Differentially Private Call Auctions and Market Impact

EMILY DIANA, University of Pennsylvania
HADI ELZAYN, University of Pennsylvania
MICHAEL KEARNS, University of Pennsylvania
AARON ROTH, University of Pennsylvania
SAEED SHARIFI-MALVAJERDI, University of Pennsylvania
JUBA ZIANI, University of Pennsylvania

We propose and analyze differentially private (DP) mechanisms for call auctions as an alternative to the complex and ad-hoc privacy efforts that are common in modern electronic markets. We prove that the number of shares cleared in the DP mechanisms compares favorably to the non-private optimal and provide a matching lower bound. We analyze the incentive properties of our mechanisms and their behavior under natural no-regret learning dynamics by market participants. We include simulation results and connections to the finance literature on market impact.

CCS Concepts: • **Theory of computation → Algorithmic game theory and mechanism design**; **Market equilibria**; **Computational pricing and auctions**; **Convergence and learning in games**; • **Security and privacy → Economics of security and privacy**.

Additional Key Words and Phrases: call auctions; mechanism design; differential privacy; learning in games

ACM Reference Format:
Emily Diana, Hadi Elzayn, Michael Kearns, Aaron Roth, Saeed Sharifi-Malvajerdi, and Juba Ziani. 2020. Differentially Private Call Auctions and Market Impact. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC '20), July 13–17, 2020, Virtual Event, Hungary*. ACM, New York, NY, USA, 43 pages. https://doi.org/10.1145/3391403.3399500

1 INTRODUCTION AND OVERVIEW OF PAPER

In modern financial markets, massive resources are directed towards what can be considered ad-hoc privacy mechanisms, intended to allow participants to cloak their trading activity and intentions. Such efforts occur both in the exchanges themselves and in the algorithmic trading services offered by large brokerages. In this work, we provide a differentially private (DP) version of classical one-shot double auctions (also known as "call auctions"). Frequent instances of DP call auctions could potentially simplify the convoluted efforts at providing trading secrecy that are rampant in today's markets while still permitting dynamic price discovery.

Current electronic exchanges offer a staggering variety of order types and mechanisms meant to provide specific types of privacy. Dark pools were introduced to allow large-volume counterparties to discover each other away from the so-called "lit" markets where high-frequency traders (HFTs) are prevalent. Order types restricting execution with small-volume counterparties are meant to provide similar protections. Hidden and "iceberg" orders in the lit exchanges provide secrecy at the expense of time priority in the standard continuous limit order book. The relatively new exchange

EC '20, July 13–17, 2020, Virtual Event, Hungary

ACM ISBN 978-1-4503-7975-5/20/07…$15.00
https://doi.org/10.1145/3391403.3399500

IEX was created to foil the latency arbitrage of HFT by introducing a "speed bump" for all incoming orders. On the brokerage side, algorithms executing large client trades attempt to minimize visibility by breaking orders up over time and across exchanges and employ randomization in both timing and sizing to avoid detectable "heartbeats."

These efforts are all ad-hoc in the sense that they each protect market participants from rather specific forms of detection or exploitation. While well-intentioned, they have contributed significantly to the complexity of modern electronic markets. At the same time, it is also widely understood that there are limits to the privacy that can be provided for large trades executed in short periods, and there is a large academic and practical literature on theories of market impact (see [11, 12] for an overview) and algorithms for minimizing it. This literature identifies a trade's *participation rate* — the ratio of its volume to that of the overall market during the trade's execution — as the key determinant of market impact.

Our main conceptual contribution is the development of DP call auctions as a mechanism providing privacy against *all* forms of attack or detection, up to the participation rate of a trade. In this formulation, we provide a per-share privacy guarantee determined by the sensitivity of the call auction, which in turn determines the amount of noise added. Trades with higher participation rates will unavoidably have less privacy than those with smaller ones, but the nature of the privacy will now be as general as possible. Repeated DP call auctions also enjoy graceful degradation of the privacy guarantee. Furthermore, we can (informally) relate our results to standard market impact theories via the shared notion of participation rate and show that, under natural conditions, DP call auctions clear a near-optimal number of shares under the predictions of the "square root law" of market impact. We analyze our DP mechanisms extensively, including its incentive properties and behavior under natural no-regret learning dynamics by market participants.

We note that (non-private) call auctions are already common in modern markets. In particular, both NYSE and NASDAQ hold call auctions (also sometimes called "crosses") to establish opening and closing prices in U.S. equities [19, 20]; in the Tokyo Stock Exchange there are additional intraday call auctions, which are also the subject of academic study (e.g. [2, 3]). The influential paper [1] (discussed at greater length in Related Work below) proposes and analyzes frequent intraday (again non-private) call auctions specifically as a defense against latency arbitrage; see also [25]. Our work can be seen as a continuation of this line of thinking, in which frequent intraday DP call auctions could provide even more general privacy guarantees to all market participants.

Outline and Summary of Results: At a high level, our results fall into three broad categories:

(1) **The development and analysis of (jointly) differentially private call auctions**. We carry this out in Section 3. We initially present this purely as an algorithm design task, abstracting away incentive properties. We prove bounds relating the *privacy* properties of the mechanism, the *number* of shares it is guaranteed to clear compared to the optimal benchmark, and the net *inventory* that the mechanisms may have to take on. (Unavoidably, jointly differentially private call auctions cannot exactly match the number of buyers and sellers and so will have to take on a net position of shares itself to clear the market — we prove that this net position is small.). We also prove a lower bound showing that our mechanisms are near optimal amongst all differentially private mechanisms. We explore the connection between our guarantees and theories of market impact in Section 3.5.

(2) **The analysis of incentive properties and learning dynamics**. Having developed our algorithms, we turn our attention to how buyers and sellers should interact with them. First, in Section 4, we show that our algorithm is ex-post individually rational and approximately dominant strategy truthful for agents who wish to trade only a small number of shares, with a guarantee that degrades gracefully in the size of the desired trade. (We note that this

is a *stronger* incentive guarantee than standard non-private call auctions.) We then study the global behavior that results when agents interact with a repeated version of one of our mechanisms using *learning dynamics*: we show that although an abstract guarantee of no-regret learning is not enough to guarantee convergence to the optimal number of trades, a small modification of the exponential weights learning algorithm (informally, a modification that still guarantees the no-regret property, but breaks ties in favor of trading whenever such ties exist) does converge to the optimal number of trades.

(3) **Simulation Results**. Finally, in Section 5, we conduct simulations in both *one-shot* and *repeated* settings, showing that in the settings considered, the realized outcomes of our mechanisms tend to be significantly better than the worst-case guarantees of our theorems.

Related Work: Our work relates to several large strands of literature. Prominently, the study of double auctions dates back to the early days of mathematical economics. [21] provides an introduction to double auctions, and a useful survey from a computer science perspective can be found in [22]. Our modeling of the strategic framework in which agents participate in the double auction is broadly consistent with this literature.

