5\,V_{\pi(n),i}(\boldsymbol{p}_i) + \frac{\kappa|Z_1|}{n^2} + \lambda|Z_1|\,\mathbb{1}(\mathcal{E}^c) \ .$$

Chaining the above relations and using $\Pr[\mathcal{E}^c] \le \lambda^{-1}$ by Proposition A.3,

$$\begin{aligned}
\mathbb{E}\left[\sum_{i\in[n]} w_{\pi(i)}(S_i \cap Z_{\pi(i)})\right] &= \sum_{i\in[n]} \mathbb{E}[W_{\pi(i)}(\boldsymbol{p}_i)] \\
&= \sum_{i\in[n]} \mathbb{E}[W_{\pi(n)}(\boldsymbol{p}_i)] \\
&\le \sum_{i\in[n-1]} \mathbb{E}\left[5\,V_{\pi(n),i}(\boldsymbol{p}_i) + \frac{\kappa|Z_1|}{n^2} + \lambda|Z_1|\,\mathbb{1}(\mathcal{E}^c)\right] + 5\,\mathbb{E}[W_{\pi(n)}(\boldsymbol{p}_n)] \\
&\le 5\,\mathbb{E}\left[\sum_{i\in[n]} w_{\pi(n)}(S_i \cap Z_{\pi(n)})\right] + \frac{\kappa|Z_1|}{n} + n|Z_1| \\
&\le 5\,\mathbb{E}[w_{\pi(n)}(Z_{\pi(n)})] + \frac{\kappa|Z_1|}{n} + n|Z_1| \\
&= 5\kappa|Z_1| + \frac{\kappa|Z_1|}{n} + n|Z_1| \\
&\le \left(5+\frac{2}{n}\right)\kappa|Z_1| \ ,
\end{aligned}$$

where the fifth line comes from additivity of $w_{\pi(n)}$ and feasibility of the allocation. □

Proof of Theorem 2.2. Lemmas 5.15 and 5.16 upper bound the expected welfare of the mechanism by $(5+3/n)\kappa|Z_1|$, while the expected optimal welfare is at least $n\kappa|Z_1|$. Thus, the approximation ratio is lower bounded by

$$\frac{n}{5+3/n} \quad \ge \quad \frac{n}{6} \quad = \quad \Theta(\log\log\log m) \ .$$ □

References

[AKS21] Sepehr Assadi, Thomas Kesselheim, and Sahil Singla. Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 653–661. SIAM, 2021. 1, 2

[AKSW20] Sepehr Assadi, Hrishikesh Khandeparkar, Raghuvansh R. Saxena, and S. Matthew Weinberg. Separating the communication complexity of truthful and non-truthful combinatorial auctions. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 1073–1085. ACM, 2020. 2

[AS19] Sepehr Assadi and Sahil Singla. Exponentially improved truthful combinatorial auctions with submodular bidders. In *Proceedings of the Sixtieth Annual IEEE Foundations of Computer Science (FOCS)*, 2019. 1, 2

[BM15] Rémi Bardenet and Odalric-Ambrym Maillard. Concentration inequalities for sampling without replacement. *Bernoulli*, 21(3):1361–1385, August 2015. 6

[CC23] José Correa and Andrés Cristi. A constant factor prophet inequality for online combinatorial auctions. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 686–697. ACM, 2023. 2

[CD26] Shuchi Chawla and Trung Dang. Single-sample prophet inequalities: A combinatorial to single-item reduction, 2026. 2

[Cla71] Edward H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11(1):17–33, 1971. 1

[DFKL17] Paul Duetting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 540–551, 2017. https://arxiv.org/pdf/1612.03161.pdf. 2

[DKL20] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. An o(log log m) prophet inequality for subadditive combinatorial auctions. *SIGecom Exch.*, 18(2):32–37, 2020. 2

[DKL+24] Paul Dütting, Thomas Kesselheim, Brendan Lucier, Rebecca Reiffenhäuser, and Sahil Singla. Online combinatorial allocations and auctions with few samples. In *65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024*, pages 1231–1250. IEEE, 2024. 1, 2

[DNS06] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful randomized mechanisms for combinatorial auctions. In Jon M. Kleinberg, editor, *Proceedings of the 38th Annual ACM Symposium on Theory of Computing, Seattle, WA, USA, May 21-23, 2006*, pages 644–652. ACM, 2006. 1, 2

[DNS10] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. *Math. Oper. Res.*, 35(1):1–13, 2010. https://www.cs.yale.edu/homes/schapira/cfjournal.pdf. 2

[Dob07] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *Proceedings of the 10th International Workshop on Approximation and the 11th International Workshop on Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 89–103, 2007. 1, 2

[Dob11] Shahar Dobzinski. An Impossibility Result for Truthful Combinatorial Auctions with Submodular Valuations. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, 2011. 2

