

Nearly Tight Regret Bounds for Profit Maximization in Bilateral Trade*

Simone Di Gregorio[†] Paul Dütting[‡] Federico Fusco[§] Chris Schwiegelshohn[¶]

Abstract

Bilateral trade models the task of intermediating between two strategic agents, a seller and a buyer, willing to trade a good for which they hold private valuations. We study this problem from the perspective of a broker, in a regret minimization framework. At each time step, a new seller and buyer arrive, and the broker has to propose a mechanism that is incentive-compatible and individually rational, with the goal of maximizing profit.

We propose a learning algorithm that guarantees a nearly tight $\tilde{O}(\sqrt{T})$ regret in the stochastic setting when seller and buyer valuations are drawn i.i.d. from a fixed and possibly correlated unknown distribution. We further show that it is impossible to achieve sublinear regret in the non-stationary scenario where valuations are generated upfront by an adversary. Our ambitious benchmark for these results is the best incentive-compatible and individually rational mechanism. This separates us from previous works on efficiency maximization in bilateral trade, where the benchmark is a single number: the best fixed price in hindsight.

A particular challenge we face is that uniform convergence for all mechanisms' profits is impossible. We overcome this difficulty via a careful chaining analysis that proves convergence for a provably near-optimal mechanism at (essentially) optimal rate. We further showcase the broader applicability of our techniques by providing nearly optimal results for the joint ads problem.

1 Introduction

Bilateral trade is a fundamental economic model where two agents—a seller and a buyer—are interested in trading a good [e.g., Myerson and Satterthwaite, 1983]. Both agents are characterized by a private valuation for the item, and their goal is to maximize their own utility. Previous work in computer science [e.g., Colini-Baldeschi et al., 2016, Brustle et al., 2017, Babaiouff et al., 2020, Blumrosen and Dobzinski, 2021, Dütting et al., 2021, Deng et al., 2022, Kang et al., 2022, Liu et al., 2023, Cai and Wu, 2023, Bernasconi et al., 2024] has mainly focused on maximizing economic efficiency, measured in terms of welfare or gains-from-trade. We consider this problem from the perspective of a broker whose goal is to maximize its profit (payment from the buyer minus

*S.D.G. and F.F. were supported by the PNRR MUR project IR0000013-SoBigData.it and the FAIR (Future Artificial Intelligence Research) project PE0000013, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence). S.D.G. was also supported by the Institute for Complex Systems (Italian National Research Council). C.S. was supported by a Google Research Award and by the Independent Research Fund Denmark (DFF) under a Sapere Aude Research Leader grant No 1051-00106B.

[†]Dept. of Computer, Control and Management Engineering, Sapienza University of Rome, Rome, Italy. Email: simone.digregorio@uniroma1.it

[‡]Google Research, Zürich, Switzerland. Email: duetting@google.com

[§]Dept. of Computer, Control and Management Engineering, Sapienza University of Rome, Rome, Italy. Email: fuscof@diag.uniroma1.it

[¶]Dept. of Computer Science, Aarhus University, Aarhus, Denmark. Email: schwiegelshohn@cs.au.dk

payment to the seller) while ensuring that (i) agents behave according to their true preferences (*dominant-strategy incentive compatibility*) and (ii) the utility for participating in the mechanism of each agent is non-negative (*individual rationality*).

Profit maximization is a canonical objective in economics (already studied in the seminal paper by Myerson and Satterthwaite [1983]), and it is a well-motivated objective in the applications (e.g., all e-commerce platforms that connect sellers to buyers). It is also gaining attention from a theoretical computer science perspective (most notably in Hajiaghayi et al. [2025]). The predominant approach considered in prior work is *Bayesian*, as it assumes that both agents’ valuations are drawn from known and typically independent distributions. In many practical applications, however, the availability of such data is a non-trivial assumption, and valuations may not be independent.

In this work, we study profit maximization in repeated bilateral trade through the lens of online learning, with the goal of designing an algorithm that minimizes the regret with respect to the best mechanism in hindsight. Note that unlike prior research on repeated bilateral trade, which has focused on gains-from-trade [e.g., Cesa-Bianchi et al., 2024a,b, Azar et al., 2024, Bernasconi et al., 2024] and has thus considered *fixed-price mechanisms*, here we have the ambitious goal of converging to the profit of a provably complex optimal mechanism. In fact, while fixed-price mechanisms to some extent *characterize* efficiency-maximizing mechanisms, they are far from optimal in our problem (we refer to Section B for further details). The nature of our benchmark also differentiates us from the vast majority of the online learning literature on economic problems, where the learning objective is intrinsically parametric, e.g., the best fixed price(s), [e.g., Blum et al., 2003, Kleinberg and Leighton, 2003, Gatmiry et al., 2024], the best reserve price(s) [e.g., Cesa-Bianchi et al., 2015, Roughgarden and Wang, 2019], the best linear contract [Zhu et al., 2023, Dütting et al., 2023] or the best fixed bid [e.g., Feng et al., 2018, Cesa-Bianchi et al., 2024c].

As it is traditional in online learning, we study two data generation models: the adversarial input model, where the agent’s valuations are decided upfront by an oblivious adversary, and the stochastic i.i.d. case, when the agent’s valuations are drawn i.i.d. from a possibly correlated but fixed unknown distribution. Note that, in the stochastic i.i.d. case, the benchmark coincides with the best Bayesian mechanism, but unlike in the typical “one-shot” Bayesian mechanism, the input distribution is not known in advance, can be correlated, and has to be learned on the fly.

1.1 Our Results

Stochastic Setting. As our main result, we show that the stochastic i.i.d. setting, with possibly correlated buyer and seller valuations, admits a (polynomial-time) learning algorithm with a nearly tight regret bound of $\tilde{O}(\sqrt{T})$ *. A particular challenge is that uniform convergence for all mechanisms’ profits is impossible. We overcome this via a careful chaining analysis that proves convergence for a provably near-optimal mechanism at (essentially) optimal rate. We provide a more detailed summary of our analysis in Section 3.

Key Result 1 (Theorem 4). For the stochastic i.i.d. setting, there is a (polynomial-time) learning algorithm that achieves a regret of $\tilde{O}(\sqrt{T})$ with respect to the best mechanism in hindsight.

We complement this result with a lower bound of $\Omega(\sqrt{T})$ on the regret of any learning algorithm (Theorem 5), showing that our algorithm is optimal up to poly-logarithmic terms. This is accomplished by embedding the standard hard instance for prediction with experts [e.g., Cesa-Bianchi and Lugosi, 2006]. Our analysis also (essentially) settles the offline sample complexity of learning an expected ε -optimal mechanism, in the spirit of the sample complexity literature [e.g., Cole

*We use $\tilde{O}(\cdot)$ to suppress poly-logarithmic terms.

and Roughgarden, 2014]. More precisely, the analysis of our $\tilde{O}(\sqrt{T})$ -regret algorithm and the corresponding lower bound implies that to recover a mechanism whose expected profit is ε -optimal, having access to approximately $1/\varepsilon^2$ is sufficient as well as necessary.

Adversarial Setting. Moving to the more challenging adversarial setting, we construct a family of hard instances that defuses any learning algorithm, thus implying that no-regret is unattainable in such regime. We actually prove a stronger result: in the adversarial setting, no learning algorithm can achieve sublinear regret with respect to a $2/3 + \varepsilon$ fraction of the optimal profit.

Key Result 2 (Theorem 6). For the adversarial setting, no learning algorithm can achieve sublinear regret with respect to a $2/3 + \varepsilon$ fraction of the optimal profit.

Our results confirm an intriguing separation between the adversarial and the stochastic model, that has been recently uncovered in many other economically motivated scenarios [e.g. Cesa-Bianchi et al., 2024a,c, Aggarwal et al., 2024]. This difference is indeed not observable in many classical online learning problems, such as prediction with experts, bandits, feedback graphs, or partial monitoring [Bartók et al., 2014, Alon et al., 2015]. A natural open question is whether there are meaningful *intermediate* data generation models beyond the stochastic setting where no-regret is achievable, as in, e.g., smoothed analysis, random order, adversarial corruption [see, e.g., Roughgarden, 2020]. A promising direction here is the notion of a σ -smooth adversary [Haghtalab et al., 2024], that has been successfully adopted to move beyond adversarial lower bounds for e.g., online classification [Block et al., 2022], bilateral trade [Cesa-Bianchi et al., 2024b] and general auctions [Durvasula et al., 2023]. We leave this question as an exciting open problem.

Joint Ads Problem. We finally exploit a novel and intriguing correspondence between bilateral trade and non-excludable mechanism design to (essentially) close an open problem of Aggarwal et al. [2024], by reducing an $O(T^{3/4})$ to an $\tilde{O}(\sqrt{T})$ regret bound for the joint ads problem. In this problem, two advertisers, e.g., a brand and a merchant, both derive a private possibly correlated value from an ad, and the mechanism’s decision is to either show the ad or not, and define payments from either party to the mechanism, with the goal of maximizing revenue.

1.2 Challenges

Consider a generic mechanism M that belongs to the set \mathcal{M} of all dominant-strategy incentive compatible (DSIC) and individually rational (IR) mechanisms. M is specified by two functions—an allocation and a payment rule—mapping bids to allocation and payments. Standard mechanism design arguments (see Section 2) enable a simple geometric interpretation of the structure of M . Its allocation rule corresponds to a subdivision of the unit square $[0, 1]^2$ into two regions: one in which the trade happens and one where this is not the case. This partition has to be “north-west” monotone so that increasing the buyer’s bid or decreasing the seller’s does not decrease the likelihood of observing a trade. The payments are then given by the “south-east” projection of the valuation pair on the boundary of the allocation region (see Figure 1a).

This characterization simplifies our task from learning two functions to a single curve/region. However, this is far from trivializing the problem: (i) the action space is (still) non-parametric, and (ii) the profit function mapping mechanisms to their induced profit may be highly non-regular. These two features defuse the standard learning approaches, as we briefly detail in the following.

Insufficiency of Fixed Discretization. When faced with a large action space, the natural approach in online learning is to discretize it (e.g., Kleinberg et al. [2019]). This method splits the problem in two: finding offline a finite subset of actions that well approximates the optimal action *uniformly over all possible inputs* and then running online a *discrete* learning algorithm on

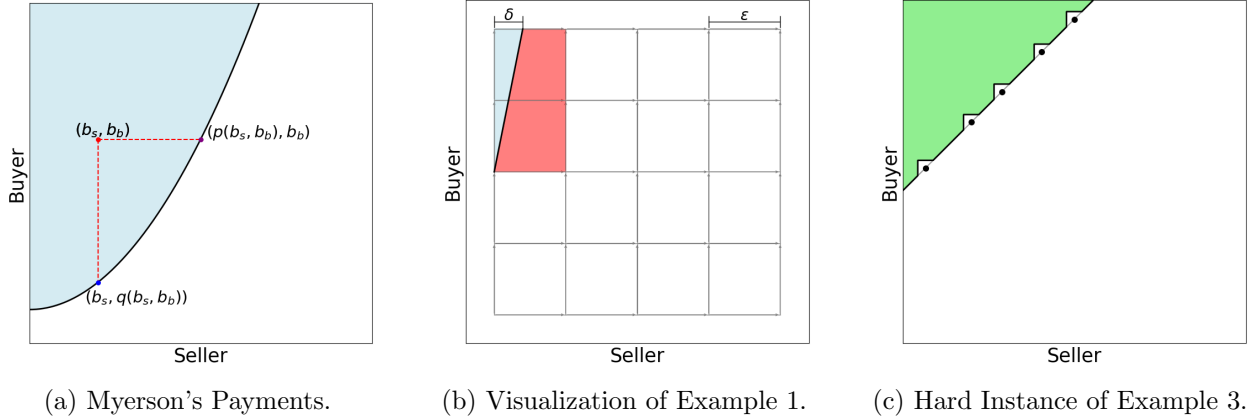


Figure 1: Visualization of Myerson's Payment, and Examples 1 and 3.

the discretized action space. The most natural discretization of our mechanism space is to cast a uniform ε -grid on the $[0, 1]^2$ square and consider all the $\binom{2/\varepsilon}{1/\varepsilon}$ mechanisms $\mathcal{M}_\varepsilon^\perp$ with allocation regions that can be described as the union of the cells of the grid.* Unfortunately, as we show in the following example, this rich class of mechanisms does not provide a uniform approximation for our problem. This is due to the non-continuity (and thus non-Lipschitzness) of our objective function.

Example 1 (A Fixed Discretization Does Not Work). *Let $\delta > 0$ be an arbitrarily small parameter, and consider the uniform distribution \mathcal{D} on the points $(0, 1/2)$ and $(\delta, 1)$. The optimal mechanism allocates in the triangle of vertices $(0, 1/2)$, $(0, 1)$, and $(\delta, 1)$, so that any valuations in the support of \mathcal{D} result in a trade, with expected revenue approximately $3/4$ (we consider δ negligible). Consider now any fixed discretization parameter $\varepsilon > \delta^\dagger$, the support of \mathcal{D} is contained in the first column of the grid, so the profit extracted by any mechanism $M \in \mathcal{M}_\varepsilon^\perp$ only depends (up to $O(\delta)$ payment to the seller) on the horizontal segment that crosses such column. The best possible profit corresponds to crossing the first column approximately at height $1/2$, for an expected utility of $1/2$, which is a constant factor away from the optimal one (see also Figure 1b).*

Fixed discretization of the action space is one of the main tools for positive results in the adversarial case when faced with a continuous action space [e.g., Kleinberg and Leighton, 2003, Kleinberg et al., 2019]. This fundamental gap between the best DSIC and IR mechanism and the best-discretized mechanism is, in fact, at the heart of our $\Omega(T)$ adversarial lower bound (Theorem 6), where we use a “needle in a haystack” construction to hide the optimal mechanism to the learner.

Failure of Local Search. A second, natural, approach consists in implementing some sort of local search in the mechanism space, as in e.g., stochastic gradient ascent. Regardless of the metric or notion of concavity we could impose on the non-parametric mechanism space \mathcal{M} , the crucial ingredient we would need for this approach to succeed is the uniqueness or “vicinity” of optimal mechanisms. Unfortunately, this is not the case, as detailed in the following example.

Example 2 (Non-Uniqueness of Optimal Mechanisms). *Consider the uniform distribution over valuations $(0, 1)$, and $(1/4, 3/4)$. Under this distribution, there are two optimal mechanisms: the*

*The family $\mathcal{M}_\varepsilon^\perp$ is a uniform ε -coverage of the mechanism space, for a suitable definition of distance between mechanisms. Similar ideas have been used successfully with some additional assumptions on the data generation model, e.g., smoothness of the underlying distributions [Durvasula et al., 2023, Aggarwal et al., 2024].

†The parameter δ can be arbitrarily small, while ε is fixed in advance by the learner. Note, for the same argument, one can adjust this construction to defuse *any* fixed grid.

mechanism M_1 that only allocates in $\{(0, 1)\}$, and M_2 that allocates in the $[0, 1/4] \times [3/4, 1]$ rectangle. These two mechanisms achieve expected utility $1/2$ and are “well-separated” (for instance, as soon as we restrict the allocation region of M_2 , we lose the trade opportunity given by $(1/4, 3/4)$, and the utility drops by a constant to $1/4$).

The possibility of having two well-separated maxima allows us, in Theorem 5, to embed the standard hard instance for prediction with two experts [e.g., Cesa-Bianchi and Lugosi, 2006] in our setting, thus providing a $\Omega(\sqrt{T})$ lower bound on the minimax regret rate for the stochastic setting.

Impossibility of Uniform Learning. The third standard strategy leverages the uniform learnability of the objective functions. Suppose for the sake of the argument that we were able to prove a statement of the following type: for any precision $\varepsilon > 0$, it is possible to reconstruct the expected profit of *all* mechanisms at the same time (up to an ε precision error) using $n(\varepsilon)$ i.i.d. samples from the underlying distribution \mathcal{D} . This would imply that a follow-the-leader strategy incurs an overall regret of order $\sum_t \varepsilon_t$, where ε_t is the precision achievable with t samples. The road-block to pursuing this strategy in our case is that the family of profit function is complex from the statistical point of view, as it exhibits unbounded pseudo-dimension and constant Rademacher complexity (even the range space of monotone allocation region has unbounded VC dimension, see Section C for further details). As it turns out, these negative clues are not a coincidence: uniform learnability *is not possible* in our problem, as we demonstrate in the following example (See also Theorem 8).

Example 3 (A Hard Distribution). *Consider the uniform distribution \mathcal{D} on the segment joining $(0, 1/2)$ and $(1/2, 1)$. For any number n , and any realization of n i.i.d. samples S from \mathcal{D} , consider mechanism M_S that allocates in the $(0, 1)$, $(0, 1/2)$ and $(1/2, 1)$ triangle, save for puncturing out the points S , as in Figure 1c. We have the following two properties: (i) the expected profit of M_S under \mathcal{D} is at least $1/4$, but (ii) the empirical profit of M_S on the samples is zero. This implies that uniform learning is not possible against \mathcal{D} .*

Our Techniques. We overcome these barriers by controlling a stochastic process that simulates a branch-and-bound strategy among sets of increasingly powerful mechanisms. The branch-and-bound strategy initially considers fixed-price mechanisms (rectangular allocation regions), then iteratively branches off multiple refined mechanisms by modifying the boundaries of simpler ones. As we refine these mechanisms, they become richer and exhibit potentially larger profits. Thus, we implicitly explore the entire search space for mechanisms, avoiding the barrier inherent in local improvements. The issue with considering richer mechanisms is that we no longer can show that their empirical profit is concentrated around the real expected profit, due to the impossibility of uniform learning. Our important insight is that the impossibility of uniform learning in some sense mainly affects the mechanisms with poor empirical profit, thus we can abstain from search paths that explore mechanisms with small profit and focus on refining profitable mechanisms.

We adapt our characterization of profitable mechanisms as we explore the search space. This addresses the lack of good fixed discretization strategies, as we are learning a (data and distribution dependent) discretization on the fly. For the profitable mechanisms, this branch-and-bound strategy mimics a chaining analysis, a powerful tool for bounding suprema of stochastic processes (see Talagrand [2014]). Chaining has been used in a wide variety of applications including generalization bounds [Audibert and Bousquet, 2007, Cohen-Addad et al., 2025b, Li et al., 2000], dimensionality reduction [Bourgain et al., 2015, Narayanan and Nelson, 2019], error correcting codes [Rudra and Wootters, 2014], graph sparsification [Jambulapati et al., 2023, Lee, 2023], data structures [Braverman et al., 2017], and compression [Cohen-Addad et al., 2022, 2025a]; see Nelson [2016] for more applications and references. The chaining analysis gives us very tight control over the profit of

good mechanisms, which allows us to learn a sequence of mechanisms with nearly optimal regret. We give a more technical overview of the key ingredients in our proof in Section 3.

1.3 Further Related Work

Efficiency and Profit in Bilateral Trade. The celebrated result of Myerson and Satterthwaite [1983] on bilateral trade (strengthening an earlier result of Vickrey [1961]) states that even for independent and regular valuation distributions, it is impossible to simultaneously achieve (a) incentive compatibility—even the weaker notion of Bayesian incentive compatibility (BIC)—(b) perfect efficiency, (i.e., trade whenever the buyer has a higher valuation than the seller), (c) individual rationality, and (d) budget balance (requiring that the broker does not run a deficit). On the positive side, they also characterize the optimal BIC mechanisms (which turn out to be DSIC) for social welfare, gains from trade, and profit under some assumptions on the valuations’ distributions.

Motivated by this impossibility result, a fruitful line of work in computer science has aimed at mechanisms that are incentive compatible and budget balanced, while achieving a provable worst-case approximation guarantee — with respect to social welfare [Liu et al., 2023, Cai and Wu, 2023, Blumrosen and Dobzinski, 2021, Dütting et al., 2021, Kang et al., 2022] or gains-from-trade [McAfee, 2008, Fei, 2022, Cai et al., 2021, Deng et al., 2022, Brustle et al., 2017]. While these results hold in the Bayesian setting, where the agents’ valuations are drawn independently from some known distribution, some of these approximation mechanisms actually make use of only a few (sometimes just a single sample) from the underlying distributions to achieve constant-factor approximations [Blumrosen and Dobzinski, 2021, Dütting et al., 2021, Kang et al., 2022].

The recent work of Hajiaghayi et al. [2025] studies the compatibility of profit and efficiency (both gains-from-trade and social welfare) in the Bayesian setting. While they demonstrate positive results for fixed-price mechanisms, they present strong impossibility results for general mechanisms.

Sample Complexity. Beyond the rich line of work on online learning, a related approach considers the *offline* learning problem of economic design problems in the spirit of PAC-learning [e.g., Cole and Roughgarden, 2014, Morgenstern and Roughgarden, 2015, Devanur et al., 2016, Cai and Daskalakis, 2017, Gonczarowski and Weinberg, 2021]. While the standard online-to-offline reduction tells that sample-complexity results imply online regret bounds, we are not aware of any prior work yielding non-trivial results for our settings. For instance, Cole and Roughgarden [2014] consider the setting where bidders’ valuations are independent and exhibit some regularity; we consider a bilateral trade with arbitrary and possibly correlated distributions. Morgenstern and Roughgarden [2015] provide results for finite pseudo-dimension; in Section C we prove that such measure is unbounded in our problem. Guo et al. [2021] show general sample complexity bounds for product distributions. Their results imply sample complexity bounds for a range of applications, via appropriate discretizations of the value space. In contrast, a feature of our problem is the lack of a viable discretization (see Example 1). They also provide results for *monotone* problems, something our setting does not enjoy.