Of particular note are [1] and [25], which both propose frequent call auctions to eliminate latency arbitrage.[1] The work of [1] first establishes the *empirical* availability of latency arbitrage opportunities for even highly traded securities, and shows moreover that competition between traders has not eliminated this opportunity over time. Instead, it has resulted in an "arms race" for speed, with arbitrage windows becoming shorter over time, but arbitrage profit per unit remaining essentially constant. The authors then propose a solution to mitigate latency arbitrage: repeated high-frequency call auctions. Using a game theoretic approach, they model how the "sniping" process results in arbitrage opportunities in the continuous limit order book; in their model, the profit opportunity (along with arms race) is an equilibrium constant, even despite improving technology. Then, using the same underlying model of firm behavior, the authors show that repeated call auctions eliminate these arbitrage opportunities and cause firms not to choose to invest in speed, ending the arms race. We follow in the spirit of [1], but note that their solution does not mitigate the problem of privacy, and in particular does not solve the issue of the proliferation of ad-hoc and increasingly complex trading algorithms. (The earlier work of [25] performs extensive simulation studies that establish the salutary effects of frequent call auctions on latency arbitrage.)

Our work is connected to, and leverages tools from, the broad literature on differential privacy [7]; for an overview, see, e.g. [8]. The most related strand of this literature is the connection between differential privacy and mechanism design, first made by [18]. In particular, they observed that differentially private mechanisms inherit strong incentive properties. For many mechanism design tasks that involve the allocation of a resource to individuals, it is not possible to satisfy differential privacy in the standard sense over allocations: in cases like this, the relevant solution concept is *joint differential privacy* [16]. This solution concept has been used in a number of mechanism design settings, including max-welfare matchings and other allocation problems [13, 14], stable matchings [15], equilibrium selection problems [5, 24], and tolling problems [23]. In particular, although joint differential privacy can be used as a tool to achieve truthfulness, not all jointly private mechanisms are approximately truthful, and more specialized arguments are needed. Finally, while [4] have shown how to privately compute near-optimal prices in double auctions, their process does not guarantee end-to-end joint differentially privacy when taking trade allocations into account, unlike this work.

[1] *Latency arbitrage* is the opportunity for traders to simultaneously buy and sell nearly or exactly identical securities on different exchanges (e.g. Chicago's Mercantile Exchange and the New York Stock Exchange) in the instant where price has changed on one exchange but remains "stale" on the other; it is described in the popular book *Flash Boys* [17].

2 MODEL AND PRELIMINARIES

2.1 Model

We consider a call auction setting with n^s sellers and n^b buyers; we let $\mathcal{S}$ be the set of sellers, $\mathcal{B}$ be the set of buyers, and $n = n^s + n^b$. Each seller $i \in \mathcal{S}$ has one unit of a security for which it has a *value* v_i^s; each buyer $j \in \mathcal{B}$ wishes to purchase one unit of the security for which it has a value of v_j^b. We let $\mathbf{v}^s = (\mathbf{v}_1^s, \ldots, \mathbf{v}_{n^s}^s)$ be the vector of all sellers' valuations and $\mathbf{v}^b = (\mathbf{v}_1^b, \ldots, \mathbf{v}_{n^b}^b)$ the vector of all buyers' valuations. We assume valuations are drawn from a discrete set P; without loss of generality, we let $P = \{1, 2, \ldots, V\}$ for some integer V.

Agents report their valuations in P directly to a mechanism $\mathcal{M}$.[2] Based on the agents' reports, the mechanism selects a clearing price $p \in P$ and an allocation vector $\mathbf{a} = (\mathbf{a}^s, \mathbf{a}^b)$, where $\mathbf{a}_i^s$ (resp. $\mathbf{a}_j^b$) is equal to 1 if seller i (resp. buyer j) is selected to participate in a trade and 0 otherwise. The mechanism concludes by buying a share at price p from every seller i with $\mathbf{a}_i^s = 1$ and selling a share at price p to every buyer with $\mathbf{a}_j^b = 1$.[3]

Privacy Constraints. The outcomes of the mechanisms we consider are functions of the agents' reports, which themselves depend on their valuations. In turn, these outcomes may leak information about the participants' valuations. This provides motivation for designing call auctions that protect the privacy of the participants. In this paper, we do so using *differential privacy* ([7]). We will design our mechanisms to release the clearing price p in a differentially private fashion, and the allocation vector $\mathbf{a} = (\mathbf{a}^s, \mathbf{a}^b)$ in a *jointly* differentially private manner [16]. Differential privacy and joint differential privacy are formally defined in Section 2.2.

Mechanism Designer's Objective. The main objective of our mechanisms for call auctions is to maximize the volume of trades between buyers and sellers. However, because of the randomization that we will need to add to achieve differential privacy, our mechanism will inevitably incur several kinds of cost. First, the *payoff* of the mechanism, given by the number of shares cleared, will generally be lower than the optimal payoff that could have been reached absent differential privacy. Second, we will have to deal with situations in which the number of sellers and the number of buyers who are selected to trade differ because of the noise added to the allocation rule for privacy concerns; this creates an *inventory* in which some of the trades must be fulfilled by the mechanism itself (when there are more sellers selected than buyers, the mechanism buys surplus shares from the sellers; when there are more buyers selected than sellers, the mechanism sells to the buyers from its own reserve of shares). We will aim to keep the inventory of our private auction mechanisms as small as possible. Formally, the payoff and the inventory of a mechanism $\mathcal{M}$ are defined as follows:

Definition 1 (Payoff and Inventory of a Mechanism $\mathcal{M}$). *For any mechanism $\mathcal{M}$ outputting a price p and an allocation vector $\mathbf{a}$, the payoff is the number of shares cleared by $\mathcal{M}$: $\Pi(\mathcal{M}) = \min\{\sum_{i\in\mathcal{S}} \mathbf{a}_i^s, \sum_{j\in\mathcal{B}} \mathbf{a}_j^b\}$, and the inventory of $\mathcal{M}$ is the number of allocations that must be fulfilled by the mechanism: $I(\mathcal{M}) = |\sum_{i\in\mathcal{S}} \mathbf{a}_i^s - \sum_{j\in\mathcal{B}} \mathbf{a}_j^b|$.*

The main benchmark we use to measure the performance of our mechanisms is the maximum number of trades that can be obtained (absent differential privacy) while setting a uniform price p

[2] We will argue in Section 4 that it is in every agent's best interest to report his valuation to the mechanism truthfully, hence our mechanisms can work with the agents' valuations without loss of generality.

[3] In principle, mechanisms can choose *non-uniform* pricing; that is, different agents could be charged different prices based on their reports. Here, we only consider *uniform* pricing mechanisms, as is commonplace in the double auction literature.

and guaranteeing every agent non-negative utility.[4] Formally, our benchmark is given by[5]

$$\text{OPT} = \max_{p \in P} \min \left\{ \sum_{i \in \mathcal{S}} \mathbf{1}\left[\mathbf{v}_i^s \leq p\right], \sum_{j \in \mathcal{B}} \mathbf{1}\left[\mathbf{v}_j^b \geq p\right] \right\}. \tag{1}$$

2.2 Differential Privacy

Let $\mathcal{D}$ be a *data universe* from which a data set D of size n is drawn. In the setting considered in this paper, $D = (\mathbf{v}^s, \mathbf{v}^b)$ contains the reported valuations of sellers and buyers in the market. The algorithms we consider in this paper have output that can naturally be partitioned across the n users who provide the inputs — namely for each agent, whether they get to participate in a trade, and at what price. Let $\mathcal{M}$ be an algorithm that takes the data set D as input and outputs $\mathcal{M}(D) \in \mathcal{R}^n$, which is a vector whose ith coordinate corresponds to the output sent to agent i. Here $\mathcal{R}$ is the output range of the algorithm for a single agent, which we will take to be $\{0, 1\} \times P$ (whether someone is chosen to participate in a trade, and a price for the trade). Informally speaking, differential privacy requires that a change in a single data entry should have little (distributional) effect on the output of the mechanism. In other words, for every pair of data sets $D, D' \in \mathcal{D}^n$ that differ in at most one entry, differential privacy requires that the distribution of $\mathcal{M}(D)$ and $\mathcal{M}(D')$ are "close" to each other where closeness is measured by the privacy parameters ε and δ.