[Dob16] Shahar Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2016, pages 940–948, New York, NY, USA, 2016. ACM. 1, 2

[DV11] Shaddin Dughmi and Jan Vondrak. Limitations of Randomized Mechanisms for Combinatorial Auctions. In *52nd Annual Symposium on Foundations of Computer Science (FOCS)*, 2011. 2

[DV12a] Shahar Dobzinski and Jan Vondrák. From query complexity to computational complexity. In *Proceedings of the 44th Symposium on Theory of Computing (STOC)*, 2012. 2

[DV12b] Shahar Dobzinski and Jan Vondrak. The Computational Complexity of Truthfulness in Combinatorial Auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 2012. 2

[DV16] Shahar Dobzinski and Jan Vondrák. Impossibility results for truthful combinatorial auctions with submodular valuations. *J. ACM*, 63(1):5:1–5:19, 2016. 2

[EHKS18] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 700–714, 2018. 2

[FGL15] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135. SIAM, 2015. 2

[Gro73] Theodore Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973. 1

[HKMT04] Ron Holzman, Noa E. Kfir-Dahav, Dov Monderer, and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, 47(1):104–123, 2004. 2

[KRTV13] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In Hans L. Bodlaender and Giuseppe F. Italiano, editors, *Algorithms - ESA*

2013 - 21st Annual European Symposium, Sophia Antipolis, France, September 2-4, 2013. Proceedings, volume 8125 of *Lecture Notes in Computer Science*, pages 589–600. Springer, 2013. 1, 2

[KV12] Piotr Krysta and Berthold Vöcking. Online mechanism design (randomized rounding on the fly). In *Automata, Languages, and Programming*, pages 636–647. Springer, 2012. 1, 2

[MNP+26] Calum MacRury, Rian Neogi, Kanstantsin Pashkovich, Sahil Singla, Siddharth M Sundaram, and Chaitanya Swamy. Efficient algorithms for online subadditive combinatorial allocations, 2026. 2

[QW24] Frederick V. Qiu and S. Matthew Weinberg. Settling the communication complexity of vcg-based mechanisms for all approximation guarantees. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 1192–1203. ACM, 2024. 2

[RS27] Aviad Rubinstein and Sahil Singla. Improved combinatorial auctions with and without samples. In *Proceedings of the 2027 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2027. To appear. 1, 2

[RTWZ24] Shiri Ron, Clayton Thomas, S. Matthew Weinberg, and Qianfan Zhang. Communication separations for truthful auctions: Breaking the two-player barrier. *CoRR*, abs/2409.08241, 2024. 2

[Vic61] William Vickrey. Counterspeculations, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961. 1

A The Full-Revelation Good Event

Lemma A.1. *$|\mathcal{A}_{t,L}| \le 2^\lambda$ for all $t \in [n]$ and $L \subseteq [\eta]$ with $|L| \le n$.*

Proof. By a direct computation,

$$|\mathcal{A}_{t,L}| \quad \le \quad (1+\varepsilon^{-1})^{\kappa|\mathcal{P}_{t,L}|} \quad \le \quad (1+\varepsilon^{-1})^{\kappa(1+\kappa)^{2^n n}} \quad \le \quad \lambda^{\kappa^\kappa} \quad \le \quad 2^\lambda \ . \qquad \square$$

Proposition A.2. *Simultaneously w.p. at least $1-\lambda^{-2}$:*

1. *For every block $t \in [n]$ and rank $k \in [\kappa]$,*

$$\Big|\mathrm{val}(k)|Z_{t,k}| - |Z_t|\Big| \quad \le \quad \varepsilon|Z_t| \ .$$

2. *For every block $t \in [n]$, rank $k \in [\kappa]$, and bidder set $I \subseteq [n]$,*

$$|Z_{t,k,I}| \quad \ge \quad |Z_t|/\lambda \ .$$

3. *For every block $t \in [n]$, ranks $k, r \in [\kappa]$, bidder sets $I \subseteq J \subseteq [n]$, and $A \in \mathcal{A}_{t,\tau_t(I)}$,*

$$\Big|\mathrm{Avg}[A(k,\cdot) \mid Z_{t,r,I}] - \mathrm{Avg}[A(k,\cdot) \mid Z_{t,r,J}]\Big| \quad \le \quad \varepsilon \ .$$

Proof. Each item is in $Z_{t,k}$ independently w.p. $1/\operatorname{val}(k)$, so (1) holds for a fixed choice w.p. at least $1-\exp(-\lambda)$ by Chernoff.

By Lemma 5.1, each item is in $Z_{t,k,I}$ independently w.p. at least $1/\sqrt{\operatorname{val}(k)} \geq \lambda^{-1/2}$, so (2) holds for a fixed choice w.p. at least $1-\exp(-\lambda)$ by Chernoff.