2 The Learning Model and Preliminaries

Profit in Bilateral Trade. In the (static) bilateral trade problem, a seller and a buyer interact with each other. The seller holds a single good. The seller’s and buyer’s (private) valuations for the good are $v_s \in [0, 1]$ and $v_b \in [0, 1]$, respectively. Seller and buyer submit a bid $b_s \in [0, 1]$ and $b_b \in [0, 1]$ (not necessarily truthful) to a broker running mechanism M . A mechanism is characterized by an allocation region $A \subseteq [0, 1]^2$, and pricing rules $p, q : [0, 1]^2 \rightarrow [0, 1]$. The trade happens if and

only if the bids (b_s, b_b) belong to the allocation region A , while the payments are made according to p and q . Payments are from the broker to the seller, and from the buyer to the broker. That is, the seller receives $p(b_s, b_b)$, while the buyer pays $q(b_s, b_b)$. We require that $p(b_s, b_b) = q(b_s, b_b) = 0$ whenever $(b_s, b_b) \notin A$, i.e., there is no trade.*

The buyer's and seller's utility with valuations v_b and v_s under bids $(b_b, b_s) \in [0, 1]^2$ are:

$$U_s(b_s, b_b) = v_s - \mathbb{1}_{\{(b_s, b_b) \in A\}} \cdot v_s + p(b_s, b_b), U_b(b_s, b_b) = \mathbb{1}_{\{(b_s, b_b) \in A\}} \cdot v_b - q(b_s, b_b)$$

We focus on dominant-strategy incentive compatible (DSIC) and individually rational (IR) mechanisms. There, each player maximizes his utility by a truthful bid regardless of the other player's bid, and the utility from participating in the mechanism is at least as high as that from not participating in the mechanism (i.e., it is at least v_s for the seller and at least 0 for the buyer). In formulas,

$$\begin{aligned} \text{DSIC: } & U_s(v_s, b_b) \geq U_s(b_s, b_b) && \forall v_s \in [0, 1], (b_s, b_b) \in [0, 1]^2 \\ & U_b(b_s, v_b) \geq U_b(b_s, b_b) && \forall v_b \in [0, 1], (b_s, b_b) \in [0, 1]^2 \\ \text{IR: } & U_s(v_s, b_b) \geq v_s, U_b(b_s, v_b) \geq 0 && \forall (v_s, v_b) \in [0, 1]^2, (b_s, b_b) \in [0, 1]^2 \end{aligned}$$

We denote the class of all DSIC and IC mechanisms by \mathcal{M} . Since we consider DSIC and IR mechanisms, we will henceforth assume that the players' bids are equal to their respective valuations. The profit earned by a broker running mechanism M is the difference between the payment the mechanism collects from the buyer and the payment that it makes to the seller. That is, $\text{profit}(M, (v_s, v_b)) = q(v_s, v_b) - p(v_s, v_b)$.

The Learning Protocol. We study the repeated bilateral trade problem, defined as follows. At each time step t , a new seller-buyer pair arrives, with private valuations $v^t = (v_s^t, v_b^t) \in [0, 1]^2$. The learner/broker proposes a mechanism M^t , and the agents bid. The trade occurs according to M^t and the bids of the agents, which then leave forever. The broker earns the difference between the price paid *by* the buyer and that paid *to* the seller, i.e., $\text{profit}(M^t, v^t)$. Formally, we consider the following learning protocol.

The Learning Protocol

for time $t = 1, 2, \dots$ **do**

a new pair of agents arrives with private valuations $v^t = (v_s^t, v_b^t) \in [0, 1]^2$
the learner declares a mechanism M^t , and then the agents declare their bids
the trade takes place according to mechanism M^t and the bids
the learner gains $\text{profit}(M^t, v^t) \in [-1, 1]$

We aim to design a learning algorithm that minimizes the regret:

$$\sup_{M \in \mathcal{M}} \mathbb{E} \left[\sum_{t=1}^T \text{profit}(M, v^t) - \sum_{t=1}^T \text{profit}(M^t, v^t) \right].$$

The typical goal is to achieve sublinear worst-case regret; this implies that on any input, the algorithm's profit converges, on average, to that of the benchmark.

We consider both the adversarial setting, where the sequence of valuations is generated up-front by an oblivious adversary, and the stochastic i.i.d. one, where the agents' valuations $v^t = (v_s^t, v_b^t)$

*The requirement that payments are bounded in $[0, 1]$ and zero when there is no trade is without loss of generality, given our focus on DSIC and IR mechanisms and the goal of maximizing the broker's profit. See Remark 4.

are drawn i.i.d. from an unknown distribution. In the latter model, the v^t are independent across time, but the seller and buyer valuations can be arbitrarily correlated.

Characterization of Mechanisms. We next observe that the allocation regions and the pricing functions of DSIC and IR mechanisms have a distinctive structure: the allocation region needs to respect a certain monotonicity property, while the allocation region uniquely induces the payments. We formally introduce monotonicity and Myerson payments [Myerson, 1981].

Definition 1 (Monotone Allocation Region). *We define the partial ordering \preceq on $[0, 1]^2$ as follows: for any two points $x = (x_1, x_2)$ and $y = (y_1, y_2)$, we say that $x \preceq y$ (y dominates x) if $x_1 \geq y_1$ and $x_2 \leq y_2$. An allocation region $A \subseteq [0, 1]^2$ is monotone (with respect to \preceq) if for any $x \in A$ and $y \in [0, 1]^2$ with $x \preceq y$ it holds that $y \in A$.*

Definition 2 (Myerson Payments). *Let A be a monotone allocation region, then the associated Myerson payments are defined as follows:*

$$\begin{aligned} p(b_s, b_b) &= \mathbb{1}_{\{(b_s, b_b) \in A\}} \cdot \max\{x \in [0, 1] \mid (x, b_b) \in A\} & \forall (b_s, b_b) \in [0, 1]^2 \\ q(b_s, b_b) &= \mathbb{1}_{\{(b_s, b_b) \in A\}} \cdot \min\{y \in [0, 1] \mid (b_s, y) \in A\} & \forall (b_s, b_b) \in [0, 1]^2 \end{aligned}$$

In words, for an allocation region to be monotone it should be “closed” in a north-west direction. That is, if a point (b_s, b_b) is in A then any point (b'_s, b_b) with $b'_s \leq b_s$ and any point (b_s, b'_b) with $b'_b \geq b_b$ should be in A as well. The payments of a point $(b_s, b_b) \in A$, in turn, correspond to the south-east projection to the allocation boundary. See Figure 1a for illustrations.

Proposition 1. *A mechanism M for bilateral trade is dominant-strategy incentive compatible and individually rational if and only if its allocation region is monotone and the payments are Myerson payments.*

The proof of Proposition 1 follows by standard arguments [Myerson, 1981, Myerson and Satterthwaite, 1983]. For completeness, we provide a proof in Section A. An important consequence of Proposition 1 is that it reduces the learning problem to learning the monotone allocation region A , with the learner’s profit given by Myerson payments. Note, though, that this is far from trivializing the problem, as we formally argue in Section C.

Profit and Properties of Optimal Mechanism. We conclude the preliminaries with some observations regarding the tradeoffs involved in choosing different (monotone) allocation regions. The first observation is that by shrinking the allocation region we loose potentially profitable trades, while we increase the profit of those trades that still happen.

Observation 1 (Trades vs. Profit Trade-off). *Consider any two mechanisms M, M' in \mathcal{M} with allocation regions A, A' such that $A \subseteq A'$. For any $v \in [0, 1]^2$ we have one of three cases:*

- $v \notin A'$, and $\text{profit}(M, v) = \text{profit}(M', v) = 0$
- $v \in A' \setminus A$, and $\text{profit}(M, v) = 0$
- $v \in A$, and $\text{profit}(M, v) \geq \text{profit}(M', v)$

This trade-off allows us to derive a useful property of the optimal mechanism: its allocation region does not contain the lower-right triangle! Denote with U the upper-left triangle, $U = \{(x, y) \in [0, 1]^2 \mid x \leq y\}$, and with \mathcal{M}^+ the subfamily of \mathcal{M} containing the mechanisms whose allocation region is contained in U . We have the following lemma.

Lemma 1. *For any mechanism $M \in \mathcal{M}$ there exists a mechanism $M^+ \in \mathcal{M}^+$ such that, for any valuation v , the following inequality holds: $\text{profit}(M^+, v) \geq \text{profit}(M, v)$.*

Proof. Fix any mechanism $M \in \mathcal{M}$ with allocation region A and payment functions p, q , and consider the mechanism M^+ defined by the allocation region $A^+ = A \cap U$, where A is the allocation region of M . The set A^+ is monotone and thus is a valid allocation region. Fix any valuation pair $v = (v_s, v_b)$; we want to prove that $\text{profit}(M^+, v) \geq \text{profit}(M, v)$. From Observation 1, we know we only need to consider the case in which $v \in A \setminus A^+$. If $v \in A \setminus A^+$, then it does not belong to U and we have the following inequality: $\text{profit}(M, v) = (q(v) - p(v)) \leq (v_b - v_s) \leq 0 = \text{profit}(M^+, v)$, where the first inequality holds by individual rationality, and the second because $v \notin U$. \square

Remark 1 (The Allocation Regions Are Closed). *For convenience, we always assume that the allocation regions are closed. This is without loss of generality; in fact, for any allocation region in U , if we consider its closure, we get a mechanism that extracts at least as much revenue. This assumption allows us to simplify and formalize some arguments; for instance, if A is closed, then the max and min operations in the definition of the payments are well defined.*

3 Proof Ideas at a Glance

In this section, we briefly summarize the ideas behind our $\tilde{O}(\sqrt{T})$ result in the stochastic setting, that are formally laid down in Section 4. In this setting, valuations v^1, v^2, \dots are drawn i.i.d. according from a fixed but unknown distribution, and our goal is to approximate the best mechanism M_{opt} on the valuation distribution.* Consider what happens at the generic time step $t + 1$, after receiving samples v^1, \dots, v^t . If we could determine a mechanism M^{t+1} such that

$$\mathbb{E} [\text{profit}(M_{\text{opt}}, v)] \leq \mathbb{E} [\text{profit}(M^{t+1}, v)] + t^{-\alpha} \quad (1)$$

for some exponent $\alpha \in (0, 1)$, then summing up over all time step, we would get an overall regret $\sum_{t=1}^T t^{-\alpha} \in O(T^{1-\alpha})$. In particular, for $\alpha = 1/2$, we would obtain the optimal rate of \sqrt{T} . From Example 1 (and Theorem 8), we know that past samples are not enough to estimate *at the same time* all mechanisms' expected profit. It is, in fact, possible that the empirical profit of some mechanism can be arbitrarily far from its true expected profit. This rules out a Follow-The-Leader approach, which would update M^t to be the empirical optimum after every sample.

Since we are ultimately only interested in Equation 1, a natural idea is to constrain the candidate mechanisms to some subset $\mathcal{M}^* \subset \mathcal{M}$ such that: (i) we can show that uniform convergence holds for all $M \in \mathcal{M}^*$ and (ii) the best mechanism in \mathcal{M}^* has nearly the same profit as M_{opt} . Indeed, given an arbitrary mechanism M and precision parameter ε , we can construct an “ ε -simple” mechanism M' that well approximates the profit of M . The simplicity of M' is captured by the fact that the boundary of its allocation region is determined by at most $O(1/\varepsilon)$ -axis parallel segments, while the approximation property is formalized by the following inequality:

$$\mathbb{E} [\text{profit}(M, v)] \leq \mathbb{E} [\text{profit}(M', v)] + \varepsilon. \quad (2)$$

Limiting the mechanisms to the set \mathcal{M}_ε of such ε -simple mechanisms turns out to be a fruitful idea as proving a uniform convergence for these mechanisms is indeed possible. Our algorithm SIMPLIFY-THE-BEST-MECHANISM computes the mechanism \hat{M}_{opt} that maximizes the empirical profit and then posts the mechanism in \mathcal{M}_ε that approximates it. While the family of ε -simple mechanisms is inspired by a similar construction in Aggarwal et al. [2024], their work only implies a

*For simplicity, we assume here that the best mechanism exists.

convergence rate of $\alpha = 1/4$ (for a regret rate of order $T^{3/4}$) via a standard VC dimension argument. As a warm-up, we next present a uniform convergence result on \mathcal{M}_ε ($\alpha = 1/3$), then we argue how to further refine our discretization argument to attain the optimal $\alpha = 1/2$.

3.1 Warm Up: A $\tilde{O}(T^{2/3})$ Regret Bound.

There are infinite mechanisms that can be described by $O(1/\varepsilon)$ axis parallel segments, so we do not have a finite bound on the size of \mathcal{M}_ε . One might be tempted to impose a fixed set of such lines as candidate boundary segments. Unfortunately, this approach does not work (Example 1) without prior knowledge of the distribution.

However, we are not interested in learning *all* ε -simple mechanisms, but only those relevant to the realized samples. While the distribution itself does not have finite support, the empirical one induced by the samples necessarily does. We then restrict our attention to the ε -grid induced by coordinates of the samples, leading to $\binom{t}{1/\varepsilon} \approx \exp(\log t/\varepsilon)$ many candidate mechanisms \mathcal{M}_ε . Using the randomness of the sample to prove concentration for mechanisms induced by the sample introduces dependencies; e.g., a mechanism's profit is highly correlated with the $O(1/\varepsilon)$ samples that determine its boundary. We address these correlations by removing these $O(1/\varepsilon)$ samples when proving a union bound, which still leaves us with $t - O(1/\varepsilon)$ fresh samples to prove concentration. With this setup, we can show via standard concentration inequalities that

$$\sup_{M \in \mathcal{M}_\varepsilon} \left| \frac{1}{t} \sum_{i=1}^t \text{profit}(M, v^i) - \mathbb{E}[\text{profit}(M, v)] \right| \lesssim \sqrt{\frac{\log |\mathcal{M}_\varepsilon|}{t}} \lesssim \sqrt{\frac{\log T}{\varepsilon t}} \quad (3)$$

The optimal mechanism M_{opt} is one single mechanism, so we can safely assume its empirical performance concentrates on past data. We have the following chain of inequalities:

$$\begin{aligned} \mathbb{E}[\text{profit}(M_{\text{opt}}, v)] &\lesssim \frac{1}{t} \sum_{i=1}^t \text{profit}(\hat{M}_{\text{opt}}, v^i) + \sqrt{\frac{\log T}{t}} \quad (M_{\text{opt}} \text{ concentrates and by design of } \hat{M}_{\text{opt}}) \\ &\lesssim \frac{1}{t} \sum_{i=1}^t \text{profit}(M^{t+1}, v^i) + \varepsilon + \sqrt{\frac{\log T}{t}} \quad (\text{Equation (2) on empirical distribution}) \\ &\lesssim \mathbb{E}[\text{profit}(M^{t+1}, v)] + \varepsilon + \sqrt{\frac{\log T}{t}} + \sqrt{\frac{\log T}{\varepsilon t}}, \end{aligned} \quad (4)$$

where the last inequality is due to Equation (3). The $\tilde{O}(T^{2/3})$ bound follows by setting $\varepsilon \approx t^{-1/3}$.

3.2 Main Result: $\tilde{O}(\sqrt{T})$ Regret

The $\tilde{O}(T^{2/3})$ provides the important insight that among the $\exp((\log t)/\varepsilon)$ many candidate mechanisms in \mathcal{M}_ε induced by the samples, there is an M^t whose expected profit is within an additive error of $t^{-1/3}$ of the expected profit of M_{opt} . To ensure that the empirical profit of M^t is close to the expected profit of M^t , the previous approach consisted in simply performing a union bound over all such candidate $\exp((\log t)/\varepsilon)$ mechanisms. This approach is very natural, as it is difficult to characterize profitable from unprofitable mechanisms a priori. At the same time, the approach is lossy: We actually do not need strong concentration for unprofitable mechanisms in \mathcal{M}_ε .

We now describe how our chaining approach, originally designed to perform very efficient union bounds for uniform convergence, can provide some intuition on how to identify profitable mechanisms on the samples. While we cannot guarantee that we estimate the profit of all mechanisms in \mathcal{M}_ε




Figure 2: Visualization of approximating sequence (in red) of mechanism M (blue allocation region).

accurately, we show that this is possible for profitable mechanisms, while unprofitable mechanisms are either identified as such, or have their profit underestimated. Suppose we have a set of mechanisms \mathcal{N}_h , such that for every $M \in \mathcal{M}$, there exists $M' \in \mathcal{N}_h$ with

$$\left| \text{profit}(M, v^i) - \text{profit}(M', v^i) \right| \leq 2^{-h} \quad \forall v^1, v^2, \dots, v^t \quad (5)$$

A crucial ingredient in chaining is to decompose the target profit $\text{profit}(M, \cdot)$ via a telescopic argument that uses a sequence of approximating mechanisms as follows

$$\text{profit}(M, v^i) = \text{profit}(\phi_0(M), v^i) + \sum_{h=0}^{\log 1/\varepsilon} \text{profit}(\phi_{h+1}(M), v^i) - \text{profit}(\phi_h(M), v^i),$$

where each $\phi_h(M)$ belongs to the corresponding \mathcal{N}_h . This sequence of approximating mechanisms is visualized in Figure 2. Our goal is to learn the expected profit of M , which translates to learning separately the terms of the above formula. The zero-th term can be addressed with standard union bound techniques by noting that $\phi_0(M)$ is a fixed-price mechanism (see Figure 2a), so we only focus on learning the increments. The idea inherent in chaining is that instead of performing a union bound on all target mechanisms in one shot, we perform a union bound on the set of successive increments. Specifically, by combining basic union bounds similar to those used in Equation (3) and a reduced variance due to Equation (5), we have

$$\sup_{\substack{(M_{h+1}, M_h) \\ \in \mathcal{N}_{h+1} \times \mathcal{N}_h}} \left| \sum_{i=1}^t \frac{1}{t} (\text{profit}(M_{h+1}, v^i) - \text{profit}(M_h, v^i)) - \mathbb{E} [\text{profit}(M_{h+1}, v) - \text{profit}(M_h, v)] \right| \lesssim \sqrt{\frac{\log |\mathcal{N}_{h+1}| |\mathcal{N}_h|}{t \cdot 2^{-2h}}}.$$

At this point, we wish to find a careful tradeoff between the magnitude of increments and the size of the families \mathcal{N}_h . Optimistically, we would attempt to use \mathcal{M}_ε as \mathcal{N}_h , for appropriate values of ε , which would bound the right-hand-side of the above inequality with

$$\frac{1}{2^h} \sqrt{\frac{\log |\mathcal{N}_{h+1}| |\mathcal{N}_h|}{t}} \lesssim \frac{1}{2^h} \sqrt{\frac{2^h \log t}{t}} \lesssim \sqrt{\frac{\log t}{t \cdot 2^h}}. \quad (6)$$

Summing these terms for all increments would then yield

$$\sup_{M \in \mathcal{M}_\varepsilon} \left| \frac{1}{t} \sum_{i=1}^t \text{profit}(M, v^i) - \mathbb{E} [\text{profit}(M, v)] \right| \lesssim \sum_{h \geq 0} \sqrt{\frac{\log t}{t \cdot 2^h}} \lesssim \sqrt{\frac{\log T}{t}},$$

achieving the desired regret of $\tilde{O}(\sqrt{T})$ (via an analysis similar to the one in Equation (4)).

Unfortunately, there is a major issue with this (ideal) analysis: mechanisms in \mathcal{M}_ε do not satisfy Equation (5), as they only provide a *one-sided* bound! This requires several major adjustments to the chaining analysis; overcoming them constitutes the core technical contribution of our work. Our first important ingredient is to show that for any mechanism M , there exists $M' \in \mathcal{M}_\varepsilon$ such that

$$\mathbf{profit}(M, v^i) - \mathbf{profit}(M', v^i) \leq \varepsilon \quad \forall i \in \{1, \dots, t\}. \quad (7)$$

Notice the absence of absolute values when comparing this guarantee with Equation (5). By itself, Equation (7) does not facilitate the chaining analysis. However, for any telescoping sequence in which the increments are strictly increasing, that is $\mathbf{profit}(\phi_{h+1}(M), v^i) \geq \mathbf{profit}(\phi_h(M), v^i)$ for all v^i , it turns out that Equation (7) is sufficient to prove concentration.

Let us interpret when this happens and more importantly when this does not happen. If $\mathbf{profit}(\phi_{h+1}(M), v^i) \geq \mathbf{profit}(\phi_h(M), v^i)$ for all v^i , then $\phi_{h+1}(M)$ extracts slightly more profit than $\phi_h(M)$ from the distribution. In this case, $\phi_{h+1}(M)$ is clearly preferable to $\phi_h(M)$ on the sample. Fortunately, Equation (7) implies $|\mathbf{profit}(\phi_{h+1}(M), v^i) - \mathbf{profit}(\phi_h(M), v^i)| \leq 2^{-h}$, which implies with the optimistic chaining analysis implies that the empirical profit of $\phi_{h+1}(M)$ is concentrated around its expectation.

Next, let us consider the more challenging case where $\mathbf{profit}(\phi_{h+1}(M), v^i) \leq \mathbf{profit}(\phi_h(M), v^i)$. In this case, a possible configuration is $\mathbf{profit}(\phi_{h+1}(M), v^i) \leq 0$ and $\mathbf{profit}(\phi_h(M), v^i) \geq 1/2$ for a few samples*. Based on a single sample v^i , it is typically not possible to determine whether $\phi_{h+1}(M)$ is more profitable than $\phi_h(M)$ as it can exclude the potential profit of v^i for more profit on other candidate samples. Unfortunately, in this case, we also have

$$|\mathbf{profit}(\phi_{h+1}(M), v^i) - \mathbf{profit}(\phi_h(M), v^i)| \geq \frac{1}{2}. \quad (8)$$

This is fatal for a chaining analysis: If this happens too many times, then the increments (i.e. the right-hand-side of Equation (6)) can only be bounded by $\sqrt{2^h/t \cdot \log t}$, and summing up over all increments yields no concentration whatsoever.

Our main insight is that if Equation (8) happens for too many samples v^i , then $\phi_{h+1}(M)$ *has less profit than M_h and cannot be the empirical optimum!* Moreover, the same also holds for any mechanism that refines $\phi_{h+1}(M)$ via Equation (2), allowing us to remove all of these mechanisms for the search space of potentially optimal mechanisms (i.e., we are breaking the chain). Thus, we can show that a (slightly weaker) form of Equation (6) holds for mechanisms with sufficiently large profit, which still allows us to obtain the desired $\tilde{O}(\sqrt{T})$ regret bound.

In a nutshell, we have the following dichotomy working in our favor (formalized in Proposition 5):

- Either all but a very tiny fraction of the increments are positive. In this case, we can continue the chaining analysis.
- Or sufficiently many increments have a large negative value. In this case, the mechanism is suboptimal, and we no longer have to prove concentration.

While the remaining analysis still has substantial technical challenges, the on-the-fly construction and truncation of chains ultimately leading towards optimal or nearly optimal mechanisms is our main novelty. We therefore believe that our work can be a starting point for proving low regret bounds for other problems lacking strong characterizations of optima. Indeed, as a consequence of our work, we also manage to extend our analysis to the Joint-Ads problem.

*These numbers are possible, but chosen arbitrarily for the sake of the argument. The main point is to emphasize that the absolute value of differences can be a fairly large constant.

4 A $\tilde{O}(\sqrt{T})$ Learning Algorithm for the Stochastic Setting

In this section, we present the learning algorithm SIMPLIFY-THE-BEST-MECHANISM and prove that it enjoys a $\tilde{O}(\sqrt{T})$ regret rate in the stochastic setting.