Definition 2. *Let $D, D' \in \mathcal{D}^n$ be two data sets of size n. We say D and D' are neighboring and write $D \sim D'$ if they differ in at most one data entry. D and D' are called i-neighbors (denoted by $D \sim_i D'$) if $D_{-i} = D'_{-i}$.*

Definition 3 ((Standard) Differential Privacy (DP) [7]). *An algorithm $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}^n$ is (ε, δ)-differentially private if for every pair of neighboring data sets $D \sim D' \in \mathcal{D}^n$, and for every subset of outputs $S \subseteq \mathcal{R}^n$,*

$$\Pr\left[\mathcal{M}(D) \in S\right] \leq e^{\varepsilon} \cdot \Pr\left[\mathcal{M}(D') \in S\right] + \delta$$

where the probability is taken with respect to the randomness of $\mathcal{M}$. if $\delta = 0$, $\mathcal{M}$ is said to be ε-DP.

We now define *joint differential privacy*. Joint differential privacy is defined in settings in which not only the inputs but also the outputs of the mechanism can be partitioned amongst the n users of the mechanism. In our setting, as in many mechanism design settings, this is the case: users report their valuations (which constitute the data) and then each receives an individual allocation. Joint differential privacy requires that an individual's input to the mechanism has little (distributional) effect on the outputs given to *others* — but allows one's own input to have a large effect on one's own output. Informally, it protects the privacy of each individual from arbitrary coalitions of other individuals using the system.

Definition 4 (Joint Differential Privacy [16]). *An algorithm $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}^n$ is (ε, δ)-joint differentially private if for every i, for every pair of i-neighbors $D \sim_i D' \in \mathcal{D}^n$, and for every $S \subseteq \mathcal{R}^{n-1}$,*

$$\Pr\left[\mathcal{M}(D)_{-i} \in S\right] \leq e^{\varepsilon} \cdot \Pr\left[\mathcal{M}(D')_{-i} \in S\right] + \delta$$

where the probability is taken with respect to $\mathcal{M}$'s randomness. If $\delta = 0$, $\mathcal{M}$ is said to be ε-joint DP.

We will use the *Laplace* and *exponential mechanisms* of differential privacy in our proposed algorithms. See Appendix A for their formal definitions, their privacy and accuracy guarantees, and a few properties of differential privacy including *post-processing* and *composition*.

[4] i.e., we only allocate agents willing to trade at price p. Agents not willing to participate at price p will opt out from the trade, thus are not taken into account by our benchmark.

[5] $\mathbf{1}[A]$, here and throughout the paper, represents the indicator function of event A.

3 PRIVATE CALL AUCTION MECHANISMS

In this section, we outline our jointly differentially private mechanisms for the call auction problem and analyze their performance guarantees. Each mechanism's performance is measured in terms of its *payoff* — that is, the total number of shares cleared — as well as its *inventory* — the net position that the mechanism must itself take on. We measure our mechanisms' payoffs against the *maximal* number of shares that could be cleared with a uniform price, *given* the agents' reports.

Throughout this section, we assume reports are truthful; we will show in Section 4 that our mechanisms are approximately dominant strategy truthful. We also highlight that taking on *some* inventory is unavoidable – if the mechanism took no net position, a coalition of agents could use the constraint that the number of buyers and sellers must be equal to circumvent joint differential privacy – but our guarantees ensure that this net position remains small with high probability.

We propose three mechanisms. The first mechanism, described in Subsection 3.1, uses the exponential mechanism (see Appendix A) to select a clearing price and then uses binomial randomization to determine who participates in a trade. In Subsection 3.2, we provide a second mechanism that again uses the exponential mechanism to select a price, but uses lottery numbers which are assigned to agents ex-ante to determine market participants. In Subsection 3.3, we describe a meta-algorithm that privately picks the mechanism with the better performance guarantee,[6] and achieves performance as good as that of the best of the first two mechanisms.

Finally, in Subsection 3.4, we show matching lower bounds (up to log factors) for the payoff of *any* (ε, δ)-joint differentially private mechanism for the call auction.

3.1 A Private Call Auction Mechanism via Coin Flipping

In this subsection, we introduce our first jointly differentially private algorithm for selecting a price and allocating buyers and sellers to trades. The algorithm uses the exponential mechanism to differentially privately select a clearing price. With a slight abuse of notation, let

$$\Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right) = \min\left\{\sum_{i\in\mathcal{S}} \mathbf{1}\left[\mathbf{v}_i^s \le p\right], \sum_{j\in\mathcal{B}} \mathbf{1}\left[\mathbf{v}_j^b \ge p\right]\right\} \tag{2}$$

be the number of trades that can happen at price p while guaranteeing every agent non-negative utility; $\Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right)$ is the utility function used by the exponential mechanism. After choosing the price, the mechanism randomly selects buyers and sellers willing to transact at the chosen price by flipping a coin with some particular bias for every agent in the market. The exchange then transacts with all selected transactors, possibly taking a net position in the asset. We formalize this mechanism in Algorithm 1. The mechanism takes data set $(\mathbf{v}^s, \mathbf{v}^b)$, privacy parameter ε, and confidence parameter α as inputs and outputs a price p and allocation vectors $\mathbf{a} = (\mathbf{a}^s, \mathbf{a}^b)$. In the algorithm description, $\exp(\cdot)$ is the exponential function, $Lap(\sigma)$ represents a mean-zero Laplace random variable with scale parameter σ, $(x)_+ := \max(x, 0)$, and $Bern(q)$ represents a Bernoulli random variable with success probability q.

We start the analysis by providing the privacy guarantees obtained by Algorithm 1:

Claim 1. *The allocation mechanism described in Algorithm 1 satisfies* 3ε *joint differential privacy.*

The full proof of this claim can be found in Appendix C. We also provide bounds on the payoff and the inventory of Mechanism 1 below:

[6] Which guarantee is best depends on the specific instance at hand.

ALGORITHM 1: Private Call Auction with Allocation via Coin Flipping ($\mathcal{M}_1$)

Input: Agents' valuations $(\mathbf{v}^s, \mathbf{v}^b)$, privacy level ε, confidence level α.
Output: Market price p, allocations $\mathbf{a} = (\mathbf{a}^s, \mathbf{a}^b)$.