Now, fix a choice of t, k, r, I, J. Condition on τ_t and $\{\psi'_{t,\gamma} : \gamma \subseteq \tau_t(I)\}$ (denote this conditioning by $\mathcal{F}$), which fixes $Z_{t,r,I}$ and $\mathcal{P}_{t,\tau_t(I)}$, allowing us to fix a choice of A. Observe that each item in $Z_{t,r,I}$ is independently included in $Z_{t,r,J}$ w.p. $\vartheta_{t,r,J}/\vartheta_{t,r,I}$. Therefore, by Corollary 3.10,

$$\begin{aligned}
&\Pr\left[\left|\operatorname{Avg}[A(k,\cdot) \mid Z_{t,r,I}] - \operatorname{Avg}[A(k,\cdot) \mid Z_{t,r,J}]\right| \geq \varepsilon \;\middle|\; \mathcal{F}\right] \\
&\qquad \leq \Pr\left[|Z_{t,r,J}| \geq |Z_t|/\lambda,\ \left|\operatorname{Avg}[A(k,\cdot) \mid Z_{t,r,I}] - \operatorname{Avg}[A(k,\cdot) \mid Z_{t,r,J}]\right| \geq \varepsilon \;\middle|\; \mathcal{F}\right] \\
&\qquad\qquad + \Pr[|Z_{t,r,J}| \leq |Z_t|/\lambda \mid \mathcal{F}] \\
&\qquad \leq \Pr[|Z_{t,r,J}| \leq |Z_t|/\lambda \mid \mathcal{F}] + 2\exp(-2\varepsilon^2|Z_t|/\lambda) \\
&\qquad \leq \Pr[|Z_{t,r,J}| \leq |Z_t|/\lambda \mid \mathcal{F}] + \exp(-\lambda^2) \ .
\end{aligned}$$

Union bound over all 2^λ choices of A, then average over the conditioning to obtain the bound $2^\lambda(\Pr[|Z_{t,r,J}| \leq |Z_t|/\lambda] + \exp(-\lambda))$. A union bound over all at most $n\kappa + n\kappa \cdot 2^n + n\kappa^2 \cdot 4^n \leq \lambda$ choices completes the proof of all three statements. □

Proposition A.3. *Simultaneously w.p. at least $1-\lambda^{-1}$: the event of Proposition A.2 holds, and for all relevant block-visit pairs (t,i), the prices $\boldsymbol{p} \in [0,\infty]^{|Z_t|}$ posted by the mechanism satisfy*

1. *For every $\gamma \in \Gamma_{t,i}$ and $A \in \mathcal{A}_{t,\gamma}$,*

$$\left|U_t(\gamma, A, B_{t,\boldsymbol{p}}) - \left(2\operatorname{Avg}[B_{t,\boldsymbol{p}}] + 2\operatorname{Avg}[AB_{t,\boldsymbol{p}}]\right)\right| \quad \leq \quad 2\varepsilon \ .$$

2. *For every $L \subseteq N_{t,i} \setminus \tau_t(i)$ with $|L| \leq n-1$,*

$$\Delta_t(L \cap \tau_t([i-1]), B_{t,\boldsymbol{p}}) - \Delta_t(L, B_{t,\boldsymbol{p}}) \quad \leq \quad 2\varepsilon \ .$$

3. *For every $k, r \in [\kappa]$,*

$$\left|\operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[i-1]}]\right| \quad \leq \quad 3\varepsilon$$

Proof. Fix a relevant block-visit pair (t,i).

The challenge is that we have no way of knowing what price vector $\boldsymbol{p}$ will be posted by the mechanism, and conditioning on $\boldsymbol{p}$ may wildly alter the distribution of the instance (since the mechanism can see $\{\psi_{t,\gamma}, \xi_{t,\gamma} : \gamma \in \Gamma\}$). We will instead prove that there does not exist *any* "good" $B : [\kappa] \times Z_t \to [0,1]$, and hence no prices posted by the mechanism can violate the above properties with non-negligible probability. In particular, we show that no B can cause a large deviation for many realizations simultaneously, and hence the probability that the posted price is good for the actual realization is low. We first formalize what we mean by "actual realization."