4.1 SIMPLIFY-THE-BEST-MECHANISM

For each time step t , the learning algorithm SIMPLIFY-THE-BEST-MECHANISM behaves in two phases. First, it computes the best mechanism M_t^* on the past data; this operation can be reduced efficiently to computing the shortest path on a suitably defined directed graph. Second, it finds a simplified version of the optimal mechanism M_t^* by adapting its allocation region to a uniform grid of step-size $\varepsilon_t \approx t^{-1/2}$. We refer to the pseudocode for further details, while we present separately in Section 4.1.1 and Section 4.1.2 the implementation details on the two phases of the algorithm.

SIMPLIFY-THE-BEST-MECHANISM

Environment: Valuations $\{(v_1^t, v_2^t)\}$ are drawn i.i.d.

Propose any mechanism M^1 and observe (v_1^1, v_2^1)

for time $t = 2, \dots, T$ **do**

Set precision parameter $\varepsilon_t = 200\sqrt{\log^3 T/t-1}$

Let M_t^* be the profit-maximizing mechanism over valuations $(v_1^1, v_2^1), \dots, (v_1^{t-1}, v_2^{t-1})$

Let M^t be the simplified mechanism with precision ε_t associated with M_t^*

Propose mechanism M^t and observe (v_1^t, v_2^t)

4.1.1 Finding the Best Mechanism

The task of finding the best mechanism is inherently offline: given a set S of n points in the $[0, 1]^2$ square, how can we compute the empirical best mechanism with respect to S ? Surprisingly, this problem can be reduced to finding the shortest path on a suitable directed graph. We have the following definition:

Definition 3 (Point-Based Grid). *Given a set $S = \{(v_s^i, v_b^i)\}$ of n points, we introduce the projections over the x and y axis. S_x (respectively S_y) contains in increasing order the values $0, 1$, and v_s^i (respectively v_b^i) for $i = 1, \dots, n$. The grid G_S associated to S is a directed graph defined as follows:*

- The nodes of G_S are represented by the points in the cartesian product $S_x \times S_y$.
- Horizontal and vertical segments represent the edges of G_S :
 - $(x_i, y) \rightarrow (x_{i+1}, y)$ for any $y \in S_y$ and $x_i \in S_x \setminus \{1\}$
 - $(x, y_j) \rightarrow (x, y_{j+1})$ for any $x \in S_x$ and $y_j \in S_y \setminus \{1\}$

Beyond the formal definition, the construction of the directed graph G_S is pretty simple: it is just an axis-aligned orthogonal grid that contains all the points in S , the origin and $(1, 1)$, and joins them via horizontal (pointing right) and vertical (pointing up) edges. We refer to Figure 3.

Remark 2. *The grid G_S — as well as the other grids used in the rest of the paper — have a double nature: they are directed graphs, but they live embedded in the $[0, 1]^2$ square. Therefore, it is natural to treat a node in G_S as a point in the square and an edge as a segment. So, we refer to the geometric properties of nodes and edges via this natural embedding.*

We introduce grids to discretize the space of allocation regions. In particular, it is natural to associate mechanisms to paths on grids, as formalized by the following definition.

Definition 4 (Complete Path and Associated Mechanism). *A path π in a point-based grid G is called complete if it starts in $(0, x_1)$ and terminates in $(x_2, 1)$, for some $x_1, x_2 \in [0, 1]$. Any complete path π univocally identifies a mechanism M_π whose allocation region is defined as follows:*

$$A_\pi = \{v \in [0, 1]^2 \text{ such that } \exists v' \in \pi \text{ with } v \succeq v'\}^*.$$

Complete paths on a point based grid G_S have a simple structure: they start from a point on the left side of the $[0, 1]^2$ square, then move either right or up in G_S , and terminate in some node on the upper side of the square. From a combinatorial perspective, it is fairly easy to *count* how many complete paths are on a point based grid.

Lemma 2 (Number of Complete Paths in G_S). *Let G_S be a point induced grid with $|S| = n \geq 5$, then there are at most $(2e)^{2n}$ distinct complete paths in G_S .*

Proof. Consider the nodes of G_S . They form an orthogonal grid in the $[0, 1]^2$ square composed by $(n+2)^2 \leq 2n^2$ nodes, where in the previous inequality we used that $n \geq 5$. There are at most $2n^2$ ways of choosing the starting and the terminal node (at most $n+2$ nodes on the left side and $n+2$ nodes of the upper side of the square) and, once these two points are fixed, a complete path is characterized by at most $\binom{2(n+1)}{n+1}$ ways in which it can arise, since its length is at most $2(n+1)$ (from $(0, 0)$ to $(1, 1)$) and a $2(n+1)$ -long complete path must go “up” $n+1$ times and “right” $n+1$ times. All in all, the number of complete paths is at most:

$$2n^2 \cdot \binom{2(n+1)}{n+1} \leq 4en^2(2e)^n \leq (2e)^{2n},$$

where we use the inequalities $\binom{n}{k} \leq (e \cdot n/k)^k$ and $2n^2 \leq (2e)^n$ for $n \geq 5$. □

As a further step, we introduce the notion of edge weights.

Definition 5 (Edge-Weights). *Given any point-induced grid, we define the weight w_e and the influence region A_e of the generic edge $e = (u; v)^\dagger$ as follows:*

$$\begin{cases} A_e = [0, x] \times [u_2, v_2) \text{ and } w_e = -x \text{ if } u_1 = v_1 = x \\ A_e = [u_1, v_1) \times [y, 1] \text{ and } w_e = y \text{ if } u_2 = v_2 = y \end{cases}$$

where we adopt the convention that if $v_2 = 1$ in the first case or $v_1 = 1$ in the second case, then the corresponding interval extreme is closed.

The edge weights are crucial in establishing a connection between mechanisms and paths.

Lemma 3. *For any set S of points and complete path π on G_S , the following equality holds true:*

$$\sum_{v \in S} \text{profit}(M_\pi, v) = \sum_{e \in \pi} w_e \cdot |A_e \cap S|.$$

*Recall, $(v_1, v_2) \succeq (v'_1, v'_2)$ if $v_1 \leq v'_1$ and $v_2 \geq v'_2$.

†To avoid confusion, we denote with (a, b) a point with coordinates a and b , while $(a; b)$ represents the edge that goes from node a to node b .




Figure 3: An instance of a point-based grid, as defined in Definition 3. The points in S are highlighted in red. The area filled with yellow is the influence region A_e for the yellow horizontal edge e ; the one filled with blue is the influence region $A_{e'}$ for the blue vertical edge e' .

Proof. Consider any complete path π and denote with E_v its vertical edges and with E_h its horizontal ones. By definition of influence regions A_e , it is immediate to verify that both $\{A_e\}_{e \in E_h}$ and $\{A_e\}_{e \in E_v}$ partition the allocation region A_π . Consider then any point $v \in S$. If $v \notin A_\pi$, then M does not extract any profit from it; similarly, v does not belong to any A_e for e edge of π , so that v does not contribute to the right-hand side of the statement.

Consider now the other case: if v belongs to A_π , then there exists exactly two edges — a vertical edge e and a vertical one e' — such that $v \in A_e \cap A_{e'}$. By definition of payments, the profit induced by v according to M_π is exactly $w_e + w_{e'}$. Summing up over all the points in S yields the statement. \square

This structural lemma allows us to argue that the optimal mechanism on a fixed and finite set S of points lies in the augmented grid G_S .

Theorem 1. *For any finite set of points S , the following equality holds true:*

$$\max_{M \in \mathcal{M}} \sum_{v \in S} \text{profit}(M, v) = \sum_{v \in S} \text{profit}(M_\pi, v),$$

for some complete path π in G_S .

Proof. For any set S of points and mechanism $M \in \mathcal{M}$, we prove that there exists a complete path π in G_S such that the profit of M_π on S is an upper bound to the profit of M on the same points.

Fix then any set of points S and mechanism M with allocation region A ; we introduce the following set A_S :

$$A_S = \{v \in [0, 1]^2 \mid \exists v' \in A \cap S \text{ such that } v' \preceq v\}.$$

The set A_S is a monotone allocation region: consider in fact any $v \in A_S$ and any $\hat{v} \in [0, 1]^2$ such that $v \preceq \hat{v}$. By definition of A_S there exists $v' \in A \cap S$ such that $v' \preceq v$, which implies that $v' \preceq \hat{v}$ (by the transitive property); therefore $\hat{v} \in A_S$ as well.

For what follows, we need to prove that there exists a complete path π in G_S such that $A_S = A_\pi$. For any $v = (v_1, v_2) \in [0, 1]^2$, we define the monotone rectangle R_v as $R_v = [0, v_1] \times [v_2, 1]$; these rectangles are useful to characterize A_S . More precisely, the set A_S is the union for $v \in S \cap A$ of R_v (see Figure 6 and the constructions in Appendix C). This means that A_S is a monotone region whose boundary is composed by the union of axis-oriented orthogonal segments. These segments - with rightward and upward orientations for horizontal and vertical edges respectively - all belong to

the grid G_S , being included in the sides of the rectangles R_v for some $v \in S$. Therefore, there exists a complete path π so that A_S is bounded by the $(0,0)$ - $(0,1)$ segment, the $(0,1)$ - $(1,1)$ segment, and by π .

Finally, $A_\pi = A_S \subseteq A$, therefore $\text{profit}(M, v) \leq \text{profit}(M_\pi, v)$ for all $v \in A \cap S$ (by Observation 1). On the other hand, both M and M_π do not extract any profit from the points in $S \setminus A$, thus concluding the proof. \square

The above theorem implies that — to find the profit maximizing mechanism on the observed data — it is enough to solve a shortest path problem.

Corollary 1 (Finding the Optimal Mechanism). *Given a set S of points, it is possible to find efficiently a mechanism $M^* \in \mathcal{M}^+$ supported on G_S such that*

$$\sum_{v \in S} \text{profit}(M^*, v) = \max_{M \in \mathcal{M}} \sum_{v \in S} \text{profit}(M, v).$$

Proof. The algorithm is simple: construct the point-induced graph G_S and find a shortest complete path $\hat{\pi}$ in G_S , where the edge weights are defined as follows:

$$w_e^S = -w_e \cdot |A_e \cap S|. \quad (9)$$

Note, the graph G_S is acyclic, so it is possible to find efficiently the shortest path using the Bellman-Ford algorithm. Moreover, to restrict the optimization to complete paths, it is enough to add a dummy source node s (with outgoing zero-weight edges pointing to all the nodes on the $(0,0)$ - $(0,1)$ segment) and a dummy target node t (with incoming zero-weight edges from all the nodes on the $(0,1)$ - $(1,1)$ segment). To force the path not to contain any point outside U , the upper-left triangle, then it is enough to remove all nodes in G_S that are not in U . As already argued in Lemma 1, this is without loss of generality.

Let π^* be the complete path satisfying the equality in Theorem 1. We have the following chain of inequalities:

$$\begin{aligned} \sum_{v \in S} \text{profit}(M_{\hat{\pi}}, v) &= \sum_{e \in \hat{\pi}} w_e |A_e \cap S| && \text{(By Lemma 3)} \\ &= - \sum_{e \in \hat{\pi}} w_e^S && \text{(By weight definition in Equation (9))} \\ &\geq - \sum_{e \in \pi^*} w_e^S && (\hat{\pi} \text{ shortest path}) \\ &= \sum_{e \in \pi^*} w_e |A_e \cap S| && \text{(By weight definition in Equation (9))} \\ &= \sum_{v \in S} \text{profit}(M_{\pi^*}, v) && \text{(By Lemma 3)} \\ &= \max_{M \in \mathcal{M}} \sum_{v \in S} \text{profit}(M, v), \end{aligned}$$

where the last equality holds by the choice of π^* . This concludes the proof. \square




Figure 4: Visualization of the approximation (red) of a mechanism M (blue allocation region), with precision $\varepsilon = 1/6$. Note, the blue points are $\partial M \cap G_\varepsilon$.

4.1.2 The Simplification Procedure

We introduce a procedure that allows us to associate to each mechanism M a “simplified” version that is adherent to a certain uniform grid.

Definition 6 (Boundary). *For any mechanism M , we denote with ∂M its boundary:*

$$\partial M = \{v \in A_M \mid \forall \varepsilon > 0, ([0, 1]^2 \setminus A_M) \cap B_\varepsilon(v) \neq \emptyset\}$$

In words, each mechanism is characterized by the curve ∂M that starts from the $(0, 0)$ - $(0, 1)$ vertical segment and terminates in the $(0, 1)$ - $(1, 1)$ horizontal one. In particular, for point-based grids, the complete path that induces a mechanism is its boundary. From a different perspective, ∂M is the (topological) boundary ∂A of the corresponding allocation region A , when considering the subspace topology on $[0, 1]^2$.

Observation 2. *The boundary ∂M of a mechanism M characterizes the allocation A as any point in A is dominated by at least a point on ∂M .*

We refer to points on the curve sorted according to the natural left-to-right and bottom-up parametrization: $v \in \partial M$ comes before $v' \in \partial M$ if $v_i \leq v'_i$ for $i = 1, 2$, where this total ordering is well posed because of the monotonicity of the allocation region. In general, ∂M may be utterly complicated (see also Section C); for instance, if we look at the profit-maximizing mechanism over n points, the corresponding ∂M may have $O(n)$ “corners”. We present a generic procedure to simplify allocation regions without affecting too much the corresponding profit.

For any $\varepsilon > 0$, we denote with G_ε the grid based on the points $(i\varepsilon, i\varepsilon)$ for $i = 1, \dots, \lfloor 1/\varepsilon \rfloor$, according to Definition 3. For any mechanism M and precision ε , we denote with S_M^ε the set of points defined as the union between the points of the form $(i\varepsilon, i\varepsilon)$ for $i = 1, \dots, \lfloor 1/\varepsilon \rfloor$ and the points at the intersections of G_ε and ∂M . Note, we are intersecting a curve ∂M with the union of segments/edges in G_ε , and we adopt the convention that if ∂M and G_ε share a segment, then the intersection only contains the first and last shared point of such segment (where first and last are defined according to the edge direction or equivalently to the ordering previously specified for the curve ∂M). We denote with G_M^ε the grid induced by the points S_M^ε , again according to Definition 3.

Constructing the Simplified Mechanism. We now explain how to construct the complete path π_M^ε in G_M^ε that defines the simplified mechanism M_ε (we refer to Figure 4 for a visualization). π_M^ε

starts from the intersection of ∂M with the $(0,0)$ - $(0,1)$ segment and terminates at the intersection between ∂M and the $(0,1)$ - $(1,1)$ segment. It contains all the points in the intersection of ∂M and G_ε , so we only need to specify how to interpolate between two consecutive points there. Based on the specified total ordering for the boundary, let $v = (v_1, v_2)$ and $v' = (v'_1, v'_2)$ be two consecutive points in the intersection of ∂M and G_ε ; we have two cases:

- If there exists an edge G_M^ε that connects them, then add this edge to the path.
- If no such edge exists, consider the node u of coordinates (v'_1, v_2) , which belongs to G_M^ε . The path π_M^ε contains the directed edges $v-u$ and $u-v'$.

Definition 7 (Simplified Mechanism). *For any mechanism $M \in \mathcal{M}$ and precision ε , we denote with M_ε the simplified mechanism induced by π_M^ε .*

Algorithmic Aspects. We have presented a procedure to compute the path π_M^ε , and we observe that this can be performed efficiently by our algorithm. From an algorithmic perspective, the only unclear passage is how to find intersections of ∂M^* and G_ε , where M^* is the optimal mechanism on the points observed so far, from Corollary 1. This can be done efficiently as ∂M^* and G_ε are (finite) unions of horizontal and vertical segments/edges of known coordinates. In particular, ∂M^* contains at most $O(n)$ segments, while G_ε has $O(1/\varepsilon^2)$ segments.

Properties of the Simplified Mechanism. Before moving on, we present some simple but crucial properties of simplified mechanisms.

Proposition 2 (Properties of Simplified Mechanisms). *Consider any mechanism $M \in \mathcal{M}$ and precision ε , then the following properties are verified:*

- (i) $A_M \subseteq A_{M_\varepsilon}$
- (ii) If $v \in A_M$, then $\text{profit}(M, v) - \text{profit}(M_\varepsilon, v) \leq 2\varepsilon$
- (iii) $S_{M_\varepsilon}^\varepsilon = S_M^\varepsilon$

Proof. Consider any point $v \in A_M$: it is dominated by a point v' on the boundary of M . If v' belongs to S_M^ε then there is nothing else to prove, as $v' \in A_{M_\varepsilon}$; otherwise, there exists two consecutive points $u = (u_1, u_2)$ and $w = (w_1, w_2)$ in S_M^ε such that v' belongs to the section of ∂M between u and w . If $u_1 \neq w_1$ and $u_2 \neq w_2$, this means that v' belongs to the $[u_1, w_1] \times [u_2, w_2]$ rectangle, which dominates (in the sense of the relation \preceq) the point (u_2, w_1) , which is in the boundary of M_ε (by construction). This proves that v belongs to A_{M_ε} when the coordinates of v' are different. In particular, this argument tells us that each such point on the boundary of M belongs to a rectangle in A_{M_ε} , such that this rectangle has sides of length at most ε and is dominated by a point on the boundary of M_ε . Also, the right and lower side of this rectangle belong to ∂M_ε , by construction. The same holds - with a segment of length at most ε instead of a rectangle - when u and w coincide for one coordinate.

Consider now point (ii), and fix any point $v = (u_1, w_2) \in A_M$. We have $\text{profit}(M, v) = q - p$, as defined in Definition 2. This means that there exists two points $u = (u_1, q)$ and $w = (p, w_2)$ in ∂M . By point (i), we know that $v \in A_{M_\varepsilon}$, and that $\text{profit}(M_\varepsilon, v) = q_\varepsilon - p_\varepsilon$, induced by the points $u_\varepsilon = (u_1, q_\varepsilon)$ and $w_\varepsilon = (p_\varepsilon, w_2)$ in ∂M_ε . By the construction used in the previous point, we have that q_ε is at most $q + \varepsilon$, while p_ε is at least $p - \varepsilon$. This implies point (ii).

We finally address point (iii). This is equivalent to proving that ∂M_ε and ∂M have the same intersection with G_ε . This is true directly by construction of M_ε . \square

4.2 Approximating Sequence of Mechanisms

In this section, we show how to iteratively apply the simplification procedure to construct a sequence of approximating mechanisms, and analyze its properties.

Definition 8 (Approximating Sequence). *Let $M \in \mathcal{M}$ be any mechanism, then the sequence of approximating mechanism $\phi_h(M)$ is defined as $\phi_h(M) = M_{2^{-h}}$, for $h \in \mathbb{N}$.*

Stated differently, the sequence $\phi_h(M)$ approximates the target mechanism M via the simplification procedure (see Definition 7) for geometrically decreasing precisions 2^{-h} . We have some immediate properties that stem from the definition of $\phi_h(M)$.

Proposition 3. *For any mechanism $M \in \mathcal{M}$, and indices $h' \geq h$, it holds that*

- (Coarser Composition Rule) $\phi_h(M) = \phi_h(\phi_{h'}(M))$
- (Monotonicity) $A_{\phi_{h'}(M)} \subseteq A_{\phi_h(M)}$

Proof. From point (iii) of Proposition 2, we know that $\partial M \cap G_{2^{-h'}} = \partial \phi_{h'}(M) \cap G_{2^{-h'}}$. Since $h' \geq h$, we have that $G_{2^{-h}} \subseteq G_{2^{-h'}}$, which, in turn, implies* that $\partial M \cap G_{2^{-h}} = \partial \phi_{h'}(M) \cap G_{2^{-h}}$; since the intersection of the boundary with the grid characterizes the mechanism, this concludes the proof of the first property: $\phi_h(M) = \phi_h(\phi_{h'}(M))$.

We move to the second property. We have that $A_{\phi_h(M)} = A_{\phi_h(\phi_{h'}(M))}$ as the two mechanisms are the same, while point (i) of Proposition 2 implies that $A_{\phi_h(\phi_{h'}(M))} \supseteq A_{\phi_{h'}(M)}$. All in all, we can conclude that $A_{\phi_{h'}(M)} \subseteq A_{\phi_h(M)}$. \square

In the following proposition, we summarize the properties of approximating sequences for \mathcal{M}^+ , the subset of \mathcal{M} whose mechanisms have allocation region contained in the upper-left triangle of $[0, 1]^2$.

Proposition 4 (Properties of the Approximating Sequence). *Let M be any mechanism in \mathcal{M}^+ , the following properties holds true for any point $v \in [0, 1]^2$:*

- (i) $\text{profit}(\phi_{h'}(M), v) - \text{profit}(\phi_h(M), v) \leq 2 \cdot 2^{-h}$ for any $h' \geq h$.
- (ii) $\text{profit}(M, v) - \text{profit}(\phi_h(M), v) \leq 2 \cdot 2^{-h}$ for any h
- (iii) If $v \notin A_{\phi_h(M)}$, then $\text{profit}(\phi_{h'}(M), v) = 0$ for all $h' \geq h$
- (iv) If $v \in A_{\phi_h(M)}$, then $\text{profit}(\phi_{h'}(M), v)$ is non-decreasing for any $h' \leq h$

Proof. We start proving point (ii), which implies, by the Coarser Composition Rule (Proposition 3) point (i). Consider any point $v \in [0, 1]^2$. We have three cases:

- If v does not belong to the allocation region of $\phi_h(M)$, then it does not belong to that of M either (by point (i) of Proposition 2 we know that $A_M \subseteq A_{\phi_h(M)}$). Therefore $\text{profit}(M, v) = \text{profit}(\phi_h(M), v) = 0$.
- If v belongs to the allocation region of M , then it also belongs to that of $\phi_h(M)$ (once again, by point (i) of Proposition 2), moreover, by point (ii) of Proposition 2, we have that:

$$\text{profit}(M, v) - \text{profit}(\phi_h(M), v) \leq 2 \cdot 2^{-h}.$$

* $A \cap C = B \cap C$ implies $A \cap C' = B \cap C'$ if $C' \subseteq C$.