Draw $p \propto \exp\left(\frac{\varepsilon \Pi(p, \mathbf{v}^s, \mathbf{v}^b)}{2}\right)$ ▷ Exponential mechanism chooses a price p privately

$\widehat{s} \leftarrow \sum_{i \in \mathcal{S}} \mathbf{1}\left[p \geq \mathbf{v}_i^s\right] + Lap(\frac{1}{\varepsilon})$ ▷ Privately estimate # of sellers willing to trade at p

$\widehat{b} \leftarrow \sum_{j \in \mathcal{B}} \mathbf{1}\left[p \leq \mathbf{v}_j^b\right] + Lap(\frac{1}{\varepsilon})$ ▷ Privately estimate # of buyers willing to trade at p

$$\mathbf{a}_i^s \leftarrow \mathbf{1}\left[p \geq \mathbf{v}_i^s\right] \cdot Bern\left(q^s = \min\left\{1, \frac{\left(\widehat{b}\right)_+}{\left(\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right\}\right) \text{ for all } i \in \mathcal{S}. \quad \triangleright \text{ Sellers' allocations}$$

$$\mathbf{a}_j^b \leftarrow \mathbf{1}\left[p \leq \mathbf{v}_j^b\right] \cdot Bern\left(q^b = \min\left\{1, \frac{\left(\widehat{s}\right)_+}{\left(\widehat{b} - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right\}\right) \text{ for all } j \in \mathcal{B}. \quad \triangleright \text{ Buyers' allocations}$$

Theorem 1 (Payoff and Inventory of Mechanism 1). *Suppose* $OPT \geq 5\ln(V/\alpha)/\varepsilon$.

(1) *Payoff: with probability* $1 - 8\alpha$,

$$\Pi(\mathcal{M}_1) \geq OPT - \frac{2\ln(V/\alpha)}{\varepsilon} - \frac{2\ln(1/\alpha)}{\varepsilon} - \sqrt{6\left(OPT + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}$$

(2) *Inventory: with probability* $1 - 6\alpha$,

$$I(\mathcal{M}_1) \leq \frac{18\ln(1/\alpha)}{\varepsilon} + 2\sqrt{6\left(OPT + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(2/\alpha)} + \frac{4\ln(2/\alpha)}{3}$$

Remark 1. *Note that we constraint* $OPT = \Omega(\ln(V/\alpha)/\varepsilon)$. *When* $OPT = O(\ln(V/\alpha)/\varepsilon)$, *the inaccuracy introduced by releasing a differentially private price via the exponential mechanism is on the order of* $OPT = \Omega(\ln(V/\alpha)/\varepsilon)$, *and we cannot hope to recover non-trivial utility guarantees.*

The proof of Theorem 1 is given in Appendix D.1. We note that our bound does not follow directly from the classical guarantees of the Laplace and exponential mechanisms; it requires a more involved analysis of the concentration of the distribution of buyers and sellers selected to trade in Algorithm 1.

3.2 A Private Call Auction Mechanism via Lottery Numbers

Here, we present a second mechanism that, rather than using independent randomization to decide who participates in a trade, uses correlated randomization to improve the payoff and reduce the inventory requirements of the mechanism. In the second mechanism, participants are given data-independent "lottery numbers", and thresholds on these lottery numbers (selected using the exponential mechanism) are used to select among willing traders on both sides of the market. This correlation allows us to remove the $\sqrt{\text{OPT}}$ term in the bounds of the previous mechanism, at the cost of introducing a logarithmic dependence on the number of agents n.

For a given valuation profile $(\mathbf{v}^s, \mathbf{v}^b)$, let $\Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right)$ be defined as in Equation 2. Assume seller i is assigned a lottery number $l_i^s \in [n^s]$ and buyer j is given $l_j^b \in [n^b]$ where we require that these lottery numbers are different for different agents. Without loss of generality, we assume $l_i^s = i$ and $l_j^b = j$. For a given price p, and profiles $(\mathbf{v}^s, \mathbf{v}^b)$, the loss of thresholds τ^s and τ^b on lottery numbers

(one for sellers and one for buyers) is expressed as follows:

$$L^s\left(\tau^s, p, \mathbf{v}^s, \mathbf{v}^b\right) = \left|\sum_{i\in\mathcal{S}} \mathbf{1}\left[p \geq \mathbf{v}_i^s, \tau^s \geq i\right] - \Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right)\right|,$$

$$L^b\left(\tau^b, p, \mathbf{v}^b, \mathbf{v}^b\right) = \left|\sum_{j\in\mathcal{B}} \mathbf{1}\left[p \leq \mathbf{v}_j^b, \tau^b \leq j\right] - \Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right)\right|$$

For a price p, these loss functions measure how far off the number of agents chosen to trade on each side of the market would be from our target number of trades, $\Pi(p, \mathbf{v}^s, \mathbf{v}^b)$, if we used thresholds τ^s and τ^b as a tie-breaking rule to select sellers and buyers who are willing to trade at price p, respectively. In Algorithm 2, just as before, we first use the exponential mechanism to select a price and then use the exponential mechanism with loss functions L^s and L^b (or utility functions: $-L^s$ and $-L^b$, based on the terminology used to describe the exponential mechanism in Appendix A) to select the thresholds on lottery numbers.

ALGORITHM 2: Private Call Auction with Allocation via Lottery Numbers ($\mathcal{M}_2$)

Input: Agents' valuations $(\mathbf{v}^s, \mathbf{v}^b)$, privacy level ε.
Output: Market price p, allocations $\mathbf{a} = (\mathbf{a}^s, \mathbf{a}^b)$.

Draw $p \propto \exp\left(\frac{\varepsilon\Pi(p,\mathbf{v}^s,\mathbf{v}^b)}{2}\right)$ ▷ Exponential mechanism to privately choose a price p

Draw $\tau^s \propto \exp\left(-\frac{\varepsilon L^s(\tau^s,p,\mathbf{v}^s,\mathbf{v}^b)}{4}\right)$ ▷ Exponential mechanism to privately choose τ^s

Draw $\tau^b \propto \exp\left(-\frac{\varepsilon L^b(\tau^b,p,\mathbf{v}^s,\mathbf{v}^b)}{4}\right)$ ▷ Exponential mechanism to privately choose τ^b

$\mathbf{a}_i^s \leftarrow \mathbf{1}\left[p \geq \mathbf{v}_i^s, \tau^s \geq i\right]$ for all $i \in \mathcal{S}$. ▷ Sellers' allocations

$\mathbf{a}_j^b \leftarrow \mathbf{1}\left[p \leq \mathbf{v}_j^b, \tau^b \leq j\right]$ for all $j \in \mathcal{B}$. ▷ Buyers' allocations

Claim 2. *The allocation mechanism described in Algorithm 2 satisfies* 3ε *joint differential privacy.*

Theorem 2 (Payoff and Inventory of Mechanism 2). *For any* $\alpha > 0$,

(1) *Payoff: with probability* $1 - 3\alpha$,

$$\Pi(\mathcal{M}_2) \geq OPT - \frac{2\ln(V/\alpha)}{\varepsilon} - \frac{4\ln(n/\alpha)}{\varepsilon}$$

(2) *Inventory: with probability* $1 - 2\alpha$,

$$I(\mathcal{M}_2) \leq \frac{8\ln(n/\alpha)}{\varepsilon},$$

The proof of Claim 2 is provided in Appendix C, and that of Theorem 2 in Appendix D.2.