Lemma A.4. *Let $s \geq 1$, and let $T \subseteq [\eta] \setminus \tau_t([i-1])$ be a set of size at most s which is fully determined by the mechanism's information prior to the response from bidder i. Then*

$$\Pr[T \cap N_{t,i} \neq \emptyset] \quad \leq \quad \frac{3s}{m} \ .$$

Proof. Condition on π, the randomness of the other blocks, and $\{\psi_{t,\gamma}, \xi_{t,\gamma} : \gamma \in \Gamma\}$. Denote this information by $\mathcal{H}$. Then since the owner of Z_t has not yet been visited, the mechanism's behavior (conditioned on $\mathcal{H}$) is solely a function of the names and name lists $\{\tau_t(j), N_{t,j} : j \in [i-1]\}$ disclosed after each visit. Further, τ_t and $\{N_{t,j} : j \in [n]\}$ are independent from $\mathcal{H}$. Consider a potential injection $\tau : [n] \to [\eta]$. Say that τ is *consistent* if τ agrees with τ_t on $[i-1]$, and $\tau([i,n]) \subseteq \bigcap_{j\in[i-1]} N_{t,j} \setminus \tau_t([i-1]) =: N$, where the empty intersection is defined to be $[\eta]$. Then the likelihood of observing the names and name lists conditioned on $\tau_t = \tau$ is

$$\begin{aligned}&\Pr[\{\tau_t(j), N_{t,j} : j \in [i-1]\} \mid \tau_t = \tau, \mathcal{H}] \\ &\qquad = \quad \mathbb{1}(\tau \text{ is consistent}) \prod_{j\in[i-1]} \left[\left(\frac{1}{m}\right)^{|N_{t,j}|-n}\left(1-\frac{1}{m}\right)^{\eta-|N_{t,j}|}\right] .\end{aligned}$$

In particular, since τ_t is independent from $\mathcal{H}$, τ_t is drawn uniformly, and the likelihood above depends on τ solely through whether τ is consistent, Bayes' rule yields

$$\Pr[\tau_t = \tau \mid \{\tau_t(j), N_{t,j} : j \in [i-1]\}, \mathcal{H}] \quad \propto \quad \mathbb{1}(\tau \text{ is consistent}) .$$

Hence, τ_t is uniformly distributed among all consistent τ, and for all $j \in [i,n]$,

$$\Pr[\tau_t(j) = \ell \mid \{\tau_t(j), N_{t,j} : j \in [i-1]\}, \mathcal{H}] \quad = \quad \frac{\mathbb{1}(\ell \in N)}{|N|} .$$

Now, observe that T is fixed, conditioned on $\{\tau_t(j), N_{t,j} : j \in [i-1]\}$ and $\mathcal{H}$. The probability that a name $\ell \in T$ is in $N_{t,i}$ is 1 if $\ell \in \tau_t([i,n])$ and $1/m$ otherwise. Hence, by a union bound,

$$\begin{aligned}&\Pr[T \cap N_{t,i} \neq \emptyset \mid \{\tau_t(j), N_{t,j} : j \in [i-1]\}, \mathcal{H}] \\ &\qquad \leq \quad \sum_{\ell\in T}\left(\frac{1}{m} + \sum_{j=i}^{n} \Pr[\tau_t(j) = \ell \mid \{\tau_t(j), N_{t,j} : j \in [i-1]\}, \mathcal{H}]\right) \\ &\qquad \leq \quad s\left(\frac{1}{m} + \frac{n}{|N|}\right) .\end{aligned}$$

We conclude by observing that each name $\ell \notin \tau_t([n])$ is included in N w.p. $m^{-(i-1)}$, so $|N| - (n - i + 1) \sim \text{Binom}(m^{2n} - n, m^{-(i-1)})$. A Chernoff bound then yields $\Pr[|N| \leq nm] \leq 1/m$. Union bound over this event and deconditioning completes the proof. □

The remainder of the proof goes as follows: we show that for any choice of $B : [\kappa] \times Z_t \to [0,1]$, there exists a small set T_B such that the desired properties are violated only if $N_{t,i} \cap T_B \neq \emptyset$. In other words, a price vector can only be "good" for a small information set, so because bidder i's name list is sparse, the price vector is unlikely to be good for bidder i. We first prove that the desired properties hold with extremely high probability for a single B, and wrap up everything at the end with a union bound over an m^{-1}-net over $[\kappa] \times Z_t \to [0,1]$.

Lemma A.5. *For all $B : [\kappa] \times Z_t \to \{0, m^{-1}, 2m^{-1}, \ldots, 1\}$, define $T_B \subseteq [\eta]$ to be a minimal set satisfying: for all $\gamma = (\ell, L)$ where $\ell \notin T_B$, and for all $A \in \mathcal{A}_{t,\gamma}$,*

$$\left|U_t(\gamma, A, B) - \Big(2\,\text{Avg}[B] + 2\,\text{Avg}[AB]\Big)\right| \quad \leq \quad \varepsilon .$$

Then $|T_B| \leq \lambda^{30}$ for all B simultaneously w.p. at least $1 - \lambda^{-2}$.