- Consider now the situation in which v belongs to the allocation region of $\phi_h(M)$, but not of M . By construction, v belongs to a rectangle with sides of length at most 2^{-h} , such that the right and lower sides of the rectangle belong to $\partial\phi_h(M)$. Therefore, under $\phi_h(M)$, the price paid to the seller is at most $v_1 + 2^{-h}$, and that paid by the buyer is at least $v_2 - 2^{-h}$, for a revenue of at least $v_2 - v_1 - 2 \cdot 2^{-h}$. If $v \in U$, the upper-left corner of $[0, 1]^2$, then $v_2 - v_1 \geq 0$, therefore $\text{profit}(\phi_h(M), v) \geq -2 \cdot 2^{-h}$, which verifies the inequality since $\text{profit}(M, v) = 0$ in this case. If v is not in U , then it cannot be too far from U , as it is contained in the allocation region of $\phi_h(M)$, and the allocation region of M is contained in U . In particular, since $M \in \mathcal{M}^+$, the lower-left and upper-right vertexes of the rectangle described above need to belong to U , since by construction the union of the lower side and the right side connects two points in $A_M \subseteq U$. This means that the lower-right vertex u , whose coordinates give the Myerson's prices $q = u_2$ and $p = u_1$ for v , is such that $q - p$ is at most -2^{-h} . This follows by considering the vertical and horizontal distance between u and the main diagonal characterizing U .

Consider now the points (iii) and (iv). We know that the set sequence $A_{\phi_h(M)}$ is monotonically decreasing with respect to the inclusion (Proposition 3), therefore $A_{\phi_{h'}(M)} \subseteq A_{\phi_h(M)}$ for any $h' \geq h$. This means that $v \notin A_{\phi_h(M)}$ implies $v \notin A_{\phi_{h'}(M)}$, and thus $\text{profit}(A_{\phi_{h'}(M)}, v) = 0$. Conversely, if $v \in A_{\phi_{h'}(M)}$, then it also belongs to $A_{\phi_h(M)}$ and, by the monotonicity of the profit function (Observation 1), we have that $\text{profit}(\phi_{h'}(M), v) \geq \text{profit}(\phi_h(M), v)$. \square

When the valuation pair is sampled from some distribution \mathcal{D} , we may use $\mathbb{E}[\text{profit}(M)]$ instead of $\mathbb{E}_{V \sim \mathcal{D}}[\text{profit}(M, V)]$ when there is no ambiguity on the distribution we are referring to. We have the following, crucial structural lemma on approximating sequences.

Proposition 5 (Local Dichotomy of Approximating Sequences). *Let M be any mechanism in \mathcal{M}^+ , and h be any integer. Then for any distribution at least one of the following two properties holds:*

- (i) $\mathbb{E}[\text{profit}(\phi_h(M)) - \text{profit}(\phi_{h+1}(M))] \geq 31 \cdot 2^{-h}$
- (ii) $\mathbb{E}[(\text{profit}(\phi_{h+1}(M)) - \text{profit}(\phi_h(M)))^2] \leq 136 \cdot 2^{-h}$

Proof. We partition the $[0, 1]^2$ square according to the revenue that the various points induce. In particular, we denote with B the points that induce significant different revenues in the two mechanisms:

$$B = \left\{ v \in [0, 1]^2 \mid \text{profit}(\phi_h(M), v) - \text{profit}(\phi_{h+1}(M), v) > \frac{2}{2^h} \right\}$$

Moreover, for $i = 1, \dots, h-1$ we define:

$$B_i = \left\{ v \in B \mid \frac{2^i}{2^h} < \text{profit}(\phi_h(M), v) \leq \frac{2^{i+1}}{2^h} \right\}.$$

Claim 1. *The sets B_i are a partition of B . Moreover for $v \in B$ the following two formulae hold: $\text{profit}(\phi_{h+1}(M), v) = 0$ and $\text{profit}(\phi_h(M), v) > 2/2^h$.*

Proof of the Claim. The revenue function is at most 1. Therefore, to prove that $\{B_i\}_{i=1}^{h-1}$ is a partition of B , we only need to prove that no $v \in B$ is such that $\text{profit}(\phi_h(M), v) \leq 2/2^h$. Clearly, if v is in B , then it cannot belong to the allocation region of both mechanisms (otherwise Proposition 4 would imply that $\text{profit}(\phi_{h+1}(M), v) \geq \text{profit}(\phi_h(M), v)$). Similarly, if v does not belong to the allocation region of $\phi_h(M)$, then it induces zero revenue for both $\phi_h(M)$ and $\phi_{h+1}(M)$, which is incompatible with being in B . The only possibility is that v belongs to the allocation region of $\phi_h(M)$ but not of $\phi_{h+1}(M)$. This implies that $\text{profit}(\phi_{h+1}(M), v) = 0$ and $\text{profit}(\phi_h(M), v) > 2/2^h$ by definition of B . \square

If we focus on the fraction of revenue generated by the points in B , we get the following:

$$\begin{aligned}
\mathbb{E}_{V \sim \mathcal{D}} [\text{profit}(\phi_h(M), V) \cdot \mathbb{1}_{\{V \in B\}}] &= \sum_{i=1}^h \mathbb{E}_{V \sim \mathcal{D}} [\text{profit}(\phi_h(M), V) \cdot \mathbb{1}_{\{V \in B_i\}}] && \text{(By Claim 1)} \\
&\geq \sum_{i=1}^h \frac{2^i}{2^h} \cdot \mathbb{P}_{V \sim \mathcal{D}}(V \in B_i) && \text{(By def. of } B_i) \\
&\geq \underbrace{\sum_{i=1}^h \frac{2^{2i}}{2^{2h}} \cdot \mathbb{P}_{V \sim \mathcal{D}}(V \in B_i)}_{(\#)}, && (10)
\end{aligned}$$

where, in the last inequality we used that $h \geq i$, and thus $2^{i-h} \geq 2^{2(i-h)}$. We now divide the analysis into two cases, according to whether the right-hand side of Equation (10) (let's denote it with $(\#)$) is larger or smaller than $33 \cdot 2^{-h}$.

First Case. If $(\#) \geq 33 \cdot 2^{-h}$, then from Equation (10) we can immediately conclude that $\mathbb{E} [\text{profit}(\phi_h(M), V) \cdot \mathbb{1}_{\{V \in B\}}] \geq 33 \cdot 2^{-h}$, and we can write the following:

$$\begin{aligned}
\mathbb{E} [\text{profit}(\phi_h(M))] &= \mathbb{E} [\text{profit}(\phi_h(M), V) \cdot \mathbb{1}_{\{V \in B\}}] + \mathbb{E} [\text{profit}(\phi_h(M), V) \cdot \mathbb{1}_{\{V \notin B\}}] \\
&\geq 33 \cdot 2^{-h} + \mathbb{E} [\text{profit}(\phi_h(M), V) \cdot \mathbb{1}_{\{V \notin B\}}] && \text{(By bound on } (\#)) \\
&\geq 33 \cdot 2^{-h} + \mathbb{E} [(\text{profit}(\phi_{h+1}(M), V) - \frac{2}{2^h}) \cdot \mathbb{1}_{\{V \notin B\}}] \\
&\geq \mathbb{E} [\text{profit}(\phi_{h+1}(M))] + 31 \cdot 2^{-h}.
\end{aligned}$$

Note, in the second to last inequality we applied point (i) of Proposition 4, while in the last passage we used the observation that $\text{profit}(\phi_{h+1}(M), v) = 0$ for $v \in B$.

Second case. If $(\#) < 33 \cdot 2^{-h}$, then we can look at the second moment of the difference:

$$\begin{aligned}
&\mathbb{E} [(\text{profit}(\phi_{h+1}(M)) - \text{profit}(\phi_h(M)))^2] \\
&= \mathbb{E} [\text{profit}(\phi_h(M), V)^2 \mathbb{1}_{\{V \in B\}}] + \mathbb{E} [(\text{profit}(\phi_{h+1}(M)) - \text{profit}(\phi_h(M)))^2 \mathbb{1}_{\{V \notin B\}}] \\
&\leq \mathbb{E} [\text{profit}(\phi_h(M), V)^2 \mathbb{1}_{\{V \in B\}}] + \frac{4}{2^{2h}} \mathbb{P}(V \notin B) \\
&\leq 4 \cdot \sum_{i=1}^h \mathbb{P}(V \in B_i) \frac{2^{2i}}{2^{2h}} + \frac{4}{2^{2h}} \leq 136 \cdot 2^{-h}. && \text{(By the case analysis and Claim 1)}
\end{aligned}$$

Note, in the first equality we use once again that $\text{profit}(\phi_{h+1}(M), v) = 0$ for $v \in B$. In the first inequality we use that for $v \notin B$ it holds that $\text{profit}(\phi_h(M), v) - \text{profit}(\phi_{h+1}(M), v) \leq 2 \cdot 2^{-h}$ which, combined with Property (i) of Proposition 4 implies that $|\text{profit}(\phi_h(M), v) - \text{profit}(\phi_{h+1}(M), v)| \leq 2 \cdot 2^{-h}$. \square

4.3 Target Net and Regret Bound

The approximating sequences and their properties play a crucial role in our results, as they allow us to control the learning guarantees of the simplified mechanism M^t from Section 4.1 that is played by the algorithm SIMPLIFY-THE-BEST-MECHANISM. In the rest of the section we forget about the online nature of the problem at hand and study the associated *sample-complexity* problem, then

concluding our argument via the standard offline-to-online reduction for online learning with full feedback. Fix then a distribution \mathcal{D} and consider a sample (multi-)set S of n points drawn i.i.d. from \mathcal{D} . For any set I of indices in $[n] = \{1, 2, \dots, n\}$, we denote with $S(I)$ the projection of S along the indices in I , i.e., all the samples corresponding to the indices in I . Note that $S(I)$ could be a multi-set. Since we consider sets of indices of cardinality $4/\varepsilon$, we make the assumption that $n \geq 4/\varepsilon$.

Definition 9 (Data-Dependent Net). *Fix any precision $\varepsilon = 2^{-H}$ for some $H \in \mathbb{N}$. Consider n i.i.d. samples S from \mathcal{D} , and let I be any subset of $4/\varepsilon$ indices in $[n]$. We denote with $G_{S(I)}^\varepsilon$ the grid induced by the points $S(I)$ and by the uniform ε -grid. Furthermore, we denote with $\mathcal{M}_{S(I)}^\varepsilon$ the set of all mechanisms induced by complete paths on $G_{S(I)}^\varepsilon$.*

Recall, the grid is defined as in Definition 3, while for complete paths we refer to Definition 4. We complement this definition by providing an upper bound on the cardinality of $\cup_I \mathcal{M}_{S(I)}^\varepsilon$, when the set of indices I span over all the subset of $4/\varepsilon$ indices in $[n]$.

Proposition 6. *Let $n > 16$, then the following holds:*

$$\log \left(\left| \bigcup_{\substack{I \subseteq [n] \\ |I|=4/\varepsilon}} \mathcal{M}_{S(I)}^\varepsilon \right| \right) \leq \frac{10}{\varepsilon} \log n.$$

Proof. There are $\binom{n}{4/\varepsilon}$ ways of choosing the indices I , and, for each such choice, each grid $G_S^\varepsilon(I)$ is induced by $5/\varepsilon$ points. We have then the following bound (by Lemma 2):

$$\left| \bigcup_{\substack{I \subseteq [n] \\ |I|=4/\varepsilon}} \mathcal{M}_{S(I)}^\varepsilon \right| \leq (2e)^{10/\varepsilon} \binom{n}{4/\varepsilon} \leq 4^{1/\varepsilon} e^{14/\varepsilon} n^{4/\varepsilon} \leq n^{10/\varepsilon},$$

where we use the inequality $\binom{n}{k} \leq (e \cdot n/k)^k$, and the assumption that $n > 16$. Applying the logarithm yields the desired bound. \square

Consider the optimal mechanism M^\star on the sample S , as constructed in Section 4.1.1 and Corollary 1; it undergoes the simplification procedure presented in Section 4.1 to generate the simplified mechanism $\phi_H(M^\star)$ (which is the mechanism actually played by the algorithm). We want to argue that, with high-probability, the profit generated by $\phi_H(M^\star)$ on the sample S is a good approximation of its expected profit with respect to a new sample from \mathcal{D} . The mechanism $\phi_H(M^\star)$ is induced by a complete path on the point induced grid $G_{M^\star}^\varepsilon$, where $\varepsilon = 2^{-H}$; moreover, it is characterized by at most $4/\varepsilon$ points in G_S , and this clearly also works as an upper bound on the number of characterizing elements from S . This means that, for some set of indices I of cardinality at most $4/\varepsilon$, the boundary of M^\star is a complete path in $G_{S(I)}^\varepsilon$. We thus have the following lemma.

Lemma 4. *Fix any precision $\varepsilon = 2^{-H}$ for some $H \in \mathbb{N}$, and consider n i.i.d. samples S from distribution \mathcal{D} . Then the simplified mechanism $\phi_H(M^\star)$ belongs to $\mathcal{M}_{S(I)}^\varepsilon$ for some set I of at most $4/\varepsilon$ indices.*

Proof. By construction from Section 4.1, the mechanism $\phi_H(M^\star)$ is characterized by $\partial M^\star \cap G_\varepsilon$. Additionally, again by construction, this time from Corollary 1, we have that each element of the intersection depends on two points of G_S , since the boundary of M^\star is the union of segments from G_S , with each element of the union thus connecting two points in the sample grid. Therefore, proving the statement is equivalent to stating that the amount of intersecting elements between boundary and ε -grid is bounded by $2/\varepsilon$.

To prove this, notice that a complete path must go from the segment $(0, 0)$ - $(0, 1)$ to the segment $(0, 1)$ - $(1, 1)$; similarly to the proof of Lemma 2, by monotonicity of the region induced by a complete path, its boundary can only enter up to $2/\varepsilon$ cells in the G_ε grid ($1/\varepsilon$ “right” movements and $1/\varepsilon$ “up” movements), with each cell visited at most once. Also, we can add an element to the intersection if and only if the path enters one of the cells by its lower or left sides, thus meaning that the cardinality of the intersection is bounded by the amount of cells that can be possibly explored, i.e. $2/\varepsilon$. \square

Unfortunately, it is unclear whether a fast enough uniform learning bound on the whole class $\cup_I \mathcal{M}_{S(I)}^\varepsilon$ is achievable. In particular, a naive union bound (using Proposition 6) would imply a suboptimal sample complexity of order $\approx 1/\varepsilon^3$. To further restrict the class of mechanism we want to learn, we introduce the family of candidate mechanisms.

Definition 10 (Candidate Mechanisms). *Fix any precision $\varepsilon = 2^{-H}$ for some level $H \in \mathbb{N}$. We say that a mechanism $M \in \mathcal{M}^+$ is a candidate if the following condition holds for all $h = 0, 1, \dots, H - 1$:*

$$\mathbb{E}[\text{profit}(\phi_h(M)) - \text{profit}(\phi_{h+1}(M))] < 4 \cdot 2^{-h}. \quad (11)$$

We denote with \mathcal{C} the set of all candidate mechanisms.

The idea underlying this definition is that a mechanism whose approximating sequence does not uniformly respect Equation (11) is clearly suboptimal (with respect to distribution \mathcal{D}), and we do not need to learn it with high precision. We are now ready to present our target net.

Definition 11 (Target Net). *For any distribution \mathcal{D} and precision parameter $\varepsilon = 2^{-H}$, consider n i.i.d. samples S from \mathcal{D} . We construct a target net $\mathcal{M}^*(S)$ as follows:*

$$\mathcal{M}^*(S) = \bigcup_{\substack{I \subseteq [n] \\ |I|=4/\varepsilon}} \mathcal{M}_{S(I)}^\varepsilon \cap \mathcal{C}.$$

In English, the target net is composed by all the candidate mechanisms in that correspond to some complete path in $G_{S(I)}^\varepsilon$, for some I of cardinality $4/\varepsilon$.

Observation 3 (The Structure of Our Net). *Before moving on, we observe that our target net $\mathcal{M}^*(S)$ is both data and distribution dependent! While the family of candidate mechanisms \mathcal{C} is distribution dependent, the families $\mathcal{M}_{S(I)}^\varepsilon$ are indeed data-dependent, as they are induced by the realized samples. In particular, each mechanism in our net only depends on a small number of samples (i.e., $4/\varepsilon$), so that the remaining samples are enough to attain concentration.*

We now want to argue that the mechanism computed by the algorithm, i.e. $\phi_H(M^*)$, where M^* is the empirical risk minimizer, belongs to the family $\mathcal{M}^*(S)$ where we want to prove uniform learnability.

Theorem 2. *Fix a precision $\varepsilon = 2^{-H}$ and a probability parameter δ . Consider n i.i.d. samples S , then $\phi_H(M^*) \in \mathcal{M}^*(S)$ with probability at least $(1 - \delta)$, as long as $n \geq \frac{25^2}{\varepsilon^2} \cdot (\log(20/\delta) + \log(20/\varepsilon))$.*

Proof. Let $M^* \in \mathcal{M}^+$ be the empirical risk minimizer on the samples S computed as in Corollary 1, i.e. the mechanism supported in G_S such that:

$$\sum_{v \in S} \text{profit}(M^*, v) \geq \sum_{v \in S} \text{profit}(M, v), \quad \forall M \in \mathcal{M}. \quad (12)$$

We have already argued in Lemma 4 that $\phi_H(M^*) \in \mathcal{M}_{S(I^*)}^\varepsilon$ for some set of indices I^* of cardinality $4/\varepsilon$, so we are left with proving that, with high probability, $\phi_H(M^*)$ also belongs to \mathcal{C} .

If we focus on the samples in $S \setminus S(I^*)$, then we get a weaker version of Equation (12), namely*:

$$\mathbf{profit}(M^*, S \setminus S(I^*)) \geq \mathbf{profit}(M, S \setminus S(I^*)) - \frac{\varepsilon}{2}, \quad \forall M \in \mathcal{M} \quad (13)$$

To justify this, we have the following chain of inequalities, for all $M \in \mathcal{M}$:

$$\begin{aligned} \sum_{p \in S \setminus S(I^*)} \mathbf{profit}(M^*, p) + \frac{4}{\varepsilon} &\geq \sum_{p \in S} \mathbf{profit}(M^*, p) && (\mathbf{profit} \leq 1 \text{ and } |I^*| = 4/\varepsilon) \\ &\geq \sum_{p \in S} \mathbf{profit}(M, p) && (\text{By Equation (12)}) \\ &\geq \sum_{p \in S \setminus S(I^*)} \mathbf{profit}(M, p) - \frac{4}{\varepsilon}. && (\mathbf{profit} \geq -1 \text{ and } |I^*| = 4/\varepsilon) \end{aligned}$$

By assumption on the sample cardinality, we have that $n \geq 32/\varepsilon^2$, which implies that the cardinality of the multi-set $S \setminus S(I)$ is at least $n/2$. Dividing both terms of the above chain of inequalities by $|S \setminus S(I)|$, we get Equation (13).

As a second step, we argue, that for all levels $h = 0, \dots, H-1$ the following inequality needs to hold:

$$\mathbf{profit}(\phi_h(M^*), S \setminus S(I^*)) - \mathbf{profit}(\phi_{h+1}(M^*), S \setminus S(I^*)) < 2 \cdot 2^{-h}. \quad (14)$$

To see this, assume towards contradiction that there exists h such that Equation (14) is not verified, then we would have the following:

$$\begin{aligned} \mathbf{profit}(\phi_h(M^*), S \setminus S(I^*)) &\geq 2/2^h + \mathbf{profit}(\phi_{h+1}(M^*), S \setminus S(I^*)) \\ &\geq 1/2^h + \mathbf{profit}(M^*, S \setminus S(I^*)) && (\text{Point (ii) of Prop. 4}) \\ &\geq 1/2^h - \varepsilon/2 + \mathbf{profit}(\phi_h(M^*), S \setminus S(I^*)). && (\text{By Eq. 13}) \end{aligned}$$

The last inequality yields a contradiction, as $1/2^h \geq 1/2^H = \varepsilon > \varepsilon/2$.

The last step in this proof consists in proving that, with high probability, all mechanisms in $\mathcal{M}_{S(I)}^\varepsilon$ that exhibit the inequality in Equation (14) at all levels h and whose allocation regions are included in U , the upper-left corner, belong to $\mathcal{M}^*(S)$.

Consider any mechanism $M \in (\mathcal{M}_{S(I)}^\varepsilon \cap \mathcal{M}^+) \setminus \mathcal{C}$, for some set of $4/\varepsilon$ indices I . If M is not a candidate, it means that there exists a level h such that $\mathbb{E}[\mathbf{profit}(\phi_h(M)) - \mathbf{profit}(\phi_{h+1}(M))] \geq 4 \cdot 2^{-h}$. We want to apply concentration on X , the sum for $v \in S \setminus S(I)$ of the random variables of the form:

$$X_v = \frac{1}{4}[\mathbf{profit}(\phi_h(M), v) - \mathbf{profit}(\phi_{h+1}(M), v) + 2 \cdot 2^{-h}].$$

The X_v lie in $[0, 1]$ by property (i) of Proposition 4 and because the revenue is at most 1. Moreover, we are assuming that $\mathbb{E}[X_v] \geq 3/2 \cdot 2^{-h}$. We also introduce the event $\mathcal{E}_{M,h}$ that is realized when X is way smaller than expected:

$$\mathcal{E}_{M,h} = \left\{ \mathbf{profit}(\phi_h(M), S \setminus S(I)) - \mathbf{profit}(\phi_{h+1}(M), S \setminus S(I)) < 2 \cdot 2^{-h} \right\}.$$

*We introduce the following notation: $\mathbf{profit}(M, T)$, where T is a set of points, denote the average profit of mechanism M on the points in T .

We have the following, where $n' = |S \setminus S(I)|$:

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_{M,h}) &= \mathbb{P}\left(X < 2^{-h} \cdot n'\right) \\
&= \mathbb{P}\left(X - \mathbb{E}[X] < 2^{-h} \cdot n' - \mathbb{E}[X]\right) \\
&\leq \mathbb{P}\left(X - \mathbb{E}[X] < -1/3 \cdot \mathbb{E}[X]\right) \\
&\leq e^{-\frac{1}{18}\mathbb{E}[X]} && \text{(By (i) of Theorem 12)} \\
&\leq e^{-\frac{\varepsilon}{24}n} && (15)
\end{aligned}$$

where in the first and last inequalities we used monotonicity of probability and that $\mathbb{E}[X] \geq n' \cdot 3/2 \cdot 2^{-h} \geq 3/4 \cdot \varepsilon n$, where this last inequality comes again from the observation that $n' > 1/2 \cdot n$ under our assumption on n .