3.3 A Meta Algorithm: Selecting the Best Mechanism Privately

Notice that the first term in the payoff bounds of both Theorems 1 and 2 are identical (as they both correspond to choosing a price using the exponential mechanism) but the remaining terms differ ($\mathcal{M}_1$ relies on binomial coin flips for tie-breaking whereas $\mathcal{M}_2$ tie-breaks via thresholds on lottery numbers). These two bounds are in general not comparable, as one depends on the maximum number of shares OPT that can be cleared, whereas the other one depends on the total number n of agents in the market. The first bound provides better guarantees (up to constants and $\ln(1/\alpha)$ terms) when $\sqrt{\text{OPT}} < \ln(n)/\varepsilon$, i.e. when the number of possible trades is significantly smaller than

the total number of agents in the market,[7] whereas the second bound provides better guarantees when $\sqrt{\text{OPT}} > \ln(n)/\varepsilon$.

We can achieve the better of the two bounds by comparing the bounds of Theorems 1 and 2 in a differentially-private manner and then running the mechanism with the better bound according to this private computation. Details of this meta-algorithm and its guarantees are given in Appendix B.

3.4 A Lower Bound

We now provide a lower bound showing that *any* algorithm which computes a price in an (ε, δ)-differentially private manner and allocates among willing participants at this price *must*, for *some* instance, suffer a loss of $\Omega(1/\varepsilon)$ (compared to the optimal number of shares that could be cleared on that instance). Because this bound applies to a broader set of mechanisms that reveal *only* the price privately (but may select the optimal allocation absent privacy), it also applies to the mechanisms considered in Section 3. We will compare the performance of any given differentially private algorithm on several input data sets. To do so, we will define an instance-dependent benchmark below, that we call OPT(D). Formally, given an input data set $D = (\mathbf{v}^s, \mathbf{v}^b)$, our benchmark is:

$$\text{OPT}(D) = \max_p \min\left\{\sum_{i\in\mathcal{S}} \mathbf{1}\left[\mathbf{v}_i^s \le p\right], \sum_{j\in S} \mathbf{1}\left[\mathbf{v}_j^b \ge p\right]\right\}.$$

Definition 5 (Loss of an algorithm). *For any (possibly randomized) algorithm $\mathcal{A} : \mathcal{D}^n \to P$ that takes a data set $D = (\mathbf{v}^s, \mathbf{v}^b)$ as an input and outputs a price p, the loss of $\mathcal{A}$ on input data set $D = (\mathbf{v}^s, \mathbf{v}^b)$ of agents valuations is defined as follows:*

$$L(\mathcal{A}, D) = OPT(D) - \mathbb{E}_{p\sim\mathcal{A}(D)}\left[\min\left\{\sum_{i\in\mathcal{S}} \mathbf{1}\left[\mathbf{v}_i^s \le p\right], \sum_{j\in S} \mathbf{1}\left[\mathbf{v}_j^b \ge p\right]\right\}\right].$$

I.e., this loss compares the number of trades that could be cleared in expectation at the price selected by $\mathcal{A}$ to the maximum number of trades when the trading price is optimally chosen. We define the worst-case expected loss of $\mathcal{A}$ as the worst-case loss over all data sets, i.e. $L(\mathcal{A}) = \sup_D [L(\mathcal{A}, D)]$.

Our lower bound will hold so long as δ is not too large in comparison with ε.[8] We note that our lower bound on the expected loss matches the $\widetilde{O}(1/\varepsilon)$ dependencies[9] of our high probability upper bounds on the loss for Mechanisms 1, 2, 5 (and consequently of any upper bound on the expected loss of these mechanisms). Finally, it is worth remarking that our lower bound for (ε, δ)-DP mechanisms matches the upper bound obtained by restricting attention to $(\varepsilon, 0)$-DP mechanisms; this implies that relaxing δ-privacy requirements of Mechanisms 1, 2, 5 will not lead to any significant improvements in terms of their accuracy guarantees.

Theorem 3. *[Lower bound on the loss of private algorithms] Pick any ε, δ such that $0 \le \varepsilon \le 1$ and $\delta = O(\varepsilon)$. There exists a range of (integer) valuations $P(\varepsilon)$ and a number of agents $n(\varepsilon)$ such that* any *(ε, δ)-DP algorithm $\mathcal{A} : \mathcal{D}^{n(\varepsilon)} \to P(\varepsilon)$ must suffer worst-case expected loss of $\Omega(1/\varepsilon)$.*

The proof of Theorem 3 relies on constructing a family of data sets $\{D_l\}_l$ such that no differentially private algorithm $\mathcal{A}$ can simultaneously suffer expected loss of $O(1/\varepsilon)$ on all of them. We do so by carefully calibrating the following trade-off: on the one hand, we require any pair of data sets in

[7] This models practical situations in repeated financial markets where sellers price a security higher than most buyers are willing to pay. In such situations, buyers may elect to wait until a new seller comes and offers a better price, while sellers may wait for a new buyer willing to buy at the current price.

[8] Typically, differentially private algorithms use $\delta << \varepsilon$.

[9] The instances we construct use $V, n \sim 1/\varepsilon$. The logarithmic dependencies of our upper bounds in n and V translate into logarithmic dependencies in $1/\varepsilon$, hence the $\widetilde{O}$ notation.

$\{D_l\}_l$ be close enough that the stability properties of differential privacy guarantee any private algorithm must pick a similar distribution of prices on both data sets. On the other hand, we require that the data sets furthest from each other are different *enough* such that no fixed distribution can incur a low loss on both.

3.5 Connections to the Market Impact Literature

As mentioned in the Introduction, it is possible to draw some informal but interesting connections between this work and the finance literature on market impact. Market impact models typically propose strong stochastic assumptions on price formation (e.g. random walk and diffusion models or martingale assumptions on limit order dynamics) and then solve for the optimal strategy to minimize trading costs and price impact. In particular, there is a large body of work on the so-called "square root law" (see, eg. [11, 12]). which predicts that the change to price inflicted by a trade of k shares scales with $\sqrt{k/\mathcal{V}}$, where $\mathcal{V}$ is the total volume of shares cleared during the trade; the ratio $k/\mathcal{V}$ is referred to as the trade's *participation rate*. As we note below, $\mathcal{V}$ is typically closely related to other measures of market activity such as the number of orders placed (as with our n) or the number of quote changes in limit order dynamics.

Our results imply that the change in the expected clearing price in our DP call auction resulting from an order of k shares is bounded by a multiplicative factor of $(e^{k\varepsilon} - 1)$. Setting this equal to $\sqrt{k/n}$ to match the square root law[10] and solving for ε approximately yields $\varepsilon \approx 1/\sqrt{kn}$ for small participation rates. Plugging this into our utility bound of Theorem 2, the shares we execute at this ε scales like $\mathrm{OPT}(1 - \sqrt{kn}/\mathrm{OPT})$. Thus as long as k is $o(n)$ and OPT scales with n,[11] asymptotically we approach OPT with the same price impact as that predicted by the square root law but with two major advantages. First, we have made *no* assumptions, stochastic or otherwise, on the orders placed by market participants. Second, we are not only bounding the price impact, we are also bounding information leakage of *any* form, as per the promises of DP.

4 STRATEGIC FRAMEWORK

In Section 3, we focused on the algorithmic form of our mechanism and provided privacy guarantees and optimality guarantees with respect to the reported valuations, without regard to whether those reports are truthful or not. In this section, we embed our mechanisms into a game theoretic framework and examine its properties, including (approximate) truthfulness. More precisely, we now assume the agents are strategic; they may decide to report a bid that differs from their valuation, or even to not participate in the mechanism in the first place. Formally, all sellers i and buyers j have quasi-linear utilities determined by their own valuations and the outcome of the mechanism: $\mathbf{u}_i^s(\mathcal{M}) = \mathbf{a}_i^s \cdot (p - \mathbf{v}_i^s)$, $\mathbf{u}_j^b(\mathcal{M}) = \mathbf{a}_i^b \cdot (\mathbf{v}_j^b - p)$, where, with a slight abuse of notation, we omit the dependency of $\mathcal{M}$ on the agents' reports.