Proof. Fix a choice of B. Suppose that the smallest set T_B were of size more than λ^{30}. Then consider the graph on T_B such that ℓ, ℓ' are connected when $\ell \in L_{\ell'}$ or $\ell' \in L_\ell$. Since $|L_\ell| \leq n$, the average degree is at most $2n$, so there exists an independent set $T'_B \subseteq T_B$ of size λ^{29} (at least half the vertices have degree at most $4n$, each such vertex only precludes at most $4n$ other vertices). Therefore, it suffices to bound the probability that there are violations on the λ^{29} witnesses $\{(\gamma_\ell = (\ell, L_\ell), A_\ell) : \ell \in T'_B\}$. Fix such a set of witnesses.

Condition on $\{\psi_{t,\gamma}, \xi_{t,\gamma} : \gamma \subseteq \bigcup_{\ell \in T'_B}(L_\ell \setminus \ell)\}$. Observe that this conditions on $\mathcal{P}_{t,L_\ell \setminus \ell}$ for all $\ell \in T'_B$ and hence fixes all A_ℓ, but our choice of T'_B means that $\{\psi_{t,\gamma_\ell}, \xi_{t,\gamma_\ell} : \ell \in T'_B\}$ remain unaffected by the conditioning and mutually independent.

Fix any $\ell \in T'_B$, and define

$$U_{t,x}(\gamma_\ell, A_\ell, B) \quad = \quad \sum_{k \in [\kappa]} \Big(2\,\mathrm{val}(k)\, \mathbb{1}(\psi_{t,\gamma_\ell}(x) = k)\, \mathbb{1}(\xi_{t,\gamma_\ell}(x) \leq 1 + A_\ell(k,x))\, B(k,x) \Big) \ .$$

Observe that the random variables $U_{t,x}(\gamma_\ell, A_\ell, B)$ are mutually independent and lie in $[0, 2\lambda]$. Also, we have

$$U_t(\gamma_\ell, A_\ell, B) \quad = \quad \frac{1}{|Z_t|} \sum_{x \in Z_t} U_{t,x}(\gamma_\ell, A_\ell, B)$$

and

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{|Z_t|} \sum_{x \in Z_t} U_{t,x}(\gamma_\ell, A_\ell, B) \right] &= \frac{1}{|Z_t|} \sum_{x \in Z_t} \sum_{k \in [\kappa]} \left(2\,\mathrm{val}(k) \frac{2}{\kappa\,\mathrm{val}(k)} \frac{1 + A_\ell(k,x)}{2} B(k,x) \right) \\
&= \frac{2}{\kappa |Z_t|} \sum_{x \in Z_t} \sum_{k \in [\kappa]} \Big(B(k,x) + A_\ell(k,x)\, B(k,x) \Big) \\
&= 2\,\mathrm{Avg}[B] + 2\,\mathrm{Avg}[A_\ell B] \ .
\end{aligned}$$

Therefore, by Hoeffding,

$$\begin{aligned}
\Pr\left[\left| U_t(\gamma_\ell, A_\ell, B) - \Big(2\,\mathrm{Avg}[B] + 2\,\mathrm{Avg}[A_\ell B] \Big) \right| \geq \varepsilon \right] &\leq 2\exp(-2\varepsilon^2 |Z_t|/(2\lambda)^2) \\
&\leq \exp(-\lambda^{77}) \ .
\end{aligned}$$

Thus, by independence, the probability that this violation occurs for all $\ell \in T'_B$ is at most $\exp(-\lambda^{106})$. Deconditioning and union bounding over $(\eta^n 2^\lambda)^{\lambda^{29}} (1+m)^{\kappa\lambda^{100}} \leq \exp(\lambda^{101})$ choices of $\{(\gamma_\ell, A_\ell) : \ell \in T'_B\}$ and B completes the proof. □

Lemma A.6. *For all $B : [\kappa] \times Z_t \to \{0, m^{-1}, 2m^{-1}, \ldots, 1\}$ and $K \subseteq [\eta]$ with $|K| \leq n-1$, define $T_{B,K} \subseteq [\eta]$ to be a minimal set satisfying: for all $L \subseteq [\eta] \setminus K$ with $|K \cup L| \leq n-1$ and $L \cap T_{B,K} = \emptyset$,*

$$\Delta_t(K, B) - \Delta_t(K \cup L, B) \quad \leq \quad \varepsilon \ .$$

Then $|T_{B,K}| \leq \lambda^{30}$ for all B, K simultaneously w.p. at least $1 - \lambda^{-2}$.

Proof. Fix B and K. Let $\mathcal{T}_{B,K}$ be the collection of L witnessing a violation, and suppose that the smallest set $T_{B,K}$ were of size more than λ^{30}. Then consider a maximal choice of disjoint sets $\mathcal{T}'_{B,K} \subseteq \mathcal{T}_{B,K}$. We must have

$$\left| \bigcup_{L \in \mathcal{T}'_{B,K}} L \right| \quad \geq \quad \lambda^{30} \ ,$$

since by maximality of $\mathcal{T}'_{B,K}$, every $L \in \mathcal{T}_{B,K}$ intersects the above object. Further, since each member set is of size at most n, $|\mathcal{T}'_{B,K}| \geq \lambda^{29}$. Therefore, it suffices to bound the probability that there are violations on a set of λ^{29} disjoint witnesses $\mathcal{T}'_{B,K}$. Fix such a set of witnesses.