We now need to union bound over all possible bad mechanisms in $\mathcal{M}_{S(I)}^\varepsilon$ and levels h for which $\mathbb{E}[\text{profit}(\phi_h(M)) - \text{profit}(\phi_{h+1}(M))] \geq 4 \cdot 2^{-h}$: thus, by Proposition 6 (and by noting that, since $\varepsilon = 2^{-H}$, there are at most $2 \log 1/\varepsilon$ levels) we have that the probability that $\phi_H(M^*)$ belongs to $\mathcal{M}^*(S)$ is at least:

$$2e^{-\frac{\varepsilon}{24}n} \cdot e^{10/\varepsilon \log n} \cdot \log \frac{1}{\varepsilon} \leq \delta,$$

by how we set the number of samples n . □

The next step consists in arguing that it is possible to learn the mechanisms in the target net $\mathcal{M}^*(S)$ fast enough in expectation. The proof of this statement is postponed until Section 4.4.

Theorem 3 (Uniform Learning in $\mathcal{M}^*(S)$). *Fix any precision $\varepsilon = 2^{-H}$ for some $H \in \mathbb{N}$. Let S be an i.i.d. sample of size $n \geq \frac{32}{\varepsilon^2} \log^3(1/\varepsilon)$. The following inequality holds:*

$$\mathbb{E}_S \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S) - \mathbb{E}[\text{profit}(M)] \right| \right] \leq 794\varepsilon.$$

We have now all the ingredients to prove our theorem: the algorithm SIMPLIFY-THE-BEST-MECHANISM achieves a nearly tight $\tilde{O}(\sqrt{T})$ regret.

Theorem 4 (Regret Bound). *Consider the profit maximization problem in bilateral trade against the stochastic i.i.d. adversary. The regret of SIMPLIFY-THE-BEST-MECHANISM is at most $C\sqrt{T \log^3 T}$, for some constant C .*

Proof. We first study the instantaneous regret of the generic time step $t + 1$. The algorithm has access to the first t samples S^t received, computes the profit-maximizing mechanism $M_{t+1}^* \in \mathcal{M}^+$ on S^t and posts M^{t+1} , the version of M_{t+1}^* adapted to the grid with precision ε_{t+1} . We denote with \mathcal{E}_1 the event that M^{t+1} belongs to $\mathcal{M}^*(S^t)$. Furthermore, let \tilde{M} be any mechanism that verifies

$$\sup_{M \in \mathcal{M}} \mathbb{E}[\text{profit}(M)] \leq \mathbb{E}[\text{profit}(\tilde{M})] + \varepsilon_{t+1}, \quad (16)$$

such mechanism exists by definition of sup. Denote with \mathcal{E}_2 the event that $\mathbb{E}[\text{profit}(\tilde{M})]$ is ε_{t+1} approximated on the t samples. We introduce the clean event $\mathcal{E} = \mathcal{E}_1 \cap \mathcal{E}_2$. By union bound, our choice of $\varepsilon_{t+1} = 200\sqrt{\log^3 T/t}$, standard Hoeffding concentrations (for \mathcal{E}_2) and Theorem 2 it holds that $\mathbb{P}(\mathcal{E}) \geq 1 - 2/T$.

Conditioning on \mathcal{E} on the samples S^t , we have the following chain of inequalities:

$$\begin{aligned}
\sup_{M \in \mathcal{M}} \mathbb{E} [\text{profit}(M)] &\leq \mathbb{E} \left[\frac{1}{t} \sum_{v \in S^t} \text{profit}(\tilde{M}, v) \middle| \mathcal{E} \right] + 2\varepsilon_{t+1} && \text{(By event } \mathcal{E}_2 \text{ and Equation (16))} \\
&\leq \mathbb{E} \left[\frac{1}{t} \sum_{v \in S^t} \text{profit}(M_{t+1}^*, v) \middle| \mathcal{E} \right] + 2\varepsilon_{t+1} && (M_{t+1}^* \text{ empirical maximizer on } S_t) \\
&\leq \mathbb{E} \left[\frac{1}{t} \sum_{v \in S_t} \text{profit}(M^{t+1}, v) \middle| \mathcal{E} \right] + 4\varepsilon_{t+1} && \text{(By Property (ii) of Proposition 4)} \\
&\leq \mathbb{E} \left[\text{profit}(M^{t+1}) \middle| \mathcal{E} \right] + 1590\varepsilon_{t+1}. && (17)
\end{aligned}$$

Note, the expectation in the left-hand-side term is with respect to a fresh sample, while that on the right-hand-side terms is with respect to the first t samples observed, conditioning on the clean event. Moreover, the last step of Equation (17) descends from Theorem 3 noting that $\mathbb{P}(\mathcal{E}) \geq 1/2$:

$$\mathbb{E} \left[\left| \text{profit}(M^{t+1}, S_t) - \mathbb{E} [\text{profit}(M^{t+1})] \right| \cdot \mathbb{1}_{\mathcal{E}} \right] \leq \mathbb{E} \left[\max_{M \in \mathcal{M}^*(S_t)} \left| \text{profit}(M, S_t) - \mathbb{E} [\text{profit}(M)] \right| \right] \leq 794\varepsilon_{t+1}$$

Note, to apply the concentration result for $\mathcal{M}^*(S)$ of the actual mechanism M^{t+1} we need to condition on the high-probability event that M^{t+1} does belong to $\mathcal{M}^*(S)$.

Equation (17) gives us the crucial ingredient to provide a bound on the expected instantaneous regret. By the law of total probability, and by the fact that \mathcal{E}^C has probability at most $2/T$ we get:

$$\sup_{M \in \mathcal{M}} \mathbb{E} [\text{profit}(M)] - \mathbb{E} [\text{profit}(M^{t+1})] \leq 1590\varepsilon_{t+1} + \frac{2}{T}$$

Summing up the instantaneous regret for $t = 0, \dots, T-1$ yields the desired regret bound. \square

4.4 The Sample Complexity of Good Mechanisms

This section is devoted to proving Theorem 3, that we restate here.

Theorem 3 (Uniform Learning in $\mathcal{M}^*(S)$). *Fix any precision $\varepsilon = 2^{-H}$ for some $H \in \mathbb{N}$. Let S be an i.i.d. sample of size $n \geq \frac{32}{\varepsilon^2} \log^3(1/\varepsilon)$. The following inequality holds:*

$$\mathbb{E}_S \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S) - \mathbb{E} [\text{profit}(M)] \right| \right] \leq 794\varepsilon.$$

We decompose the proof of the theorem in many lemmata. For the coupling argument of this section, let S' be a (multi-)set of n new i.i.d. samples from \mathcal{D} , independent of S . Also, let $I_M \subset [n]$ be a set of indexes satisfying $M \in \mathcal{M}_{S(I_M)}^\varepsilon$. We have the following lemma.

Lemma 5 (Adding the Ghost Samples S'). *The following inequality holds:*

$$\begin{aligned}
&\mathbb{E}_S \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S) - \mathbb{E} [\text{profit}(M)] \right| \right] \\
&\leq \mathbb{E}_{S, S'} \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \right] + \varepsilon
\end{aligned}$$

Proof. Let $M \in \mathcal{M}^*(S)$. We start by rewriting $\mathbb{E}[\text{profit}(M)]$ as the expectation of $\text{profit}(M, S')$ over the distribution of S' . Fix any realization of the samples S , we have the following:

$$\begin{aligned}
& |\text{profit}(M, S) - \mathbb{E}_{S'}[\text{profit}(M, S')]| \\
& \leq \frac{1}{|S|} \left[\sum_{i \in I_M} |\text{profit}(M, p_i) - \mathbb{E}_{p'_i}[\text{profit}(M, p'_i)]| + \mathbb{E}_{S'} \left[\left| \sum_{i \notin I} (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \right| \right] \right] \\
& \leq \frac{2|I_M|}{|S|} + \frac{|S| - |S(I_M)|}{|S|} \mathbb{E}_{S'} \left[\left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \right] \\
& \leq \frac{8}{\varepsilon n} + \mathbb{E}_{S'} \left[\left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \right],
\end{aligned}$$

where the second inequality holds because the revenue function's image is contained in $[-1, 1]$. The third one follows from Definition 11: every mechanism M in $\mathcal{M}^*(S)$ is characterized by a set of indices I_M of size $4/\varepsilon$. So far we have considered any mechanism in $\mathcal{M}^*(S)$, we can thus apply the max:

$$\begin{aligned}
& \max_{M \in \mathcal{M}^*(S)} |\text{profit}(M, S) - \mathbb{E}_{S'}[\text{profit}(M, S')]| \\
& \leq \frac{8}{\varepsilon n} + \max_{M \in \mathcal{M}^*(S)} \mathbb{E}_{S'} \left[\left| \text{profit}(M, S \setminus S(I)) - \text{profit}(M, S' \setminus S'(I)) \right| \right] \\
& \leq \frac{8}{\varepsilon n} + \mathbb{E}_{S'} \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S \setminus S(I)) - \text{profit}(M, S' \setminus S'(I)) \right| \right]
\end{aligned}$$

The last inequality is Jensen's. The result then follows by taking the expectation with respect to S , and by the assumption on n , which implies $n \geq 8/\varepsilon^2$. \square

The next step consists in symmetrizing the random variables we are trying to learn. In particular, for this step to work, we need that the random variables we are taking the max over to be independent of the family of functions that are contained in the max. Stated differently, the family $\mathcal{M}^*(S)$ depends on the samples $S(I)$, therefore, we cannot use those samples in the symmetrization argument! Moreover, to maintain symmetry, we cannot restrict our attention to the sole mechanisms generated by the samples S . We introduce the symmetric family $\mathcal{M}^*(S, S')$ defined as follows:

$$\mathcal{M}^*(S, S') = \bigcup_{\substack{I \subseteq [n] \\ |I|=4/\varepsilon}} \bigcup_{J \subseteq I} \mathcal{M}_{S(I \setminus J) \cup S'(J)}^\varepsilon.$$

Stated differently, a mechanism in $\mathcal{M}^*(S, S')$ is supported on the standard ε -grid, augmented with some points with indexes J of S' and some points with indexes $I \setminus J$ of S . For a mechanism $M \in \mathcal{M}^*(S, S')$ we define I_M in the same way as before. We also adopt the convention that n' denotes $n - 4/\varepsilon$.

Lemma 6 (Symmetrization). *The following inequality holds:*

$$\begin{aligned}
& \mathbb{E}_{S, S'} \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \right] \\
& \leq 2 \mathbb{E}_{S, S', \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot r_p \right| \right],
\end{aligned}$$

where \mathbf{r} is a n -dimensional Rademacher vector.

Proof. We introduce the vector \mathbf{r} of n Rademacher random variables, indexed by the points p (or, alternatively, by the indices in $[n]$). Let I be a set of indexes of cardinality $4/\varepsilon$ and $M \in \mathcal{M}_{S(I)}^\varepsilon$; therefore, $I_M = I$. For every point $p_i \in S \setminus S(I)$ and $p'_i \in S' \setminus S'(I)$, the following equality holds in distribution*:

$$\text{profit}(M, p_i) - \text{profit}(M, p'_i) \stackrel{(d)}{=} (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \cdot r_i.$$

Similarly, if we now take any mechanism $M \in \mathcal{M}_{S(I)}^\varepsilon$, the following equality holds in distribution:

$$\left| \sum_{i \notin I} (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \right| \stackrel{(d)}{=} \left| \sum_{i \notin I} (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \cdot r_i \right|.$$

We now would like to argue that, if we take the max over all mechanisms in $\mathcal{M}^*(S)$, then we maintain the equal distribution. However, this is not the case: the role of S and S' is not symmetric! The sample p_i determines the structure of the mechanisms in $\mathcal{M}_{S(I)}^\varepsilon$ for I that contains i , *but this is not the case for p'_i !* This is why we need to consider the wider family $\mathcal{M}^*(S, S')$. Note, $\mathcal{M}^*(S) \subset \mathcal{M}^*(S, S')$, therefore:

$$\begin{aligned} & \max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \\ & \leq \max_{M \in \mathcal{M}^*(S, S')} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right|. \end{aligned}$$

We then focus on upper bounding the expected value (with respect to S and S') of the right-hand side term. To avoid the cumbersome dependence of the family of mechanisms involved in the max on the sampled points, we introduce the following, equivalent stochastic procedure to extract the two independent samples S and S' : first, two independent and identical samples T_1 and T_2 are drawn, then for each index, an independent Bernoulli X_i is drawn: if $X = 1$, then p_i is the i^{th} sample in T_1 and p'_i is the corresponding sample in T_2 , otherwise, the converse happens. This procedure induces the random split of the observations into S and S' .

Fix any outcome of the samples T_1 and T_2 , they deterministically induce the family of mechanisms $\mathcal{M}^*(S, S') = \mathcal{M}^*(T_1, T_2)$. So we can focus on coupling the impact of X_i with that of the Rademacher vector \mathbf{r} . For any mechanism $M \in \mathcal{M}^*(S, S')$ and any index $i \notin I_M$, we have the following equality, in distribution:

$$(\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \stackrel{(d)}{=} (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \cdot r_i.$$

Note: the randomness on the left-hand side is *only* with respect to X_i , as the values of the i^{th} samples from T_1 and T_2 are fixed[†]; the right-hand side also depends on the Rademacher variable r_i . Let m be the maximum cardinality of $\mathcal{M}^*(S, S')$ when $|S| = n$. Therefore, any vector indexed by $M \in \mathcal{M}^*(S, S')$ and $i \notin I_M$ belongs to $[-1, 1]^{m \cdot (n - 4/\varepsilon)}$, adopting the convention that exceeding entries are null whenever the realization of $\mathcal{M}^*(S, S')$ - which we are keeping fixed - has cardinality lower than m . We want to prove that - for any fixed realization of T_1 and T_2 , the following holds (for the joint distributions!):

$$\begin{aligned} & \left\{ (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \right\}_{M \in \mathcal{M}^*(S, S'), i \notin I_M} \\ & \stackrel{(d)}{=} \left\{ (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \cdot r_i \right\}_{M \in \mathcal{M}^*(S, S'), i \notin I_M} \end{aligned}$$

*This follows by the fact that p_i and p'_i are independent and have the same distribution; therefore, applying the same measurable function on both and taking the difference renders a random variable with law which is symmetrical about zero.

[†]In other words, the equality is in the stronger sense of *conditional* distributions, when conditioning on (T_1, T_2) .

To see why this is indeed the case, consider any Borel set $B \subseteq [-1, 1]^{m \cdot (n-4/\varepsilon)}$. We have the following:

$$\begin{aligned} & \mathbb{P}_{\mathbf{X}, \mathbf{r}} \left((r_i \cdot (\text{profit}(M, p_i) - \text{profit}(M, p'_i)))_{i, M} \in B \right) \\ &= \frac{1}{2^n} \sum_{\bar{\mathbf{r}} \in \{-1, 1\}^n} \mathbb{P}_{\mathbf{X}} \left((\bar{r}_i \cdot (\text{profit}(M, p_i) - \text{profit}(M, p'_i)))_{i, M} \in B \right) \\ &= \mathbb{P}_{\mathbf{X}} \left((\text{profit}(M, p_i) - \text{profit}(M, p'_i))_{i, M} \in B \right), \end{aligned}$$

where the last equality follows by symmetry and independence of the random variables.

Given the equal distribution, we can apply any measurable function* to the two random vectors and still retain the same distribution. Therefore, for any fixed realizations of T_1 and T_2 we get:

$$\begin{aligned} & \max_{M \in \mathcal{M}^*(S, S')} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \\ & \stackrel{(d)}{=} \max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{i \notin I_M} r_i \cdot (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \right|. \end{aligned}$$

Note, once again, that the above equality holds for any fixed realizations of T_1 and T_2 , while the randomness is with respect to the \mathbf{X} and \mathbf{r} . If we now take the expectation with respect to the realizations of T_1 and T_2 , by the law of total probability, we get:

$$\begin{aligned} & \mathbb{E}_{S, S'} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \right] \\ &= \mathbb{E}_{S, S', \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{i \notin I_M} r_i \cdot (\text{profit}(M, p_i) - \text{profit}(M, p'_i)) \right| \right] \\ &\leq 2 \mathbb{E}_{S, S', \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{i \notin I_M} r_i \cdot \text{profit}(M, p_i) \right| \right], \end{aligned}$$

where the last inequality follows by triangle inequality and by the fact that S and S' are distributed independently and in the same way. \square

Lemma 7 (From Rademacher to Gaussian). *The following inequality holds:*

$$\begin{aligned} & \mathbb{E}_{S, S', \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot r_p \right| \right] \\ &\leq \sqrt{\frac{\pi}{2}} \mathbb{E}_{S, S', \mathbf{g}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot g_p \right| \right] \end{aligned}$$

where \mathbf{r} is a n -dimensional Rademacher vector and \mathbf{g} is a n -dimensional standard gaussian vector.

Proof. Fix any realization of the samples S' and S . They induce the set of mechanisms $\mathcal{M}^*(S, S')$, and the corresponding revenues. For any mechanism $M \in \mathcal{M}^*(S, S')$ and point $p \in S \setminus S(I_M)$, we have that the following two random variables share the same distribution:

$$\text{profit}(M, p) \cdot g_p \stackrel{(d)}{=} \text{profit}(M, p) \cdot |g_p| \cdot r_p, \quad (18)$$

*The composition of the averaging over $i \notin I_M$ for every mechanism and the maximum over the resulting vector.

where r_p are independent Rademacher random variables. Equation (18) jointly holds for any point $p \in S \setminus S(I_M)$, and any mechanism in $M \in \mathcal{M}^*(S, S')$. This follows by an analogous reasoning to the one in the proof of Lemma 6, and like in that case we have:

$$\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot g_p \right| \stackrel{(d)}{=} \max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot r_p \cdot |g_p| \right|. \quad (19)$$

We now have the final chain of inequalities:

$$\begin{aligned} & \mathbb{E}_{S, S', \mathbf{g}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot g_p \right| \right] \\ &= \mathbb{E}_{S, S', \mathbf{g}, \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I)} \text{profit}(M, p) \cdot r_p \cdot |g_p| \right| \right] \quad (\text{Equation (19)}) \\ &\geq \mathbb{E}_{S, S', \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I)} \text{profit}(M, p) \cdot r_p \cdot \mathbb{E}_{\mathbf{g}} [|g_p|] \right| \right] \quad (\text{Jensen}) \\ &\geq \sqrt{\frac{2}{\pi}} \mathbb{E}_{S, S', \mathbf{r}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I)} \text{profit}(M, p) \cdot r_p \right| \right]. \quad (\mathbb{E}_{\mathbf{g}} [|g_p|] = \sqrt{2/\pi}) \end{aligned}$$

The statement of the lemma follows by multiplying both terms of the above inequality by $\sqrt{\pi/2}$. \square

Lemma 8. *Let $g_1 \dots g_N$ be N arbitrary Gaussian random variables, with $\mu = 0$ and $\sigma_i^2 \leq \sigma^2$. Then*

$$\mathbb{E} \left[\max_{i \in [N]} |g_i| \right] \leq 2\sqrt{\sigma^2 \log N}.$$

Proof. This is a standard result, see e.g., Lemma 2.3.4. of Talagrand [2014]. We provide here a proof for completeness. For any positive t , we have the following chain of inequalities:

$$\begin{aligned} \exp \left(t \cdot \mathbb{E} \left[\max_{i \in [N]} |g_i| \right] \right) &\leq \mathbb{E} \left[\exp \left(\max_{i \in [N]} t \cdot |g_i| \right) \right] \quad (\text{By Jensen's Inequality}) \\ &= \mathbb{E} \left[\max_{i \in [N]} e^{t \cdot |g_i|} \right] \quad (\text{Monotonicity of the exponential}) \\ &\leq \mathbb{E} \left[\max_{i \in [N]} (e^{t \cdot g_i} + e^{-t \cdot g_i}) \right] \quad (\text{Because } e^x \geq 0 \text{ for all } x) \\ &\leq \sum_{i=1}^N \mathbb{E} [e^{t \cdot g_i} + e^{-t \cdot g_i}] \\ &= 2 \sum_{i=1}^N \mathbb{E} [e^{t \cdot g_i}] \quad (g_i \text{ and } -g_i \text{ identically distributed}) \\ &\leq 2N e^{\frac{t^2 \sigma^2}{2}}. \quad (\text{Gaussian moment generating function}) \end{aligned}$$

At this point, it is enough to apply the logarithm to both sides of the above inequality and upper bounding $\log(2N)$ with $\log N^2$:

$$\mathbb{E} \left[\max_{i \in [N]} |g_i| \right] \leq \frac{t\sigma^2}{2} + \frac{2}{t} \log(N) = 2\sqrt{\sigma^2 \log N},$$

where the last equality holds by setting $t = 2\sqrt{\log N/\sigma^2}$ \square

Lemma 9. *Consider any realization of the samples S and S' . We have the following inequality:*

$$\begin{aligned} & \mathbb{E}_{\mathbf{g}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot g_p \right| \right] - 38\varepsilon \\ & \leq \sum_{h=0}^{H-1} \mathbb{E}_{\mathbf{g}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} (\text{profit}(\phi_{h+1}(M), p) - \text{profit}(\phi_h(M), p)) \cdot g_p \right| \right] \end{aligned}$$

Proof. Note, the revenue of any mechanism $M \in \mathcal{M}_{S(I \setminus J), S'(J)}^\varepsilon$ on any point $p \in S \setminus S(I)$ can be written via the following telescopic sum:

$$\text{profit}(M, p) = \text{profit}(\phi_0(M), p) + \sum_{h=0}^{H-1} (\text{profit}(\phi_{h+1}(M), p) - \text{profit}(\phi_h(M), p)).$$

In fact, since $M \in \mathcal{M}_{S(I \setminus J), S'(J)}^\varepsilon$ it holds that $\phi_H(M) = M$. We study separately the term that depends on $\phi_0(M)$ and the summation. In particular, we only need to address the latter term, as the former one already appears in the statement (simply by taking the summation out of the absolute value and the expected value). We want to prove the following inequality, which concludes the proof of the lemma.

$$\mathbb{E}_{\mathbf{g}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(\phi_0(M), p) \cdot g_p \right| \right] \leq 38\varepsilon \quad (20)$$

The mechanisms of the form $\phi_0(M)$ have a precise structure: they are of the form $[0, v_1] \times [v_2, 1]$ for some node (v_1, v_2) in one of the grids that define $\mathcal{M}^*(S, S')$. In particular, there are at most $3n^2$ such points, as $n \geq 8/\varepsilon$. This means that when approximating through $\phi_0(M)$, given the tuple (S, S') , there are at most $3n^2$ distinct mechanism. We have the following inequality for any $M \in \mathcal{M}^*(S, S')$ (recall, $n' = n - 4/\varepsilon$)

$$\begin{aligned} & \max_{M \in \mathcal{M}^*(S, S')} \frac{1}{n'} \left| \sum_{p \in S \setminus S(I_M)} \text{profit}(\phi_0(M), p) \cdot g_p \pm \sum_{p \in S(I_M)} \text{profit}(\phi_0(M), p) \cdot g_p \right| \\ & \leq \underbrace{\max_{M \in \mathcal{M}^*(S, S')} \frac{1}{n'} \left| \sum_{p \in S} \text{profit}(\phi_0(M), p) \cdot g_p \right|}_{(A)} + \underbrace{\max_{M \in \mathcal{M}^*(S, S')} \frac{1}{n'} \left| \sum_{p \in S(I_M)} \text{profit}(\phi_0(M), p) \cdot g_p \right|}_{(B)} \end{aligned}$$

In term (A) we are taking the maximum over the absolute values of at most $3n^2$ centered gaussian terms, each with variance $n/(n')^2 \leq 4/n$. Using Lemma 8, we get that (A) is bounded in expectation by $8\sqrt{\log n/n}$. For what concerns (B), by Proposition 6, we are taking the maximum over the absolute values of at most $e^{10 \log n/\varepsilon} \cdot 2^{4/\varepsilon} \leq e^{14 \log n/\varepsilon}$ gaussian terms, each with variance at most $\frac{16}{n^2 \varepsilon}$. Again by Lemma 8, we have that (B) is bounded in expectation by $30\sqrt{\frac{\log n}{\varepsilon^2 n^2}} \leq 30\sqrt{\frac{\log n}{n}}$, under $n \geq 1/\varepsilon^2$.