Buyers and sellers aim to maximize their utility from participating (or not participating) in the mechanism. In the face of strategic behavior, we will require our mechanisms to be (approximately) truthful and individually rational; i.e., it should never be in an agent's best interest to misreport his valuation, and an agent should always have a strategy that guarantees non-negative utility from participating in the mechanism and so would rather participate than not. Individual rationality and (approximate) truthfulness are formally defined below:

Definition 6 (Ex-Post Individual Rationality). *We say a double-auction mechanism $\mathcal{M}$ satisfies ex-post individual rationality if, for every seller $i \in \mathcal{S}$, there exists a bid $\mathbf{r}_i^s$ for agent i such that*

[10] Here we are assuming that the number of orders n in our model plays the role of $\mathcal{V}$ above; see subsequent footnote.

[11] This scaling is broadly consistent with recent data from electronic exchanges. For instance, the ratio of shares traded to quote changes (a common measure of market activity) across 3443 U.S. equities averaged 0.16 with standard deviation 0.09.

for every possible set of bids $\mathbf{r}_{-i}$ submitted by all agents but i, and every realization of the randomness of the mechanism $\mathcal{M}$, $\mathbf{u}_i^s\left(\mathcal{M}\left(\mathbf{r}_i^s, \mathbf{r}_{-i}\right)\right) \geq 0$, and similarly for every buyer j, there exists a bid $\mathbf{r}_j^b$ for agent j such that for every possible set of bids $\mathbf{r}_{-j}$ submitted by all agents but j, and every realization of the randomness of the mechanism $\mathcal{M}$, $\mathbf{u}_j^b\left(\mathcal{M}\left(\mathbf{r}_j^b, \mathbf{r}_{-j}\right)\right) \geq 0$.

Definition 7 (Approximate Dominant-Strategy Truthfulness). *We say a double-auction mechanism $\mathcal{M}$ satisfies γ-approximate dominant-strategy truthfulness if, for every seller $i \in \mathcal{S}$, every possible bid $\mathbf{r}_i^s$ submitted by i, and every possible set of bids $\mathbf{r}_{-i}$ submitted by all agents but i,*

$$\mathbb{E}_{\mathcal{M}}\left[\mathbf{u}_i^s\left(\mathcal{M}\left(\mathbf{r}_i^s, \mathbf{r}_{-i}\right)\right)\right] \leq \mathbb{E}_{\mathcal{M}}\left[\mathbf{u}_i^s\left(\mathcal{M}\left(\mathbf{v}_i^s, \mathbf{r}_{-i}\right)\right)\right] + \gamma$$

and similarly for every buyer j, for every possible bid $\mathbf{r}_j^b$ submitted by j and every possible set of bids $\mathbf{r}_{-j}$ submitted by all agents but j,

$$\mathbb{E}_{\mathcal{M}}\left[\mathbf{u}_j^b\left(\mathcal{M}\left(\mathbf{r}_j^b, \mathbf{r}_{-j}\right)\right)\right] \leq \mathbb{E}_{\mathcal{M}}\left[\mathbf{u}_j^b\left(\mathcal{M}\left(\mathbf{v}_j^b, \mathbf{r}_{-j}\right)\right)\right] + \gamma$$

where expectations are taken with respect to the randomness of $\mathcal{M}$.

In Section 4.1, we show that our mechanisms are *individual rational* and (unlike in the standard call auction) approximately *dominant-strategy truthful*. While our results assume that agents wish to trade a single share, we show how our per-share guarantees translate into (gracefully degrading) per-player guarantees in more general setting in which agents can trade multiple shares.

Then, in Sections 4.2.1-4.2.2, we consider *learning dynamics* under both the standard call auction and our mechanism and show that a system in which agents use a modified exponential weights algorithm (which we call "Social" Exponential Weights) to learn to bid will eventually converge to the optimal number of shares cleared. While it is true that truthfulness implies that agents cannot do better than bidding their true values, one might consider learning dynamics for two reasons. First, if agents do not trust the mechanism designer (or share their assumptions), applying a no-regret learning algorithm is a plausible response to guarantee good performance. Second, good outcomes obtained in the presence of decentralized, distributed, and selfish algorithms are compelling evidence of the robustness and quality of our mechanism. To our knowledge, the use of no-regret learning algorithms by all agents in a call auction setting has not been studied before, and these results may be of independent interest.

4.1 Individual Rationality and Truthfulness Properties of Our Algorithms

In this section, we discuss the incentive properties of our proposed algorithms. To do so, we note that $(\mathbf{v}^s, \mathbf{v}^b)$ are the true valuations of sellers and buyers and denote their revealed bids by $(\mathbf{r}^s, \mathbf{r}^b)$, respectively. To study truthfulness, we assume seller i (buyer j) can submit a bid $\mathbf{r}_i^s$ ($\mathbf{r}_j^b$) that may not be equal to their valuation $\mathbf{v}_i^s$ ($\mathbf{v}_j^b$), and show that it is approximately never in agent i's (resp. j's) best interest to do so. We start by noting that our mechanisms are individually rational:

Claim 3 (Individual Rationality). *The mechanisms described in Algorithms 1, 2, and the Meta Algorithm 5 (described in Appendix B) are ex-post individually rational.*

Proof. We prove the result for Mechanism 1; proofs for the other mechanisms are similar. It suffices to show that there exists a strategy for any seller (resp. any buyer) that guarantees him non-negative utility. For any seller i, setting $\mathbf{r}_i^s = \mathbf{v}_i^s$ is a strategy that ensures that whenever i is allocated a trade (i.e. $\mathbf{a}_i^s = 1$), it must be that $p \geq \mathbf{r}_i^s = \mathbf{v}_i^s$; this immediately guarantees i gets non-negative utility—independently of how other agents bid and of the randomness of the mechanism. A similar proof holds for buyers. □

We also show that differential privacy guarantees approximate truthfulness in the dominant-strategy sense: i.e., it does not allow agents (sellers and buyers) to gain too much profit by submitting a bid different than their true valuation, *no matter what the realized bids of the other agents are*[12].

Claim 4 (Approximate Truthfulness). *The mechanisms described in Algorithms 1 and 2 satisfy γ-approximate dominant-strategy truthfulness for $\gamma = (e^{3\varepsilon} - 1)V$; the mechanism described in Algorithm 5 (in Appendix B) satisfies γ-approximate dominant-strategy truthfulness for $\gamma = (e^{7\varepsilon} - 1)V$.*

We defer the full proof to Appendix F. In the proof, we first observe that since the market price is chosen subject to differential privacy, individual agents cannot significantly change it by misreporting their valuations. However, this is not enough to argue truthfulness, as under *joint* differential privacy, an agent's allocation may heavily depend on his report. To complete the proof, we show that the function by which the mechanism determines transactors is a best-response for an agent with the reported valuation given the output of the differentially private mechanism.