Condition on $\{\psi_{t,\gamma}, \xi_{t,\gamma} : \gamma \subseteq K\}$. This fixes the partition $\mathcal{P}_{t,K}$ and $\Delta_t(K, B)$, while leaving the partitions $\{\mathcal{P}_{t,K\cup L} : L \in \mathcal{T}'_{B,K}\}$ conditionally mutually independent.

Fix any $L \in \mathcal{T}'_{B,K}$, and choose a function $f : [\kappa] \times \{0, 1, \ldots, \kappa\}^{\{\gamma \subseteq K\cup L\}} \to \{0, \varepsilon, 2\varepsilon, \ldots, 1\}$. Observe that for every realization of $\mathcal{P}_{t,K\cup L}$, the choices of f exhaust the choices of $A \in \mathcal{A}_{t,K\cup L}$, since f chooses a value for every rank and equivalence class under $\mathcal{P}_{t,K\cup L}$. Let $A_f \in \mathcal{A}_{t,K\cup L}$ denote this random object, whose randomness depends on $\mathcal{P}_{t,K\cup L}$.

Define

$$Y_{f,x} \quad = \quad \frac{1}{\kappa} \sum_{k\in[\kappa]} (A_f(k, x) - B(k, x))^2 \ .$$

$A_f(k, x)$ is a random variable in $[0, 1]$ depending on the labels that x is assigned by $\{\psi'_{t,\gamma} : \gamma \subseteq K \cup L\}$. Also, $A_f(k, x)$ are independent across items, and hence $Y_{f,x}$ are mutually independent. Additionally,

$$\mathrm{Avg}[(A_f - B)^2] \quad = \quad \frac{1}{|Z_t|} \sum_{x\in Z_t} Y_{f,x} \ .$$

Fix some $k \in [\kappa]$ and $P \in \mathcal{P}_{t,K}$. Observe that for any labels $\omega \in \{0, 1, \ldots, \kappa\}^{\{\gamma\subseteq K\cup L\}}$, the probability that an item $x \in P$ receives the labels ω is some fixed probability ρ_ω. Therefore,

$$\begin{aligned}
\mathbb{E}\Big[\mathrm{Avg}[(A_f - B)^2 \mid k \times P]\Big] &= \frac{1}{|P|} \sum_{x\in P} \sum_{\omega} \Big(\rho_\omega (f(k, \omega) - B(k, x))^2\Big) \\
&= \sum_{\omega} \Big(\rho_\omega \,\mathrm{Avg}[(f(k, \omega) - B)^2 \mid k \times P]\Big) \\
&\geq \min_{a\in\{0,\varepsilon,2\varepsilon,\ldots,1\}} \mathrm{Avg}[(a - B)^2 \mid k \times P] \ .
\end{aligned}$$

Averaging, we have

$$\begin{aligned}
\mathbb{E}\Big[\mathrm{Avg}[(A_f - B)^2]\Big] &\geq \frac{1}{\kappa|Z_t|} \sum_{k\in[\kappa]} \sum_{P\in\mathcal{P}_{t,K}} \Big(|P| \min_{a\in\{0,\varepsilon,2\varepsilon,\ldots,1\}} \mathrm{Avg}[(a - B)^2 \mid k \times P]\Big) \\
&= \min_{A\in\mathcal{A}_{t,K}} \mathrm{Avg}[(A - B)^2] \\
&= \Delta_t(K, B) \ .
\end{aligned}$$

In other words, the expected error of the predictor A_f is at least the error of the optimal predictor in $\mathcal{A}_{t,K}$. Thus, by a union bound over all 2^λ choices of f and Hoeffding,

$$\begin{aligned}
\Pr\Big[\Delta_t(K, B) - \Delta_t(K \cup L, B) \geq \varepsilon\Big] &= \Pr\Big[\Delta_t(K, B) - \min_f \mathrm{Avg}[(A_f - B)^2] \geq \varepsilon\Big] \\
&\leq \sum_f \Pr\Big[\Delta_t(K, B) - \mathrm{Avg}[(A_f - B)^2] \geq \varepsilon\Big] \\
&\leq \sum_f \Pr\Big[\mathbb{E}[\mathrm{Avg}[(A_f - B)^2]] - \mathrm{Avg}[(A_f - B)^2] \geq \varepsilon\Big] \\
&\leq \sum_f \exp(-2\varepsilon^2|Z_t|) \\
&\leq \exp(-\lambda^{80}) \ .
\end{aligned}$$