By combining these two bounds, we get $38\sqrt{(\log n)/n}$ as an upper bound for the sum of (A) and (B). Our assumption implies $n \geq 4/\varepsilon^2 \log(1/\varepsilon)$, so we conclude* by bounding $\sqrt{(\log n)/n}$ with ε . \square

*Notice that $\log n = \varepsilon^2 n$ has no elementary solution in n , therefore the bound follows by upper bounding terms involving the Lambert W function.

We now define the clean event \mathcal{E} as follows: for all mechanisms $M \in \mathcal{M}^*(S, S')$, we require that the following inequality is verified for all $h = 0, \dots, H - 1$:

$$\text{profit}(\phi_h(M, S \setminus S(I_M))) - \text{profit}(\phi_{h+1}(M, S \setminus S(I_M))) < 31 \cdot 2^{-h}. \quad (21)$$

Lemma 10 (Clean Event). *The clean event \mathcal{E} happens with probability at least $(1 - \frac{1}{n \log 1/\epsilon})$.*

Proof. Fix any mechanism $M \in \mathcal{M}^*(S, S')$. This means that M is a candidate in some family $\mathcal{M}_{S(I \setminus J) \cup S'(J)}^\epsilon$, recalling Definition 10. We denote with $\mathcal{E}_{M,h}$ the following event:

$$\mathcal{E}_{M,h} = \left\{ \text{profit}(\phi_h(M, S \setminus S(I_M))) - \text{profit}(\phi_{h+1}(M, S \setminus S(I_M))) \geq 31 \cdot 2^{-h} \right\}.$$

Similarly to what is done in Theorem 2, we introduce the random variables X_p and study their sum X for $p \notin S(I_M)$:

$$X_p = \frac{1}{4} [\text{profit}(\phi_h(M), p) - \text{profit}(\phi_{h+1}(M), p) + 2/2^h].$$

We know that $X_p \in [0, 1]$, and, since we assume that M is a candidate, we are under the assumption* that $\mathbb{E}[X] < 3/2 \cdot 2^{-h} \cdot n'$ (as $\mathbb{E}[\text{profit}(\phi_h(M)) - \text{profit}(\phi_{h+1}(M))] \leq 4/2^h$, and $n' = |S \setminus S(I)|$). We then have the following chain of inequalities, under our assumptions on n :

$$\begin{aligned} \mathbb{P}(\mathcal{E}_{M,h}) &= \mathbb{P}\left(X \geq 33 \cdot 2^{-h} \cdot n'\right) && \text{(By def. of } X) \\ &\leq 2^{-33 \cdot 2^{-h} \cdot n'} && \text{(By Theorem 12)} \\ &\leq 2^{-\frac{33}{2} \epsilon n} && \text{(As } 2^{-h} \geq \epsilon \text{ and } n' \geq n/2.) \end{aligned}$$

Note, we applied the second inequality in Theorem 12, using the fact that

$$2e\mathbb{E}[X] \leq 3e \cdot 2^{-h} \cdot n' \leq 33n' \cdot 2^{-h}.$$

The only missing ingredient is now to take the union bound over all possible levels and mechanisms in $\mathcal{M}^*(S, S')$. The levels are at most $2 \log 1/\epsilon$ and by Proposition 6 the number of mechanisms is at most $e^{10 \log n/\epsilon}$, multiplied by the possible way of choosing $J \subset I$, i.e. $2^{4/\epsilon}$, the cardinality of the power set of J . So the probability of failure is at most:

$$2 \cdot \log 1/\epsilon \cdot 2^{-\frac{33}{2} \epsilon n + \frac{4}{\epsilon} \frac{10 \log n}{\epsilon}} \leq \frac{1}{n \log 1/\epsilon},$$

as long as $n \geq \frac{8 \log 1/\epsilon}{\epsilon^2}$. □

Lemma 11. *Consider any realization of the samples S and S' that is compatible with the clean event \mathcal{E} . Then for any $h = 0, \dots, H - 1$, we have the following inequality:*

$$\mathbb{E}_g \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I)} (\text{profit}(\phi_{h+1}(M), p) - \text{profit}(\phi_h(M), p)) \cdot g_p \right| \right] \leq 116 \frac{\epsilon}{\log 1/\epsilon}$$

Proof. We are under the clean event, therefore we know that, for any $M \in \mathcal{M}^*(S, S')$ we have that Equation (21) is verified. Fix any such M , it belongs to some $\mathcal{M}_{S(I \setminus J), S'(J)}^\epsilon$ for some I and $J \subseteq I$. We want to apply Proposition 5 to the empirical distribution induced by the points in $S \setminus S(I)$.

*This is the opposite situation with respect to the argument in the proof of Theorem 2.

By Equation (21), we can then conclude that, for all $h = 0, \dots, H - 1$ the following inequality is verified:

$$\frac{1}{n'} \sum_{p \in S \setminus S(I)} (\text{profit}(\phi_{h+1}(M)) - \text{profit}(\phi_h(M)))^2 \leq 136 \cdot 2^{-h}. \quad (22)$$

Therefore, if we consider the random variables that appear in the absolute value in the statement of this theorem, they are all centered gaussian variables with variance bounded as follows:

$$\begin{aligned} & \mathbb{V} \left[\frac{1}{n'} \sum_{p \in S \setminus S(I)} (\text{profit}(\phi_{h+1}(M), p) - \text{profit}(\phi_h(M), p)) \cdot g_p \right] \\ &= \frac{1}{(n')^2} \sum_{p \in S \setminus S(I)} (\text{profit}(\phi_{h+1}(M), p) - \text{profit}(\phi_h(M), p))^2 \mathbb{V}[g_p] \leq \frac{136}{n'} \cdot 2^{-h}, \end{aligned}$$

where we used the property of the variance, the fact that the g_p are i.i.d. standard gaussians, and the bound provided by Equation (22).

The last ingredient we need before applying the result in Lemma 8 is to provide an upper bound on the number of mechanisms that appear in the max. The points in S and S' are fixed, moreover, each mechanism of the form $\phi_h(M)$ for $M \in \mathcal{M}^*(S, S')$ is supported on the grid $G_{S \cup S'}^{2^{-h}}$, recalling Definition 9, and has at most 2^{h+1} ‘‘corners’’, by the same potential argument used in the second part of the proof of Lemma 4. Therefore, using monotonicity of the allocation regions, the cardinality N of the set onto which the max is taken over is (note, each term contains both ϕ_h and ϕ_{h+1}):

$$N \leq \binom{2n + 2/\varepsilon}{2^{h+1}} \cdot \binom{2n + 2/\varepsilon}{2^{h+2}} \leq \binom{4n}{2^{h+1}} \cdot \binom{4n}{2^{h+2}} \leq (64^2 n^6 e^6)^{2^h}.$$

Note, in the above inequality we used that $\binom{n}{k} \leq ((en)/k)^k$. We can then conclude via Lemma 8 to claim that the left-hand side term in the statement is upper-bounded by:

$$\begin{aligned} 2\sqrt{\sigma^2 \log N} &\leq 2\sqrt{\frac{140}{n/2} \cdot 2^{-h} \cdot \log((4en)^{6 \cdot 2^h})} \\ &\leq 2\sqrt{\frac{140}{n/2} \cdot 2^{-h} \cdot 6 \cdot 2^h \log(11n)} \leq 116\sqrt{\frac{\log n}{n}} \end{aligned}$$

The statement follows since $\sqrt{\frac{\log n}{n}} \leq \frac{\varepsilon}{\log 1/\varepsilon}$ when we assume that $n \geq 32/\varepsilon^2 \log^3(1/\varepsilon)$. \square

Lemma 12. *Consider any realization of the samples S that is compatible with the clean event \mathcal{E} . We have the following inequality:*

$$\mathbb{E}_{\mathbf{g}} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{|S \setminus S(I_M)|} \sum_{p \in S \setminus S(I)} \text{profit}(M, p) \cdot g_p \right| \right] \leq 270\varepsilon$$

Proof. It follows by combining Lemma 11 with Lemma 9, and noting that $H = \log_2 1/\varepsilon \leq 2 \log 1/\varepsilon$. \square

Proof of Theorem 3. The result follows by combining the above inequalities:

$$\begin{aligned}
& \mathbb{E}_S \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S) - \mathbb{E}[\text{profit}(M)] \right| \right] - \varepsilon \\
& \leq \mathbb{E}_{S, S'} \left[\max_{M \in \mathcal{M}^*(S)} \left| \text{profit}(M, S \setminus S(I_M)) - \text{profit}(M, S' \setminus S'(I_M)) \right| \right] \quad (\text{By Lemma 5}) \\
& \leq 2\mathbb{E}_{S, S', r} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot r_p \right| \right] \quad (\text{Lemma 6}) \\
& \leq \sqrt{2\pi} \mathbb{E}_{S, S', g} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot g_p \right| \right] \quad (\text{Lemma 7}) \\
& \leq \sqrt{2\pi} \left(\mathbb{E}_{S, S', g} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} \text{profit}(M, p) \cdot g_p \right| \middle| \mathcal{E} \right] \mathbb{P}(\mathcal{E}) + 38\varepsilon \right. \\
& \quad \left. + \sum_{h=0}^{H-1} \mathbb{E}_{S, S', g} \left[\max_{M \in \mathcal{M}^*(S, S')} \left| \frac{1}{n'} \sum_{p \in S \setminus S(I_M)} (\text{profit}(\phi_{h+1}(M), p) - \text{profit}(\phi_h(M), p)) \cdot g_p \right| \middle| \mathcal{E}^c \right] \mathbb{P}(\mathcal{E}^c) \right) \\
& \hspace{15em} (\text{By Lemma 9 and Law of Total Probability}) \\
& \leq \sqrt{2\pi} \left(308\varepsilon + \frac{1}{n \log(1/\varepsilon)} \cdot 2 \log(1/\varepsilon) \cdot \sqrt{14 \log n} \right) \quad (\text{By Lemmas 8, 10 and 12}) \\
& \leq 793\varepsilon.
\end{aligned}$$

In the second to last inequality, we can bound the first term by Lemma 12, while the second can be bounded with Lemma 10 and by noticing that, when conditioning on (S, S') , we are taking the maximum over at most $\exp(10 \log n/\varepsilon) \cdot 2^{4/\varepsilon} \leq \exp(14 \log n/\varepsilon)$ centered gaussian terms with variance at most $1/n' \leq 2/n$. By Lemma 8, we have that the inner expectation w.r.t. the gaussian terms is bounded by $\sqrt{14 \log n}$. \square

5 Lower Bounds

In this section, we prove that the positive result in Section 4 is, essentially, tight. In particular, in Section 5.1 we prove that no learning algorithm can achieve better than \sqrt{T} regret in the stochastic setting. Finally, in Section 5.2 we prove that no learning can happen in the adversarial setting.

5.1 Stochastic Lower Bound

We derive the \sqrt{T} lower bound in the stochastic setting by embedding into our framework the hard instance for prediction with two experts. In particular, we consider two stochastic instances that are statistically close but admit two well-separated optimal mechanisms.

Consider the probability measure \mathbb{P}^0 under which V is drawn uniformly at random in the set $\{(0, 1), (1/4, 3/4)\}$. We use \mathbb{P}^0 as “baseline distribution” to compare the following two measures \mathbb{P}^1 and \mathbb{P}^2 . Let $\varepsilon > 0$ a precision parameter we set later, under measure \mathbb{P}^1 , the random valuation V is drawn according to the following distribution:

$$\begin{cases} (0, 1) & \text{with probability } (1+\varepsilon)/2 \\ (1/4, 3/4) & \text{with probability } (1-\varepsilon)/2. \end{cases}$$

Under \mathbb{P}^2 , the distribution is the same, but with flipped signs, so that $\mathbb{P}^2(V = (0, 1)) = (1-\varepsilon)/2$ and $\mathbb{P}^2(V = (1/4, 3/4)) = (1+\varepsilon)/2$.

Remark 3. *To avoid overloading the notation, we do not consider “different” random variables but vary the underlying probability measure. We denote with \mathbb{E}^i , for $i = 0, 1, 2$ the corresponding expectation.*

Under these three probability measures, there is a clear way to partition the mechanism space, according to whether the allocation region contains $(1/4, 3/4)$ or not. To this end, denote with M_1 the mechanism that allocates only in the point $(0, 1)$, and with M_2 the one that allocates in the rectangle $[0, 1/4] \times [3/4, 1]$. Clearly, M_1 or M_2 dominate any mechanism, as formalized by the following lemma.

Lemma 13 (Domination). *Consider any mechanism M with allocation region A , then one of the following is verified:*

- *If $(1/4, 3/4) \notin A$, then $\mathbb{E}^i[\text{profit}(M_1)] \geq \mathbb{E}^i[\text{profit}(M)]$ for $i = 0, 1, 2$*
- *If $(1/4, 3/4) \in A$, then $\mathbb{E}^i[\text{profit}(M_2)] \geq \mathbb{E}^i[\text{profit}(M)]$ for $i = 0, 1, 2$*

Proof. Consider the first case: $(1/4, 3/4) \notin A$. If A is empty, then clearly M_1 dominates M , as its expected profit is approximately $1/2$ under \mathbb{P}^i for $i = 0, 1, 2$. If A is not empty, then it contains $(0, 1)$, therefore the profit it extracts from it is at most that extracted by M^1 , whose allocation region is contained in A (Observation 1).

We move our attention to any mechanism M whose allocation region A contains $(1/4, 3/4)$. The smallest monotone allocation region that contains $(1/4, 3/4)$ is the rectangle $[0, 1/4] \times [3/4, 1]$, i.e., the allocation region of M_2 . Therefore, the profit extracted by M_2 is at least that extracted by M (again, by Observation 1). \square

We have all the ingredients to prove our first lower bound.

Theorem 5 (Lower Bound in the Stochastic Case). *Consider the profit maximization problem in bilateral trade against the stochastic i.i.d. adversary. Then any algorithm \mathcal{A} suffers at least $1/16\sqrt{T}$ regret.*

Proof. We prove that any learning algorithm \mathcal{A} suffers $\Omega(\sqrt{T})$ regret rate against at least one between \mathbb{P}^1 and \mathbb{P}^2 . For $i = 0, 1, 2$, denote with \mathbb{P}_V^i the push-forward measure on $[0, 1]^2$ (equipped with the Borel σ -algebra) induced by the random variable V when the underlying probability measure is \mathbb{P}^i . We start formalizing the intuition that the \mathbb{P}^1 and \mathbb{P}^2 are “close” to \mathbb{P}^0 in a statistical sense by studying their Kullback-Leibler divergence KL.

Claim 2. *If $\varepsilon \in (0, 1/2)$, it holds that $\max\{\mathcal{D}_{\text{KL}}(\mathbb{P}_V^0, \mathbb{P}_V^1), \mathcal{D}_{\text{KL}}(\mathbb{P}_V^0, \mathbb{P}_V^2)\} \leq 2\varepsilon^2$.*

Proof of the Claim. We bound the KL divergence between \mathbb{P}^1 and \mathbb{P}^0 ; the other term is equal.

$$\begin{aligned} \mathcal{D}_{\text{KL}}(\mathbb{P}_V^0, \mathbb{P}_V^1) &= \frac{1}{2} \log\left(\frac{1}{1+\varepsilon}\right) + \frac{1}{2} \log\left(\frac{1}{1-\varepsilon}\right) && \text{(By definition)} \\ &= -\frac{1}{2} \log(1-\varepsilon^2) \leq 2\varepsilon^2, \end{aligned}$$

where the last inequality holds for any $\varepsilon \in (0, 1/2)$. \square

Under the stochastic instance represented by \mathbb{P}^0 , both M_1 and M_2 are optimal (by Lemma 13), and achieve expected profit: $\mathbb{E}^0[\text{profit}(M_1)] = \mathbb{E}^0[\text{profit}(M_2)] = \frac{1}{2}$. When the instance is induced by \mathbb{P}^i , the domination result in Lemma 13 tells us that the optimal mechanism is M_i , with the performance of M_j for $j \neq i$ that is $\Theta(\varepsilon)$ far:

$$\begin{cases} \mathbb{E}^1[\text{profit}(M_1)] = \frac{1+\varepsilon}{2}, & \mathbb{E}^1[\text{profit}(M_2)] = \frac{1}{2} \\ \mathbb{E}^2[\text{profit}(M_1)] = \frac{1-\varepsilon}{2}, & \mathbb{E}^2[\text{profit}(M_2)] = \frac{1}{2} \end{cases} \quad (23)$$

Fix now any deterministic learning algorithm \mathcal{A} , we prove that its expected regret with respect to the randomized instance that is either \mathbb{P}^1 or \mathbb{P}^2 with equal probability is $\Omega(\sqrt{T})$. Note, the randomized instance is constructed by first choosing \mathbb{P}^i uniformly at random and then sampling from it, *not* by changing distribution at each time step. We remark that proving this result (restricting to deterministic algorithms) is enough to prove the theorem, as the randomized instance we use is independent of the algorithm used by the learner.

We denote with N_i the number of times that algorithm \mathcal{A} plays the mechanism M_i , for $i = 1, 2$. Note, there is no point for \mathcal{A} to play any other mechanism, by Lemma 13. Denote with V_1, V_2, \dots the i.i.d. samples of V , and denote with $\mathbb{P}_{V_1, \dots, V_t}^i$ the push-forward to $[0, 1]^{2t}$. We have the following chain of inequalities:

$$\begin{aligned} |\mathbb{E}^i[N_i] - \mathbb{E}^0[N_i]| &\leq \sum_{t=1}^T |\mathbb{P}^i(M^t = M_i) - \mathbb{P}^0(M^t = M_i)| \\ &\leq \sum_{t=2}^T \|\mathbb{P}_{V_1, \dots, V_t}^i - \mathbb{P}_{V_1, \dots, V_t}^0\|_{\text{TV}} && \text{(By definition of total variation)} \\ &\leq \sum_{t=2}^T \sqrt{\frac{t}{2} \mathcal{D}_{\text{KL}}(\mathbb{P}_{V_t}^i, \mathbb{P}_{V_t}^0)} && \text{(By Pinsker Inequality and i.i.d. design)} \\ &\leq \varepsilon T \sqrt{T}, \end{aligned} \quad (24)$$

where the last inequality follows by Claim 2.

We are ready to bound the expected performance of \mathcal{A} against the randomized instance constructed by picking either \mathbb{P}^1 or \mathbb{P}^2 with equal probability. Note, under \mathbb{P}^i , the regret of \mathcal{A} is induced by the N_j time steps in which it plays mechanism M_j for $j \neq i$, multiplied by the corresponding instantaneous regret of $\varepsilon/2$ (by Equation (23)):

$$\begin{aligned} \mathbb{E}[R_T(\mathcal{A})] &= \frac{\varepsilon}{4} (\mathbb{E}^1[N_2] + \mathbb{E}^2[N_1]) \\ &\geq \frac{\varepsilon}{4} (\mathbb{E}^0[N_1 + N_2] - 2\varepsilon T \sqrt{T}) && \text{(By Equation (24))} \\ &= \frac{\varepsilon}{2} T (1 - 2\varepsilon \sqrt{T}) && \text{(As } N_1 + N_2 = T) \\ &= \frac{1}{16} \sqrt{T}. && \text{(By setting } \varepsilon = 1/4\sqrt{T}) \end{aligned}$$

This concludes the proof. \square

5.2 Adversarial Lower Bound

This section is devoted to our adversarial lower bound. We not only prove that it is impossible to achieve sublinear regret in the adversarial case, but argue that no sublinear α -regret is possible, for any constant $\alpha \in [1, 3/2)$.*

*For $\alpha = 1$ the definition of α -regret (as in the statement of Theorem 6) collapses to that of regret.

Theorem 6 (Lower Bound in the Adversarial Case). *Consider the profit maximization problem in bilateral trade in the adversarial setting. For any constant $\varepsilon \in (0, 1]$ and any learning algorithm, there exists an adversarial sequence such that the following $(3/2 - \varepsilon)$ -regret bound holds:*

$$\sup_{M \in \mathcal{M}} \sum_{t=1}^T \text{profit}_t(M, v^t) - \left(\frac{3}{2} - \varepsilon\right) \cdot \sum_{t=1}^T \mathbb{E} [\text{profit}_t(M_t, v^t)] \geq \frac{\varepsilon}{4} T.$$

Proof. We prove this lower bound via Yao's minimax principle: we construct a randomized (adversarial) instance against which any deterministic algorithm suffers expected linear $(3/2 - \varepsilon)$ regret. Our hard instance is based on an auxiliary sequence constructed in Theorem 4 of Aggarwal et al. [2024].

Let δ be a small auxiliary parameter we set later. We construct our randomized instance while maintaining two auxiliary sequences, a_t and b_t , that are initialized as $a_1 = \delta/3$ and $b_1 = 2\delta/3$; we also denote with Δ_t the difference $b_t - a_t$, which is maintained at $\delta \cdot 3^{-t}$ throughout. The sequence of valuations (v_s^t, v_b^t) is generated iteratively as follows:

- with probability $1/2$, the valuations are $(v_s^t, v_b^t) = (b_t, 1/2)$, and the auxiliary sequence is updated as follows: $(a_{t+1}, b_{t+1}) = (b_t + \Delta_{t+1}, b_t + 2\Delta_{t+1})$
- with probability $1/2$, the valuations are $(v_s^t, v_b^t) = (a_t, 1)$, and the next pair in the auxiliary sequence is $(a_{t+1}, b_{t+1}) = (a_t - 2\Delta_{t+1}, a_t - \Delta_{t+1})$.