Note that, in general, call auctions are *not* dominant-strategy truthful, since even small bidders may impact the price selected by a mechanism acting on reported bids. This is a consequence of the fact that in the simple call auction (as well as in continuous order book mechanisms) the optimal price is, in general, *not* stable [10]. Importantly, we note that the truthfulness guarantees are a function of ε; as OPT grows larger, ε can be made smaller with less and less relative cost. Consequently, the truthfulness guarantee can be made stronger for a given level of privacy as the number of optimal trades cleared increases.

We highlight that because our strategic framework assumes each bidder controls a single share, our guarantees are at the *per-share* level. Our privacy guarantees generalize, however, to the case where bidders control at most k shares by expanding ε by a factor of k.[13] Our truthfulness guarantees also follow by expanding ε by a factor of k.

4.2 Learning in Repeated Call Auctions

In this section, we consider a *repeated* call auction. Agents are initially unaware of each other's valuations and behavior and run simple learning algorithms to learn how to bid. In each time step t, each seller i (respectively buyer j) reports a bid $\mathbf{r}^s_{i,t}$ (resp. $\mathbf{r}^b_{j,t}$), which may differ from his valuation, to the mechanism. Given this input $(\mathbf{r}^s_t, \mathbf{r}^b_t)$, the mechanism computes and publicly releases a price p_t and assigns an allocation $\mathbf{a}_{i,t}$ to each seller i (respectively $\mathbf{a}_{j,t}$ to each buyer j). We will consider two versions of this mechanism, one that is non-private and is inspired by standard call auctions, in Section 4.2.1, and one that is private and is based on Mechanism 1, in Section 4.2.2. The agents then update their bidding strategies based on the quantities outputted by the mechanism, via a simple no-regret algorithm (Exponential Weights).

We highlight that our agents are *naive* in that they do not compute a counterfactual price p_t and allocation vector $\mathbf{a}_t$ given alternative bids they could have made. Instead, they only update their bidding strategies with respect to how much better off they could have been by bidding differently, *assuming they had no effect on the price*. The motivation for this is two-fold: first, counterfactual reasoning would require the agents to know the bids of other agents, which are not released by the mechanism (and typically not available in many real-life call auctions). Second, when agents are small relative to the total market, they may believe that their actions *do not* greatly affect these

[12] Truthfulness is desirable not only because it makes computing equilibrium strategies and predicting equilibrium behavior simpler, but also because knowing the *true* valuations allows the mechanism designer to clear the most shares.

[13] An ε-differentially private mechanism with respect to a single share is $k\varepsilon$-differentially private with respect to the data of a bidder who controls k shares; intuitively, this is because an agent that misreports his valuation over k shares creates a dataset that is a k-neighbor of the dataset in which they had bid truthfully. Such a bidder can affect the distribution of prices by an amount of at most $e^{k\varepsilon}$, and so their own expected utility by $(e^{k\varepsilon} - 1)V$.

quantities. We note that differential privacy makes this belief into a *property* of our mechanism rather than a naive assumption. Thus, small bidders using naive updates will have a *real* regret guarantee when interacting with a differentially private call auction.

4.2.1 Learning in the Absence of Privacy. In this section, we focus on learning dynamics when the mechanism runs a standard call auction, absent privacy; this non-private setting will serve as a natural point of comparison for dynamics with respect to our private mechanism. At every time step t, agents submit bids that may differ from their valuations. In response, the mechanism computes a price and allocation, with the goal of maximizing traded shares among willing participants. We denote the agents' reports as $\mathbf{r}^s_{i,t}$ and $\mathbf{r}^b_{j,t}$ for seller i and buyer j, respectively, at time t. The mechanism chooses a price p_t to maximize

$$\Pi\left(p, \mathbf{r}^s_t, \mathbf{r}^b_t\right) \triangleq \min\left\{\sum_{i\in\mathcal{S}} \mathbf{1}\left[\mathbf{r}^s_{i,t} \le p\right], \sum_{j\in\mathcal{B}} \mathbf{1}\left[\mathbf{r}^b_{j,t} \ge p\right]\right\},$$

which is the number of shares the mechanism will trade at price p, assuming that sellers will only agree to trade when the price is above their reported bid and buyers when the price is below their reported bid. To compute the allocation $\mathbf{a}_t = (\mathbf{a}^s_t, \mathbf{a}^b_t)$, the mechanism must choose among these sellers and buyers it believes (based on the reports) are willing to trade at the chosen price p_t. When there are an equal number of sellers and buyers willing to trade at price p_t, the mechanism allocates a trade to all of them; otherwise, the mechanism randomly selects a subset of $\Pi(p, \mathbf{r}^s_t, \mathbf{r}^b_t)$ agents from the side with excess number of willing participants. Formally, the mechanism computes $q^b_t = \Pi(p_t, \mathbf{r}^s_t, \mathbf{r}^b_t)/\sum_{j\in\mathcal{B}} \mathbf{1}[\mathbf{r}^b_{j,t} \ge p_t]$ and $q^s_t = \Pi(p_t, \mathbf{r}^s_t, \mathbf{r}^b_t)/\sum_{i\in\mathcal{S}} \mathbf{1}[\mathbf{r}^s_{i,t} \le p_t]$; these probabilities will be less than 1 on the excess side of the market and exactly 1 on the short side. We assume the mechanism publicly releases p_t, q^s_t, and q^b_t to all agents in the market, and communicates to each seller i (resp. buyer j) his own allocation $\mathbf{a}^s_{i,t}$ (resp. $\mathbf{a}^b_{j,t}$).

Agents learn via Exponential Weights: A natural no-regret (*regret* here is the classic notion of performance in online learning) algorithm for updating bidding strategies is the Exponential Weights mechanism. We describe the classical *Exponential Weights Update* rule for buyers (buyer j) in Algorithm 3, and note that this update is defined symmetrically for the sellers.

ALGORITHM 3: Exponential Weights

Input: Learning rate η.

Set $\mathbf{w}^b_{j,1}(k) \leftarrow \frac{1}{\mathbf{v}^b_j}$ for $k = 1, \ldots, \mathbf{v}^b_j$ ▷ Initialize uniform weights.

for $t \in 1 \ldots T$ **do**

$\mathbf{r}^b_{j,t} \sim \mathbf{w}^b_{j,t}$ ▷ Draw bid from distribution

$\mu^b_{j,t}(k) \leftarrow q^b_t(\mathbf{v}^b_j - p_t)\mathbf{1}[k \ge p_t]$ for $k = 1, \ldots, \mathbf{v}^b_j$ ▷ Observe payoff of each bid k

$\mathbf{w}^b_{j,t+1}(k) \leftarrow \frac{\exp(\eta\mu^b_{j,t}(k))}{\sum_j \mathbf{w}^b_{j,t}(j)\exp(\eta\mu^b_{j,t}(j))} \cdot \mathbf{w}^b_{j,t}(k)$ for $k = 1, \ldots, \mathbf{v}^b_j$ ▷ Update the weights.

end

Informally, the updates work as follows. Initially, we assume every seller bids uniformly above their value and every buyer bids uniformly below their value.[14] Then, in each round t, for every

[14] A buyer j cannot improve his utility by bidding over his valuation (as increasing his bid cannot decrease p_t nor increase his probability of allocation), and risks obtaining negative utility by doing so, if $\mathbf{v}^b_{j,t} < p_t \le \mathbf{r}^b_{j,t}$. Hence, bidding above his valuation is a dominated strategy for the buyer. Similarly, bidding under his value is a dominated strategy for a seller. The assumption of this prior knowledge can be relaxed at the price of slower convergence.

possible $k \in P$, agents compute what their expected payoff would have been had they reported k as their valuation, given the *current* price p_t and the allocation probabilities q_t^s and q_t^b. They use these expected payoffs to update their distribution of bids, in a way that puts exponentially more weight on bids with higher expected utilities; the speed at which these updates happen is controlled by the learning rate parameter η, taken here to be constant. For appropriate choices of learning rate η, this algorithm is known to be no-regret.