Thus, by independence, the probability that this violation occurs for all $L \in \mathcal{T}'_{B,K}$ is at most $\exp(-\lambda^{109})$. Deconditioning and union bounding over all at most $(\eta^n)^{\lambda^{29}}(1+m)^{\kappa\lambda^{100}}\eta^n \le \exp(\lambda^{101})$ choices of $\mathcal{T}'_{B,K}$, B, and K completes the proof. □

Lemma A.7. *For all $B : [\kappa] \times Z_t \to \{0, m^{-1}, 2m^{-1}, \dots, 1\}$ and partial injections $\sigma : [i-1] \to [\eta]$, define $T_{B,\sigma} \subseteq [\eta]$ to be a minimal set satisfying: for all injections $\tau : [n] \to [\eta]$ extending σ where $\tau([i,n]) \cap T_{B,\sigma} = \emptyset$, and for all $k, r \in [\kappa]$,*

$$|Z_{t,r,[n]}(\tau)| \;\ge\; |Z_t|/\lambda \qquad \Longrightarrow$$
$$\Big|\mathrm{Avg}[B(k,\cdot) \mid Z_{t,r,[n]}(\tau)] - \mathrm{Avg}[B(k,\cdot) \mid Z_{t,r,[i-1]}(\sigma)]\Big| \;\le\; \varepsilon \;.$$

Then $|T_{B,\sigma}| \le \lambda^{30}$ for all B, σ simultaneously w.p. at least $1 - \lambda^{-2}$.

Proof. Fix B and σ. Suppose that $|T_{B,\sigma}| \ge \lambda^{30}$. The same argument as in Lemma A.6 allows us a choice of injections $\mathfrak{T}'_{B,\sigma}$ of size λ^{29} such that the sets $\{\tau([i,n]) : \tau \in \mathfrak{T}'_{B,\sigma}\}$ are disjoint. Therefore, it suffices to bound the probability that there are violations on the set of λ^{29} witnesses $\{(\tau, k_\tau, r_\tau) : \tau \in \mathfrak{T}'_{B,\sigma}\}$. Fix such a choice of witnesses.

Define $Z_{t,r,[i-1]}(\sigma)$ to be the value of $Z_{t,r,[i-1]}$ if $\tau_t|_{[i-1]} = \sigma$, and define $Z_{t,r,[n]}(\tau)$ to be the value of $Z_{t,r,[n]}$ if $\tau_t|_{[i-1]} = \sigma$ and $\tau_t|_{[i,n]} = \tau|_{[i,n]}$. Condition on $\{\psi'_{t,\gamma} : \gamma \subseteq \sigma([i-1])\}$, which fixes $Z_{t,r,[i-1]}(\sigma)$ for all $r \in [\kappa]$. Then observe that the conditioning leaves the sets $\{Z_{t,r_\tau,[n]}(\tau) : \tau \in \mathfrak{T}'_{B,\sigma}\}$ conditionally mutually independent by disjointness of $\{\tau([i,n]) : \tau \in \mathfrak{T}'_{B,\sigma}\}$.

Fix any $\tau \in \mathfrak{T}'_{B,\sigma}$. Observe that on the conditioning, each item in $Z_{t,r_\tau,[i-1]}(\sigma)$ is independently included in $Z_{t,r_\tau,[n]}(\tau)$ w.p. $\vartheta_{t,r_\tau,[n]}/\vartheta_{t,r_\tau,[i-1]}$. Therefore, by Corollary 3.10,

$$\Pr\Big[|Z_{t,r_\tau,[n]}(\tau)| \ge |Z_t|/\lambda,\ \Big|\mathrm{Avg}[B(k_\tau,\cdot) \mid Z_{t,r_\tau,[n]}(\tau)] - \mathrm{Avg}[B(k_\tau,\cdot) \mid Z_{t,r_\tau,[i-1]}(\sigma)]\Big| \ge \varepsilon\Big]$$
$$\le \quad 2\exp(-2\varepsilon^2|Z_t|/\lambda) \quad \le \quad \exp(-\lambda^{79}) \;.$$

Thus, by independence, the probability that this violation occurs for all $\tau \in \mathfrak{T}'_{B,\sigma}$ is at most $\exp(-\lambda^{108})$. Deconditioning and union bounding over all at most $(\eta^n\kappa^2)^{\lambda^{29}}(1+m)^{\kappa\lambda^{100}}\eta^n \le \exp(\lambda^{101})$ choices of $\{(\tau, k_\tau, r_\tau) : \tau \in \mathfrak{T}'_{B,\sigma}\}$, B, and σ completes the proof. □