Arguing similarly to Aggarwal et al. [2024], it is possible to verify that this random sequence is well defined, as all realizations of the pairs (v_s^t, v_b^t) belong to the $[0, 1]^2$ square. Crucially, *for any realization*, there exists a threshold τ that separates the valuations of the form $(b_t, 1/2)$ from those of the form $(a_t, 1)$. Formally, denote with L (for left) the time steps characterized by a $(a_t, 1)$ valuation, and with R (for right) the ones corresponding to $(b_t, 1/2)$ valuations, the following proposition holds (this is a restatement of Proposition 4 of Aggarwal et al. [2024], we provide a more direct proof in Section E for completeness).

Proposition 7. *For any realization of the valuations and auxiliary sequence, we have:*

- (i) $a_t \in [0, \delta]$ and $b_t \in [0, \delta]$, for all $t = 1, \dots, T$.
- (ii) there exists a value $\tau \in (0, \delta)$ such that $b_t < \tau$ for all $t \in R$, while $a_t > \tau$ for all $t \in L$.

The threshold value τ can be used to design a mechanism M^* that allocates in $[0, \tau] \times [1/2, 1]$ united with the $(0, 1)$ - $(\delta, 1)$ segment. M^* accepts all the bids of the form $(x, 1/2)$ for $x \leq \tau$ (with profit at least $1/2 - \delta$), and all the bids of the form $(x, 1)$ with $x \in [\tau, \delta]$ (with profit at least $1 - \delta$). In each iteration, the expected profit that M^* extracts is thus at least

$$\mathbb{E} [\text{profit}_t(M^*)] \geq \frac{1}{2} (1 - \delta) + \frac{1}{2} (\frac{1}{2} - \delta) = \frac{1}{2} (\frac{3}{2} - 2\delta). \quad (25)$$

Consider now any deterministic algorithm and its expected revenue at time t . Fix the history of the sequence up to time t (excluded); there are two options for the next valuations, with equal probability: either $(b_t, 1/2)$ or $(a_t, 1)$. Since the latter dominates the former, any mechanism M^t that the learning algorithm posts can do one of three things: it never trades, it makes the trade happen for both possible valuations $(b_t, 1/2)$ and $(a_t, 1)$ (with an expected profit of at most $1/2$), or only for $(b_t, 1/2)$ (with an expected profit of at most $1/2$). All in all, the expected revenue is at most

$$\mathbb{E} [\text{profit}_t(M^t)] \leq \frac{1}{2}. \quad (26)$$

Combining Equation (25) and Equation (26), we see that the expected $3/2 - \varepsilon$ instantaneous regret at the generic time step t is at least:

$$\mathbb{E}[\text{profit}_t(M^*)] - \left(\frac{3}{2} - \varepsilon\right) \mathbb{E}[\text{profit}_t(M^t)] \geq \frac{1}{2} \left(\frac{3}{2} - 2\delta\right) - \left(\frac{3}{2} - \varepsilon\right) \frac{1}{2} = \frac{\varepsilon}{4},$$

where in the last equality we set $\delta = \varepsilon/4$. Summing up over all the time steps t , yields the desired lower bound on the $(3/2 - \varepsilon)$ -regret. \square

6 The Joint Ads Problem

In this section, we showcase the applicability of our chaining analysis from Section 4 by providing an upper bound of \sqrt{T} on the regret for the joint ads problem [Aggarwal et al., 2024], in the i.i.d. stochastic setting. This result, together with the $\Omega(\sqrt{T})$ lower bound from Aggarwal et al. [2024] provides a tight (up to poly-logarithmic factors) characterization of the minimax regret rate for the problem. For completeness, we recall that an $\Omega(T)$ lower bound in the adversarial setting is provided in Aggarwal et al. [2024].

6.1 The Joint Ads Problem

In the Joint Ads problem, two buyers participate in an auction for a non-excludable good (e.g., a public good or a shared ad slot that is either allocated to *both* agents or to none). The goal of the mechanism designer is to maximize the revenue extracted from the trade. We now briefly introduce the problem; for further details we refer to Aggarwal et al. [2024].

The Learning Protocol. Formally, at each time step $t = 1, 2, \dots, T$, a new pair of agents arrives, characterized by private valuations (v_1^t, v_2^t) in the $[0, 1]^2$ square. Independently, the learner proposes a mechanism M^t to the agents, who then declare bids (b_1^t, b_2^t) . A mechanism M^t is composed by a non-excludable allocation rule $x^t : [0, 1]^2 \rightarrow \{0, 1\}$ (either both agents win the auction or both lose) and two payment functions $p_1^t, p_2^t : [0, 1]^2 \rightarrow [0, 1]$ that map bids to allocation and payment, respectively. The revenue of mechanism M^t when the bids are b_1^t and b_2^t is defined as

$$\text{rev}(M^t, b_1^t, b_2^t) = x^t(b_1^t, b_2^t) \cdot [p_1^t(b_1^t, b_2^t) + p_2^t(b_1^t, b_2^t)].$$

Structure of Incentive Compatible Mechanisms. As in the bilateral trade problem, the agents behave strategically, meaning that they strive to maximize their own quasi-linear utility $u_i^t(b_1^t, b_2^t) = v_i^t \cdot x^t(b_1^t, b_2^t) - p_i^t(b_1^t, b_2^t)$ for $i \in \{1, 2\}$. Note that, differently from the bilateral trade problem, here the two agents are symmetric, as they are both interested in purchasing a good. We require the mechanisms M^t to be DSIC and IR. The family of mechanisms \mathcal{M} that respect these two properties can be easily characterized because this model falls within the framework of single parameter auctions [Myerson, 1981], as each agent is characterized by a single value representing its value for the good for sale. This tells us that (i) an allocation rule is implementable if and only if it is monotone and (ii) for each such allocation rule there exists a unique payment rule that completes it to a DSIC and IR mechanism.* For this reason, we only talk about valuations in the rest of the section, as the agents are incentivized to behave truthfully.

*As in bilateral trade, payments are unique only up to constant shifts. Since our goal is to maximize revenue, we adopt the maximum payments to achieve DSIC and IR.




Figure 5: Visualization of the allocation region of a joint ad mechanism.

Differences with Bilateral Trade. There are two main differences with bilateral trade. First, the two agents are symmetric; this induces a different notion of monotonicity: valuations (v_1, v_2) dominate (v'_1, v'_2) if $v_i \geq v'_i$ for $i = 1, 2$. Stated differently, if at least one of the two agents raises its bid, then the likelihood of having a trade does not decrease. Second, in the joint ads problem, the objective of the learner, i.e., the revenue, lies in $[0, 2]$, as opposed to the profit in bilateral trade, which lives in $[-1, 1]$. This is a *crucial* distinction: while in both problems the event “no trade” induces zero gain to the learner, in bilateral trade, there is a notion of “bad trade” that induces *negative* utility for the learner. On the contrary, in the joint ads problem, every trade guarantees at least non-negative payments *to* the learner.

6.2 The Stochastic Algorithm

In this section, we argue that the algorithm SIMPLIFY-THE-BEST-MECHANISM easily adapts to this new problem and achieves the desired $\tilde{O}(\sqrt{T})$ regret bound in the stochastic setting. From the technical point of view, the main difference with respect to the bilateral trade problem is that the revenue function lies in the $[0, 2]$ interval, so that (i) observing “no trade” is the worst occurrence and (ii) the gap between the best revenue and no-trade is 2 instead of 1. Point (i) implies that we do not want to focus our attention only on mechanisms living in the bottom left triangle (note, with the notion of monotonicity induced by this problem, the “low-value” corner of the square is the one close to the origin), so that we do not restrict the mechanism space as in Lemma 1.

For the joint ads problem, SIMPLIFY-THE-BEST-MECHANISM collapses to the algorithm used in Aggarwal et al. [2024] (AUGMENT-THE-BEST-MECHANISM). There is, however, one big difference: our chaining-based analysis allows us to prove a greatly improved sample complexity of $1/\varepsilon^2$, as opposed to the sample complexity of $1/\varepsilon^4$ implicit in Aggarwal et al. [2024]. This difference is observable in the different choice of precision parameter: we set $\varepsilon_t = 200\sqrt{\log^3 T/t-1}$, as opposed to approximately $t^{-1/4}$.

Theorem 7. *Consider revenue-maximization in the joint ad problem against the stochastic i.i.d. adversary. The regret of AUGMENT-THE-BEST-MECHANISM is at most $C\sqrt{T \log^3 T}$, for some constant C .*

Proof. The proof of this result follows verbatim from the proof of Theorem 4 by replacing the profit

function with `rev`, and by considering the different notion of monotonicity (this implies, for instance, that the approximating sequences approximate mechanisms from the bottom-left direction, and not from the bottom-right, see for instance Figure 5). Note, we do not need to worry about negative revenue (this simplifies some of the arguments), while we get an extra multiplicative 2 in the regret rate as `rev` $\in [0, 2]$, while `|profit|` $\in [0, 1]$. \square

References

- Gagan Aggarwal, Ashwinkumar Badanidiyuru, Paul Dütting, and Federico Fusco. Selling joint ads: A regret minimization perspective. In *EC*, pages 164–194. ACM, 2024.
- Noga Alon, Nicolò Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *COLT*, volume 40, pages 23–35, 2015.
- Jean-Yves Audibert and Olivier Bousquet. Combining pac-bayesian and generic chaining bounds. *J. Mach. Learn. Res.*, 8:863–889, 2007.
- Yossi Azar, Amos Fiat, and Federico Fusco. An α -regret analysis of adversarial bilateral trade. *Artif. Intell.*, 337:104231, 2024.
- Moshe Babaioff, Kira Goldner, and Yannai A. Gonczarowski. Bulow-klemperer-style results for welfare maximization in two-sided markets. In *SODA*, pages 2452–2471. ACM-SIAM, 2020.
- Gábor Bartók, Dean P. Foster, Dávid Pál, Alexander Rakhlin, and Csaba Szepesvári. Partial monitoring - classification, regret bounds, and algorithms. *Math. Oper. Res.*, 39(4):967–997, 2014.
- Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Federico Fusco. No-regret learning in bilateral trade via global budget balance. In *STOC*, pages 247–258. ACM, 2024.
- Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin. Smoothed online learning is as easy as statistical learning. In *COLT*, volume 178, pages 1716–1786, 2022.
- Avrim Blum, Vijay Kumar, Atri Rudra, and Felix Wu. Online learning in online auctions. In *SODA 2003*, pages 202–204. ACM-SIAM, 2003.
- Liad Blumrosen and Shahar Dobzinski. (Almost) efficient mechanisms for bilateral trading. *Games Econ. Behav.*, 130:369–383, 2021.
- Jean Bourgain, Sjoerd Dirksen, and Jelani Nelson. Toward a unified theory of sparse dimensionality reduction in euclidean space. In *STOC*, pages 499–508. ACM, 2015.
- Vladimir Braverman, Stephen R. Chestnut, Nikita Ivkin, Jelani Nelson, Zhengyu Wang, and David P. Woodruff. Bptree: An ℓ_2 heavy hitters algorithm using constant memory. In *PODS*, pages 361–376. ACM, 2017.
- Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. Approximating gains from trade in two-sided markets via simple mechanisms. In *EC*, pages 589–590. ACM, 2017.
- Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *FOCS*, pages 516–527. IEEE, 2017.
- Yang Cai and Jinzhao Wu. On the optimal fixed-price mechanism in bilateral trade. In *STOC*, pages 737–750. ACM, 2023.

- Yang Cai, Kira Goldner, Steven Ma, and Mingfei Zhao. On multi-dimensional gains from trade maximization. In *SODA 2021*, pages 1079–1098. ACM-SIAM, 2021.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, Cambridge, UK, 2006.
- Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. Regret minimization for reserve prices in second-price auctions. *IEEE Trans. Inf. Theory*, 61(1):549–564, 2015.
- Nicolò Cesa-Bianchi, Tommaso Cesari, Roberto Colomboni, Federico Fusco, and Stefano Leonardi. Bilateral trade: A regret minimization perspective. *Math. Oper. Res.*, 49(1):171–203, 2024a.
- Nicolò Cesa-Bianchi, Tommaso Cesari, Roberto Colomboni, Federico Fusco, and Stefano Leonardi. Regret analysis of bilateral trade with a smoothed adversary. *J. Mach. Learn. Res.*, 25:234:1–234:36, 2024b.
- Nicolò Cesa-Bianchi, Tommaso Cesari, Roberto Colomboni, Federico Fusco, and Stefano Leonardi. The role of transparency in repeated first-price auctions with unknown valuations. In *STOC*, pages 225–236. ACM, 2024c.
- Vincent Cohen-Addad, Kasper Green Larsen, David Saulpic, and Chris Schwiegelshohn. Towards optimal lower bounds for k-median and k-means coresets. In *STOC*, pages 1038–1051. ACM, 2022.
- Vincent Cohen-Addad, Andrew Draganov, Matteo Russo, David Saulpic, and Chris Schwiegelshohn. A tight VC-dimension analysis of clustering coresets with applications. In *SODA*, pages 4783–4808. SIAM, 2025a.
- Vincent Cohen-Addad, Silvio Lattanzi, and Chris Schwiegelshohn. Almost optimal PAC learning for k-means. In *STOC*, pages 2019–2030. ACM, 2025b.
- Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *STOC 2014*, pages 243–252. ACM, 2014.
- Riccardo Colini-Baldeschi, Bart de Keijzer, Stefano Leonardi, and Stefano Turchetta. Approximately efficient double auctions with strong budget balance. In *SODA*, pages 1424–1443. ACM-SIAM, 2016.
- Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. Approximately efficient bilateral trade. In *STOC*, pages 718–721. ACM, 2022.
- Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *STOC*, pages 426–439. ACM, 2016.
- Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, Cambridge, UK, 2009.
- Naveen Durvasula, Nika Haghtalab, and Manolis Zampetakis. Smoothed analysis of online non-parametric auctions. In *EC 2023*, pages 540–560. ACM, 2023.
- Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *STOC*, pages 1452–1465. ACM, 2021.

- Paul Dütting, Guru Guruganesh, Jon Schneider, and Joshua Ruizhi Wang. Optimal no-regret learning for one-sided lipschitz functions. In *ICML 2023*, pages 8836–8850, 2023.
- Yumou Fei. Improved approximation to first-best gains-from-trade. In *WINE*, pages 204–218. Springer, 2022.
- Zhe Feng, Chara Podimata, and Vasilis Syrgkanis. Learning to bid without knowing your value. In *EC 2018*, pages 505–522. ACM, 2018.
- Khashayar Gatmiry, Thomas Kesselheim, Sahil Singla, and Yifan Wang. Bandit algorithms for prophet inequality and pandora’s box. In *SODA 2024*, pages 462–500, 2024.
- Yannai A. Gonczarowski and S. Matthew Weinberg. The sample complexity of up-to- ϵ multi-dimensional revenue maximization. *J. ACM*, 68(3):15:1–15:28, 2021.
- Chenghao Guo, Zhiyi Huang, Zhihao Gavin Tang, and Xinzhi Zhang. Generalizing complex hypotheses on product distributions: Auctions, prophet inequalities, and pandora’s problem. In *COLT 2021*, pages 2248–2288, 2021.
- Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis with adaptive adversaries. *J. ACM*, 71(3):19, 2024.
- Ilya Hajiaghayi, MohammadTaghi Hajiaghayi, Gary Peng, and Suho Shin. Gains-from-trade in bilateral trade with a broker. In *SODA*, pages 4827–4860. SIAM, 2025.
- Arun Jambulapati, Yang P. Liu, and Aaron Sidford. Chaining, group leverage score overestimates, and fast spectral hypergraph sparsification. In *STOC*, pages 196–206. ACM, 2023.
- Zi Yang Kang, Francisco Pernice, and Jan Vondrák. Fixed-price approximations in bilateral trade. In *SODA*, pages 2964–2985. ACM-SIAM, 2022.
- Robert Kleinberg, Aleksandrs Slivkins, and Eli Upfal. Bandits and experts in metric spaces. *J. ACM*, 66(4):30:1–30:77, 2019.
- Robert D. Kleinberg and Frank Thomson Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In *FOCS 2003*, pages 594–605. IEEE, 2003.
- James R. Lee. Spectral hypergraph sparsification via chaining. In *STOC*, pages 207–218. ACM, 2023.
- Yi Li, Philip M. Long, and Aravind Srinivasan. Improved bounds on the sample complexity of learning. In *SODA*, pages 309–318. ACM-SIAM, 2000.
- Zhengyang Liu, Zeyu Ren, and Zihe Wang. Improved approximation ratios of fixed-price mechanisms in bilateral trades. In *STOC*, pages 751–760. ACM, 2023.
- R Preston McAfee. The gains from trade under fixed price mechanisms. *Appl. Econ. Res. Bull.*, 1(1):1–10, 2008.
- Jamie Morgenstern and Tim Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *NIPS 2015*, pages 136–144, 2015.
- Roger Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, 1981.

- Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *J. Econ. Theory*, 29(2):265–281, 1983.
- Shyam Narayanan and Jelani Nelson. Optimal terminal dimensionality reduction in euclidean space. In *STOC*, pages 1064–1069. ACM, 2019.
- Jelani Nelson. Chaining introduction with some computer science applications. *Bull. EATCS*, 120, 2016.
- Tim Roughgarden. *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, Cambridge, UK, 2020.
- Tim Roughgarden and Joshua R. Wang. Minimizing regret with multiple reserves. *ACM Trans. Econ. Comput.*, 7(3):17:1–17:18, 2019.
- Atri Rudra and Mary Wootters. Every list-decodable code for high noise has abundant near-optimal rate puncturings. In *STOC*, pages 764–773. ACM, 2014.
- Michel Talagrand. *Upper and lower bounds for stochastic processes*, volume 60. Springer, Berlin, Heidelberg, Germany, 2014.
- William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance*, 16(1): 8–37, 1961.
- Banghua Zhu, Stephen Bates, Zhuoran Yang, Yixin Wang, Jiantao Jiao, and Michael I. Jordan. The sample complexity of online contract design. In *EC 2023*, page 1188. ACM, 2023.

A Characterization of DSIC and IR Mechanisms

In this section, we prove the characterization of DSIC and IR mechanisms.

Proposition 1. *A mechanism M for bilateral trade is dominant-strategy incentive compatible and individually rational if and only if its allocation region is monotone and the payments are Myerson payments.*

Proof. Consider a mechanism with allocation region A and payment functions $q(b_s, b_b)$ and $p(b_s, b_b)$. Let us use the shorthand $x(b_1, b_2) = \mathbb{1}_{\{(b_s, b_b) \in A\}}$ to denote the mechanism's decision whether a pair of bids b_s, b_b should trade.

First assume that A is monotone and that payments are Myerson payments. We want to show that then M is both DSIC and IR. We take into consideration the perspective of the buyer. We fix a bid b_s for the seller and a valuation v_b for the buyer such that bidding truthfully enables a trade, i.e., $(b_s, v_b) \in A$. In this case, overbidding with $b'_b > v_b$ leaves the utility invariant, since a trade still happens and $q(b_s, b'_b) = q(b_s, v_b)$. Specifically, in this case, $U_b(b_s, b'_b) = U_b(b_s, v_b) = v_b - \min\{y \in [0, 1] \mid (b_s, y) \in A\}$, which is non-negative because $\min\{y \in [0, 1] \mid (b_s, y) \in A_M\} \leq v_b$. Underbidding renders the same scenario, up until we reach bids b'_b such that $(b_s, b'_b) \notin A_M$; this leads to the trade not happening and the buyer obtaining zero utility, thus losing the surplus described before. If, on the other hand, the buyer's value v_b is such that with a truthful bid no trade happens, i.e., $(b_s, v_b) \notin A$, then any bid b'_b such that $(b_s, b'_b) \notin A$ keeps the utility at zero, while any bid b'_b such that $(b_s, b'_b) \in A$ will entail a payment $> v_b$ and thus result in negative utility. Therefore, bidding truthfully is the dominant strategy independently of the value of v_b . For what concerns individual rationality, notice that we have already considered the relevant cases for the value of v_b and argued that the induced utility $U_b(b_s, v_b)$ is non-negative. The specular reasoning for the seller, shows that the mechanism is DSIC and IR also for the seller.

We next establish the reverse implication. To this end, we show that DSIC and IR imply both monotonicity of A and Myerson payments. Let us consider the perspective of the seller: we fix a bid b_b for the buyer and two real numbers $0 \leq y \leq z$, with $(y, z) \in [0, 1]^2$. At first, we consider y as the real valuation of the seller, and z a possibly non-truthful bid; incentive compatibility gives the inequality $p(y, b_b) - x(y, b_b)y \geq p(z, b_b) - x(z, b_b)y$. From now on, we disregard the b_b argument since we are keeping it fixed. Conversely, if we let z be the valuation and y the bid, we have: $p(z) - x(z)z \geq p(y) - x(y)z$. By simply factoring out z in the second inequality and y in the first, we get both an upper bound and a lower bound for price changes:

$$z(x(z) - x(y)) \leq p(z) - p(y) \leq y(x(z) - x(y)) \quad (27)$$

Notice that the assumption of x increasing immediately leads to a contradiction with the inequalities above, so x is weakly decreasing when taking as argument only the bid for the seller. An analogous argument applies to the the buyer, and shows that x is weakly increasing when taking as argument only the bid of the buyer. We have thus established monotonicity.

To show the need for Myerson payments, let us focus again on the seller. Notice that $x(\cdot, b_b)$ can only take two values: either we are in the allocation region and the trade happens, or not. This immediately implies that the payment $p(\cdot, b_b)$ can only take two different values, one for the case where seller and buyer trade and one for the case where there is no trade. We next argue that if the discontinuity of $x(\cdot, b_b)$ happens at y (i.e., $x(y, b_b) = 1$ while for any $z > y$ we have $x(z, b_b) = 0$), then the difference between these two payments must be y . To this end, note that the lower and upper bounds from Equation (27) both converge to $-y$ as $z \rightarrow y$ from above, meaning that $\lim_{z \rightarrow y^+} p(z) - p(y) = -y$. We have thus shown that $p(b_s, b_b) = y + c_{b_b}$ for $b_s \leq y$ and $p(b_s, b_b) = c_{b_b}$

otherwise, where $c_{b_b} \geq 0$ is a non-negative constant that may depend on the buyer’s bid. By our requirement that $p(b_s, b_b) = 0$ when $(b_s, b_b) \notin A$ and there is no trade, we must have $c_{b_b} = 0$. An analogous reasoning for the buyer completes the proof. \square

Remark 4. *The requirement that payments are bounded in $[0, 1]$ and zero when there is no trade is without loss of generality, when focusing on DSIC and IR mechanisms and the goal of maximizing the broker’s profit. This can be seen by considering general payment functions $p, q : [0, 1]^2 \rightarrow \mathbb{R}$. Then an argument analogous to the one used to establish Proposition 1 shows that the payments must be Myerson payments, up to constant terms c_{b_b} and c_{b_s} that are added to the seller’s and buyer’s payment, respectively. By IR these constant must be such that $c_{b_b} \geq 0$ and $c_{b_s} \leq 0$, and, to maximize profit, we want $c_{b_b} = c_{b_s} = 0$.*

B Efficiency Maximization and Budget Balance

In this section, we discuss the efficiency maximization version of the bilateral trade problem. Let M be a mechanism with allocation region A , the social welfare and gain from trade induced by M on agents with valuations v_s and v_b are defined as follows:

$$\begin{aligned} \text{GFT}(M, v_s, v_b) &= \mathbb{1}_{\{(v_s, v_b) \in A\}}(v_b - v_s) && \text{(Gain-From-Trade)} \\ \text{SW}(M, v_s, v_b) &= \mathbb{1}_{\{(v_s, v_b) \in A\}}(v_b - v_s) + v_s && \text{(Social Welfare)} \end{aligned}$$

Note that social welfare is an additive factor away from gain from trade; therefore, from the regret perspective, these two objectives are exactly equivalent.

Budget Balance. Beyond DSIC and IR, another natural property required to efficiency maximizing mechanisms for bilateral trade is budget balance: we want in fact to avoid situations in which the mechanism subsidizes the market. There are two main notions of budget balance in the literature [see e.g., Cesa-Bianchi et al., 2024a, Colini-Baldeschi et al., 2016, Bernasconi et al., 2024]:

$$\begin{aligned} \text{profit}(M, v_s, v_b) &= 0 \quad \forall (v_s, v_b) \in [0, 1]^2 && \text{(Strong Budget Balance)} \\ \text{profit}(M, v_s, v_b) &\geq 0 \quad \forall (v_s, v_b) \in [0, 1]^2 && \text{(Weak Budget Balance)} \end{aligned}$$

We have all the notions to formally recall the impossibility result by Myerson and Satterthwaite [1983]: no mechanism can, at the same time, achieve full efficiency (e.g., maximize social welfare/gain from trade), individual rationality, incentive compatibility and (weak) budget balance.

Fixed-Price Mechanisms. An important class of mechanisms is composed by *fixed-price mechanisms*, where the broker simply posts two fixed prices p to the seller and q to the buyer, and the trade happens if and only if both agents agree. In the rest of the section, we clarify what we mean in the introduction by claiming that fixed price mechanisms “characterize” efficiency maximizing mechanisms. We start by recalling a folklore result on strong budget balance mechanisms.

Proposition 8 (e.g., Colini-Baldeschi et al. [2016]). *A mechanism M for bilateral trade is dominant-strategy incentive compatible, individually rational, and strongly budget balanced if and only if it is a fixed price mechanism for some price p , i.e., if its allocation region is $[0, p] \times [p, 1]$.*

Weak budget balance has a nice geometric interpretation following the characterization of Proposition 1. For any non-trivial mechanism M (i.e., with non-empty allocation region), let \bar{s} be the price paid to the seller and \underline{b} be the price paid by the buyer, under valuation $(0, 1)$. From a

geometrical point of view, these two numbers correspond to the intersections of the boundary of M with the left and top side of the $[0, 1]^2$. We have the following result.

Proposition 9 (Characterization of Weak-Budget Balanced Mechanism). *A mechanism M for bilateral trade that is dominant-strategy incentive compatible and individually rational enforces weak budget balance if and only if $\bar{s} \leq \underline{b}$. Moreover, the social welfare and gain from trade of any such mechanism is dominated by the fixed price mechanism with allocation region $[0, \bar{s}] \times [\underline{b}, 1]$.*

Proof. Clearly, if a mechanism does not respect $\bar{s} \leq \underline{b}$, then it violates budget balance under valuation $(0, 1)$. The other implication follows by noting that, for any valuation (v_s, v_b) that results in a trade, we have: $p(v_s, v_b) \leq \bar{s} \leq \underline{b} \leq q(v_s, v_b)$.

Consider now the last part of the statement. Increasing the allocation region while retaining budget balance and monotonicity always (weakly) increases efficiency. Therefore taking the maximal monotone region compatible with \underline{b} and \bar{s} is clearly a (weak) improvement. \square

Budget Balance and Profit. Since the focus of this paper is on profit maximization, there is no reason to constrain the optimization procedure to DSIC and IR mechanisms that also enforce weak budget balance (clearly, the profit of any strongly budget balance mechanism is always 0). In fact, there are simple examples in which there is a large gap between the expected profit of the best mechanism and the best budget balance one.

Example 4 (Budget Balance is Suboptimal). *Consider the uniform distribution over $(0, 1/4)$ and $(3/4, 1)$. The optimal mechanism allocates in the $(0, 1/4)$, $(3/4, 1)$, $(0, 1)$ triangle, for an expected profit of $1/4$. Such mechanism is not budget balanced, as $\bar{s} = 3/4$ and $\underline{b} = 1/4$. The allocation region of any budget balance mechanism can either contain $(0, 1/4)$ or $(3/4, 1)$, for an expected profit of most $1/8$.*

Beyond Dominant-Strategy Incentive Compatibility. The impossibility results of Myerson and Satterthwaite [1983] hold in the more general setting where mechanisms only need to enforce incentive compatibility, individual rationality, and budget balance *in expectation*. These mechanisms clearly do not respect the characterization provided in Propositions 1, 8 and 9, but have also been studied. In particular, the first mechanism to achieve a constant factor approximation of the optimal gain from trade in the Bayesian setting [Deng et al., 2022] is BIC but not DSIC.

C The Learning Complexity of \mathcal{M}

In this section, we argue that it is not possible to uniformly learn the expected profit of *all* mechanisms in \mathcal{M} for bilateral trade. To provide some further intuition, we also explicitly prove that standard statistical-complexity measures are not helpful (VC dimension, pseudo-dimension, and Rademacher complexity). Formally, we have the following theorem.

Theorem 8. *There exists a distribution \mathcal{D} such that for any number n of i.i.d. samples S , there exists a mechanism $M \in \mathcal{M}$ that respects the following two properties:*

- *The expected profit of M under \mathcal{D} is constant: $\mathbb{E}[\text{profit}(M)] \geq 1/4$*
- *The empirical profit of M on the samples S is zero.*

Proof. This proof formalizes the argument seen in Example 3. Let \mathcal{D} be the uniform distribution on the $(0, 3/4)$ - $(1/4, 1)$ segment, and consider any realization of the n samples S . Denote with $\delta > 0$ a small parameter we set later, and with $B_\delta(v)$ the generic ℓ^∞ open ball around v of radius δ .




Figure 6: Visualization of the proofs of Theorem 9 (left) and Theorems 10 and 11 (right).

Consider the mechanism M whose allocation region is the $(0, 3/4), (1/4, 1), (0, 1)$ triangle minus the balls $B_\delta(v)$ for v in S .

By construction, the empirical profit of M on the samples is zero, while the expected profit on \mathcal{D} is at least $1/4$, for δ small enough. In fact, S is finite, so it is possible to take δ small enough so that at least half of the support of \mathcal{D} falls in the allocation region of M . For such δ , a trade happens with probability at least $1/2$, with an expected profit of at least $1/4$. \square

Definition 12 (VC Dimension for Range Spaces). *Let \mathcal{X} be a set, and \mathcal{G} a collection of subsets of \mathcal{X} . The tuple $(\mathcal{X}, \mathcal{G})$ is called a range space. For each finite set $S \subseteq X$ we define the projection of \mathcal{G} onto S as $\mathcal{G}_S = \{G \cap S : G \in \mathcal{G}\}$; we say that S is shattered by \mathcal{G} if $|\mathcal{G}_S| = 2^{|S|}$, with \mathcal{G}_S thus the power set of S . The VC-dimension of the range space $(\mathcal{X}, \mathcal{G})$ is the cardinality of the largest S that can be shattered by \mathcal{G} . If sets of arbitrary high cardinality can be shattered, we say that the VC-dimension is unbounded.*

Theorem 9 (Unbounded VC Dimension). *The VC dimension of the collection of sets induced by monotone allocations regions is unbounded.*

Proof. Let \mathcal{A} be the family of all monotone allocation regions. Consider any set S of n points on the diagonal $\{(x, y) : x = y\} \cap [0, 1]^2$, we prove that S can be shattered by \mathcal{A} . As an auxiliary construction, for any point (x_i, x_i) in S , define the rectangle

$$R_i = [0, x_i] \times [x_i, 1]. \quad (28)$$

For any subset S' of S , consider the region $A' = \bigcup_{(x_i, x_i) \in S'} R_i$; to conclude the proof it is enough to argue that the following three properties are verified: (i) A' is monotone, (ii) A' contains all the points in S' , and (iii) A' does not contain any point in $S \setminus S'$.

We start proving monotonicity. Let \hat{v} be any point in $[0, 1]^2$ that dominates some v' in A' , by definition v' belongs to some R_i with $(x_i, x_i) \in S'$, implying that $(x_i, x_i) \preceq v'$. By transitivity, \hat{v} dominates (x_i, x_i) , meaning that it belongs to R_i (since each R_i is monotone), thus belonging to A' . We are left with proving points (ii) and (iii). The monotone region A' is the union of the R_i for $i : (x_i, x_i) \in S'$, and each R_i only touches the main diagonal in the corresponding point (x_i, x_i) . Thus, the intersection of A' with the main diagonal is exactly S' . \square

Definition 13 (Pseudo-dimension). *For a class of real-valued functions, we say that a set of points, here $S = \{v_i\}_{i=1}^m$, is pseudo-shattered if \exists a vector $r \in \mathbb{R}^m$ such that $\forall b \in \{\pm 1\}^n \exists f_b \in \mathcal{F}$ such that $\text{sign}(f_b(v_i) - r_i) = b_i \forall v_i \in C$. The pseudo-dimension is the maximum cardinality of a set which is pseudo-shattered by \mathcal{F} .*

Theorem 10 (Unbounded Pseudo-dimension). *The pseudo-dimension of the class of profit functions \mathcal{F} is unbounded.*

Proof. Fix any set $S \subset C = \{(x, y) : x = y - \frac{1}{2}\} \cap [0, 1]^2$, with $|S| = m$. If S is a set of valuations, we have that each valuation vector lies on C , meaning that the difference between its entries is always $\frac{1}{2}$. With respect to Definition 13, here $\mathcal{F} = \{\text{profit}(M, \cdot), M \in \mathcal{M}\}$, i.e. the profit functions induced by valid mechanisms, as characterized in Proposition 1.

Let us fix an arbitrary vector of signs b : this induces a partition of S into S' and $S \setminus S'$. We can build a mechanism M' and a monotone allocation region A' such that the corresponding profit function $\text{profit}(M', \cdot) \in \mathcal{F}$ assigns $1/2$ to points in S' and 0 otherwise. More precisely, almost identically to Theorem 9, $A' = \cup_{i: v_i \in S'} R_i$, where the rectangles $R_i = [0, x_i] \times [y_i, 1]$ and $v_i = (x_i, y_i)$. If we take $r \in \mathbb{R}^m$ to be the constant vector with entries equal to $1/4$, we immediately have that $\text{sign}(\text{profit}(M', v_i) - r_i) = b_i$. No assumption was made on b , so the same holds for any choice of b . Therefore, S is pseudo-shattered by \mathcal{F} . The procedure is valid $\forall m \in \mathbb{N}$: the pseudo-dimension of \mathcal{F} is thus unbounded. \square

Definition 14 (Rademacher Complexity). *Let $\{v_i\}_{i=1}^m$ be a i.i.d. sample of size m , such that each $v_i \sim \mathcal{D}$, where \mathcal{D} is a distribution over \mathbb{R}^n . Let \mathcal{F} be generic class of real-valued functions on \mathbb{R}^n . Let $\{\sigma_i\}_{i=1}^n$ be a set of i.i.d Rademacher variables (i.e., $\sigma_i \in \{-1, +1\}$ with equal probability). The Rademacher complexity of the tuple $(\mathcal{D}, \mathcal{F})$ is defined as follows:*

$$R_m(\mathcal{D}, \mathcal{F}) = \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i f(v_i) \right| \right].$$

Theorem 11 (Lower Bound on Rademacher Complexity). *There exists a distribution over $[0, 1]^2$ such that the Rademacher complexity for the class \mathcal{F} of profit functions is at least $1/4$.*

Proof. We use the same construction of Theorem 10. Consider the uniform distribution \mathcal{D} over the segment $C = \{(x, y) : x = y - \frac{1}{2}\} \cap [0, 1]^2$. The segment C is characterized by a constant difference of $1/2$ between the buyer and the seller valuations.

First, we fix a realization of the i.i.d. sample $S = \{v_i\}_{i=1}^m$, where $v_i \sim \mathcal{D}$. Here we abuse notation using the same symbol for both a random variable and its realization. Therefore, as for the proof of Theorem 10, $S \subset C$; also, $\mathcal{F} = \{\text{profit}(M, \cdot), M \in \mathcal{M}\}$. Note, we assume without loss of generality that the points in S are disjoint: this happens almost surely.

Fix any realization of the Rademacher random variables and denote with S' the subset of S associated to points whose σ_i is 1. Without loss of generality, we assume that $|S'| \geq m/2$ (if this is not the case, then the same argument holds for $S \setminus S'$). Consider now the mechanism associated with the allocation region $A' = \cup_{i: v_i \in S'} R_i$, where the rectangles R_i are defined as in the proof of Theorem 10. A' contains all the points in S' and no point in $S \setminus S'$. Moreover, the associated mechanism M' extracts exactly profit $1/2$ from each point in S' and 0 from the points not in S . We then have:

$$\sup_{M \in \mathcal{M}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \cdot \text{profit}(M, v_i) \right| \geq \left| \frac{1}{m} \sum_{v_i \in S'} \sigma_i \cdot \text{profit}(M', v_i) \right| = \frac{1}{m} \sum_{v_i \in S'} \frac{1}{2} \geq \frac{1}{4},$$

where in the last inequality we use the assumption $|S'| \geq m/2$.

If we now take the expectation with respect to the Rademacher random variables, we get a $1/4$ lower bound. This means that the same bound also holds when taking an outer expectation with respect to the distribution of S , since it holds for almost every realization of S . We have then proved the desired statement:

$$\mathbb{E}_{S,\sigma} \left[\sup_{M \in \mathcal{M}} \left| \frac{1}{m} \sum_{i=1}^m \sigma_i \cdot \text{profit}(M, v_i) \right| \right] \geq \frac{1}{4},$$

Note, our bound crucially *does not* depend on the number m of samples. □

D Concentration Inequality

In the main body, we use the following version of multiplicative Chernoff–Hoeffding bound that we use (restatement of Theorem 1.1. of Dubhashi and Panconesi [2009] for variables in $[0, 1]$).

Theorem 12. *Let X be defined as the sum of m random variables X_i , independently distributed in $[0, 1]$. Let $\varepsilon \in (0, 1)$ and $t > 2\varepsilon\mathbb{E}[X]$ be any precision parameters, then the following inequalities hold:*

- (i) $\mathbb{P}(X - \mathbb{E}[X] < -\varepsilon\mathbb{E}[X]) \leq \exp(-\frac{1}{2}\varepsilon^2\mathbb{E}[X])$
- (ii) $\mathbb{P}(X > t) \leq 2^{-t}$

E Proof of Proposition 7

For completeness, we report here a direct proof of Section E from Aggarwal et al. [2024].

Proposition 7. *For any realization of the valuations and auxiliary sequence, we have:*

- (i) $a_t \in [0, \delta]$ and $b_t \in [0, \delta]$, for all $t = 1, \dots, T$.
- (ii) there exists a value $\tau \in (0, \delta)$ such that $b_t < \tau$ for all $t \in R$, while $a_t > \tau$ for all $t \in L$.

Proof. Consider any realization of the random sequence (a_t, b_t) . We start by providing a more direct characterization of its elements, summarizing the “left” and “right” decisions in an auxiliary variable γ_t , which has value -1 if at time t the sequence goes left, and value 1 otherwise. We also introduce the middle point m_t of the segment $[a_t, b_t]$, and rewrite a_t and b_t accordingly:

$$m_t = \frac{\delta}{2} + \sum_{s=1}^{t-1} \gamma_s \frac{\delta}{3^s}, \quad a_t = m_t - \frac{\Delta_t}{2}, \quad b_t = m_t + \frac{\Delta_t}{2}. \quad (29)$$

Claim 3. *The new definition is equivalent to the old one.*

Proof of the Claim. We prove this result by induction. The base case is simple to verify, as $m_1 = 1/2$, $a_1 = 1/3$, and $b_1 = 2/3$ in both cases. Consider now the generic time step $t + 1$ and assume that the

two definitions are equivalent up to time t ; we prove that Equation (29) is equivalent to the one in the main body for $t \in L$ and $\gamma_t = -1$ (the other case is analogous):

$$\begin{aligned} a_{t+1} &= m_{t+1} - \frac{\Delta_{t+1}}{2} = m_t - \frac{\delta}{3^t} - \frac{\Delta_{t+1}}{2} = a_t + \frac{\Delta_t}{2} - \frac{\delta}{3^t} - \frac{\Delta_{t+1}}{2} = a_t - 2\Delta_{t+1} \\ b_{t+1} &= m_{t+1} + \frac{\Delta_{t+1}}{2} = m_t - \frac{\delta}{3^t} + \frac{\Delta_{t+1}}{2} = a_t + \frac{\Delta_t}{2} - \frac{\delta}{3^t} + \frac{\Delta_{t+1}}{2} = a_t - \Delta_{t+1}. \end{aligned}$$

Note, all equalities follow by Equation (29), except the third equalities in both formulae that are due to the inductive hypotheses. \square

The proof of point (i) is fairly direct; we prove it for a_t , the other case is analogous. Any γ_t lives in $\{-1, 1\}$, therefore we have:

$$\begin{aligned} a_t &= m_t - \frac{\Delta_t}{2} = \frac{\delta}{2} + \sum_{s=1}^{t-1} \gamma_s \frac{\delta}{3^s} - \frac{\Delta_t}{2} \geq \frac{\delta}{2} - \sum_{s=1}^t \frac{\delta}{3^s} \geq \frac{\delta}{2} - \sum_{s=1}^{\infty} \frac{\delta}{3^s} = 0 \\ a_t &= m_t - \frac{\Delta_t}{2} = \frac{\delta}{2} + \sum_{s=1}^{t-1} \gamma_s \frac{\delta}{3^s} - \frac{\Delta_t}{2} \leq \frac{\delta}{2} + \sum_{s=1}^{\infty} \frac{\delta}{3^s} = \delta. \end{aligned}$$

As a preliminary step to tackle point (ii), we argue that the a_t and b_t are monotone *when* restricted to L and R , respectively.

Claim 4. *The sequence a_t is decreasing for $t \in L$, while b_t is increasing for $t \in R$.*

Proof of the Claim. We prove the Claim for a_t , the other case follows by negating the expression in Equation (30). Consider any consecutive $t_1 \leq t_2$ in L (meaning that all integer in (t_1, t_2) are in R , and so the corresponding γ_t are positive), we have the following:

$$a_{t_1} - a_{t_2} = \sum_{s=t_1+1}^{t_2-1} \frac{\delta}{3^s} - \frac{\delta}{3^{t_1}} - \frac{\Delta_{t_1}}{2} + \frac{\Delta_{t_2}}{2} \geq \sum_{s=t_1+1}^{t_2-1} \frac{\delta}{3^s} + \frac{\Delta_{t_2}}{2} > 0 \quad (30)$$

\square

Given the result of the last Claim, in order to prove point (ii) of the Proposition, we only need to argue that $a_{t_L} \geq b_{t_R}$, where t_L , respectively t_R , are the last “left”, respectively “right” rime step. Once we have this result, it is enough to set τ in (a_{t_L}, b_{t_R}) .

We prove the desired inequality for the case in which $t_L > t_R$ (and thus $\gamma_t = -1$ for every integer t in $(t_R, t_L]$).

$$\begin{aligned} b_{t_R} - a_{t_L} &= - \sum_{s=t_R}^{t_L-1} \gamma_s \frac{\delta}{3^s} + \frac{\Delta_{t_L}}{2} + \frac{\Delta_{t_R}}{2} && \text{(By Equation (29))} \\ &= - \frac{\delta}{3^{t_R}} + \sum_{s=t_R+1}^{t_L-1} \frac{\delta}{3^s} + \frac{\delta}{2 \cdot 3^{t_L}} + \frac{\delta}{2 \cdot 3^{t_R}} && \text{(As } \gamma_t = -1 \text{ after } t_R \text{ and } \gamma_{t_R} = 1) \\ &< \sum_{s=t_R+1}^{t_L} \frac{\delta}{3^s} - \frac{\delta}{2 \cdot 3^{t_R}} < \sum_{s=t_R+1}^{\infty} \frac{\delta}{3^s} - \frac{\delta}{2 \cdot 3^{t_R}} \leq 0, \end{aligned}$$

where the last inequality follows by the closed-form expression of the tail sum of a geometric series.

When $t_R > t_L$ we have (recall that in this case $\gamma_t = 1$ for every integer $t \in (t_L, t_R]$):

$$\begin{aligned}
b_{t_R} - a_{t_L} &= \sum_{s=t_L}^{t_R-1} \gamma_s \frac{\delta}{3^s} + \frac{\Delta_{t_L}}{2} + \frac{\Delta_{t_R}}{2} && \text{(By Equation (29))} \\
&= -\frac{\delta}{3^{t_L}} + \sum_{s=t_L+1}^{t_R-1} \frac{\delta}{3^s} + \frac{\delta}{2 \cdot 3^{t_L}} + \frac{\delta}{2 \cdot 3^{t_R}} && \text{(As } \gamma_t = 1 \text{ after } t_L \text{ and } \gamma_{t_L} = -1) \\
&< \sum_{s=t_L+1}^{t_R} \frac{\delta}{3^s} - \frac{\delta}{2 \cdot 3^{t_L}} < \sum_{s=t_L+1}^{\infty} \frac{\delta}{3^s} - \frac{\delta}{2 \cdot 3^{t_L}} \leq 0,
\end{aligned}$$

□