We are now ready to finish the proof. Fix a relevant block-visit pair (t,i) with prices $\boldsymbol{p}$. Let $\widehat{B} : [\kappa] \times Z_t \to \{0, m^{-1}, 2m^{-1}, \dots, 1\}$ be the discretization of $B_{t,\boldsymbol{p}}$. Define

$$T_{\boldsymbol{p}} \quad = \quad \left[T_{\widehat{B}} \cup \left(\bigcup_{K \subseteq \tau_t([i-1])} T_{\widehat{B},K}\right) \cup T_{\widehat{B},\tau_t|_{[i-1]}}\right] \setminus \tau_t([i-1])$$

and

$$T'_{\boldsymbol{p}} \quad = \quad \begin{cases} T_{\boldsymbol{p}} & |T_{\boldsymbol{p}}| \le \lambda^{31} \\ \emptyset & |T_{\boldsymbol{p}}| > \lambda^{31} \end{cases} \;.$$

By Lemmas A.5 to A.7, $\Pr[|T_{\boldsymbol{p}}| \ge \lambda^{31}] \le 3\lambda^{-2}$. Since $T'_{\boldsymbol{p}}$ is fully determined by the mechanism's information prior to the response from bidder i, Lemma A.4 gives us that

$$\Pr[T_{\boldsymbol{p}} \cap N_{t,i} \ne \emptyset] \quad \le \quad \Pr[T'_{\boldsymbol{p}} \cap N_{t,i} \ne \emptyset] + \Pr[|T_{\boldsymbol{p}}| \ge \lambda^{31}] \quad \le \quad 4\lambda^{-2} \;.$$

On the event $T_{\boldsymbol{p}} \cap N_{t,i} = \emptyset$ and the event of Proposition A.2, we have:

1. For every $\gamma = (\ell, L) \in \Gamma_{t,i}$, we have $\ell = \tau_t(i) \in N_{t,i}$, so Lemma A.5 gives that for all $A \in \mathcal{A}_{t,\gamma}$,

$$\left| U_t(\gamma, A, \widehat{B}) - \left(2 \operatorname{Avg}[\widehat{B}] + 2 \operatorname{Avg}[A\widehat{B}] \right) \right| \quad \leq \quad \varepsilon \ .$$

2. For every $L \subseteq N_{t,i} \setminus \tau_t(i)$ with $|L| \leq n-1$, we have $L \setminus \tau_t([i-1]) \subseteq N_{t,i}$, so Lemma A.6 gives

$$\Delta_t(L \cap \tau_t([i-1]), \widehat{B}) - \Delta_t(L, \widehat{B}) \quad \leq \quad \varepsilon \ .$$

3. $\tau_t([n]) \subseteq N_{t,i}$, so Lemma A.7 and Proposition A.2 give that for all $k, r \in [\kappa]$,

$$\left| \operatorname{Avg}[\widehat{B}(k,\cdot) \mid Z_{t,r,[n]}] - \operatorname{Avg}[\widehat{B}(k,\cdot) \mid Z_{t,r,[i-1]}] \right| \quad \leq \quad \varepsilon$$

Since the above functions of $\widehat{B}$ are $(2\lambda+4)$-Lipschitz, 4-Lipschitz, and 2-Lipschitz in the sup-norm, and $\|\widehat{B} - B_{t,\boldsymbol{p}}\|_\infty \leq m^{-1}$, replacing $\widehat{B}$ by $B_{t,\boldsymbol{p}}$ only increases the error terms from ε to 2ε. This proves (1) and (2) of Proposition A.3 for the block-visit pair (t,i).

To prove (3), we additionally need to show that w.h.p.,

$$\left| \operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[n]}] \right| \quad \leq \quad \varepsilon \ .$$

To do so, condition on all the randomness of the instance except the draws of $Z_{t,1}, \ldots, Z_{t,\kappa}$. This fixes $B_{t,\boldsymbol{p}}$ (because the owner of Z_t has not been visited yet) and $Z_{t,r,[n]}$. Since each item in $Z_{t,r,[n]}$ is independently included in $Z_{t,r}$ w.p. $1/(\vartheta_{t,r,[n]} \operatorname{val}(r))$, Corollary 3.10 and Proposition A.2 yield

$$\begin{aligned} \Pr\left[|Z_{t,r}| \geq |Z_t|/(2\lambda), \ \left| \operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r}] - \operatorname{Avg}[B_{t,\boldsymbol{p}}(k,\cdot) \mid Z_{t,r,[n]}] \right| \geq \varepsilon \right] \\ \leq \quad 2\exp(-2\varepsilon^2 |Z_t|/(2\lambda)) \quad \leq \quad 2\lambda^{-2} \ . \end{aligned}$$

Deconditioning and union bounding over the κ^2 choices of k, r and the event of Proposition A.2 proves (3). One last union bound over block-visit pairs shows that the desired event holds w.p. at least $1 - 3n^2\kappa^2\lambda^{-2} \geq 1 - \lambda^{-1}$. □