One may hope these dynamics converge to clearing OPT shares with probability going to 1, where OPT is defined as in Equation 1. However, this may not be the case when agents update their weights according to Algorithm 3. This stems from the fact that agents are indifferent between trading at their valuation, and not trading at all, as both net a payoff of zero. This is reflected in the exponential weight update, and buyers learn to put a significant amount of weight on bids that are strictly less than their valuation (as trades for those bids are strictly profitable). When clearing OPT trades requires many agents to bid exactly at their valuation, the number of shares cleared is bounded away from the benchmark. We show instead that the dynamics will clear the following benchmark, which only considers trades that are strictly profitable for both sides of the market:

Definition 8 (Optimal Jointly Profitable Trades). *We let OPT' be the maximum number of trades achievable for a given $(\mathbf{v}^s, \mathbf{v}^b)$, such that all trading buyers and sellers get strictly positive utility. Formally,*

$$OPT' = \max_p \min \left\{ \sum_{i \in \mathcal{S}} \mathbf{1}\left[\mathbf{v}_i^s < p\right], \sum_{j \in \mathcal{B}} \mathbf{1}\left[\mathbf{v}_j^b > p\right] \right\}.$$

We call this benchmark the "Optimal Jointly Profitable Trades with Uniform Pricing" benchmark.

The statement showing that the mechanism will converge in probability to clearing at least OPT′ shares is formalized below:

Theorem 4 (Convergence to (at least) OPT′). *Suppose buyers and sellers update their bid distributions according to Algorithm 3 (with any $\eta > 0$). Further, at any time t, suppose p_t is chosen uniformly at random among the set of optimal prices at time t. Then, the number of shares cleared at time t satisfies the following statement:* $\lim_{t\to\infty} \Pr\left[\Pi\left(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b\right) \geq OPT'\right] = 1$.

We also provide a variant of the Exponential Weights algorithm, that we will show converges to OPT shares cleared. This variant is described in Algorithm 4.

ALGORITHM 4: Social Exponential Weights

Input: Learning rate η, "fake" utility ξ.
$\mathbf{w}_{j,1}^b(k) \leftarrow 1/\mathbf{v}_j^b$ for $k = 1, \ldots, \mathbf{v}_j^b$ ▷ Initialize with uniform weights.
for $t = 1, \ldots, T$ **do**
  $\mathbf{r}_{j,t}^b \sim \mathbf{w}_{j,t}^b$ ▷ Draw from the current weights.
  if $\mathbf{v}_j^b \neq p_t$ **then**
    $\mu_{j,t}^b(k) \leftarrow q_t^b(\mathbf{v}_j^b - p_t)\mathbf{1}[k \geq p_t]$ for $k = 1, \ldots, \mathbf{v}_j^b$ ▷ expected utility for bid k.
  else
    $\mu_{j,t}^b(k) \leftarrow q_t^b \xi \mathbf{1}[k = \mathbf{v}_j^b]$ for $k = 1, \ldots, \mathbf{v}_j^b$ ▷ Agent pretends getting utility ξ from trading.
  end
  $\mathbf{w}_{j,t+1}^b(k) \leftarrow \frac{\exp(\eta \mu_{j,t}^b(k))}{\sum_l \mathbf{w}_{j,t}^b(l) \exp(\eta \mu_{j,t}^b(l))} \cdot \mathbf{w}_{j,t}^b(k)$ for $k = 1, \ldots, \mathbf{v}_j^b$ ▷ Update the weights.
end

Algorithm 4 is a modification of the classic Exponential Weights algorithm. In particular, when the price is equal to agent's valuation, the algorithm assigns a nonzero utility $q_t^b \xi$ to reporting the agent's valuation, for ξ arbitrarily small; this can be seen as agents updating their weights as if they strictly preferred trading to not trading, even when their trade would make no profit. In other words, it implements a preference to break ties (in utility) in favor of trading over not trading. We call this "Social" Exponential Weights because incorporating this modified utility allows the system as a whole to reach a better social outcome (one with more shares traded) than otherwise. Crucially, despite this modification, Algorithm 4 remains no-regret for a fixed horizon T with appropriate choices of learning rate, η, and "fake" utility, ξ:

Lemma 1 (No-regret). *Algorithm 4 is no-regret ($O\left(\sqrt{T}\right)$ cumulative regret) for $\eta, \xi = O\left(1/\sqrt{T}\right)$.*

The proof is almost identical to that of the no-regret guarantees of traditional exponential weights, and is deferred to Appendix G.1. We highlight that we define regret with respect to the single best action in hindsight given the *fixed* sequence of prices observed; that is, we do not consider the notion of *Stackelberg* regret, which is calculated with respect to the best fixed action *given* that the mechanism picks a sequence of prices in response to the selected actions (see, e.g. [6]). If agents are small enough that their actions do not greatly affect the mechanism's responses, then the standard notion of regret and Stackelberg regret do not greatly differ; if, moreover, a mechanism is differentially private, then (for small enough agents) these notions of regret coincide, because differential privacy ensures that agents placing small orders have little impact on the price.

Under Algorithm 4, the number of shares cleared converges to OPT with probability that tends to 1 as t grows large. We make this statement formally below:

Theorem 5 (Convergence to OPT). *Suppose buyers and sellers that update their bidding strategies according to Algorithm 4 (with any $\eta, \xi > 0$). Further, suppose p_t is chosen uniformly at random among the set of optimal prices at time t. Then, the number of shares cleared at time t satisfies the following statement:* $\lim_{t\to\infty} \Pr\left[\Pi\left(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b\right) = OPT\right] = 1$.

To prove this result, we show that with a small, constant probability (in t) in any given round, all agents bid their valuations. In such cases, the mechanism picks an optimal price, and at least OPT buyers (resp. sellers) increase their probability of bidding above (resp. below) this price. When the number of rounds goes to infinity, this event is repeated infinitely often for some optimal price $p^\star$, and OPT buyers (resp. sellers) bid above (resp. below) $p^\star$ with probability that tends to 1. The full proof is given in Appendix G.3. A similar argument is used to prove Theorem 4, in Appendix G.2.

4.2.2 Learning in Repeated Call Auctions with Differential Privacy. We now consider the same dynamic setting as before, with the difference that the centralized designer now computes the price p_t and the allocation $\mathbf{a}_t$ at time t in a joint-differentially private fashion. For simplicity of exposition, we pick the private mechanism used by the designer to be Mechanism 1, which picks a price via the exponential mechanism and picks agents to allocate from the smaller side of the market via binomial coin flips. We show that when agents play according to the exponential weights (resp. Social EW) algorithm, the dynamics converge to clearing at least OPT (resp. OPT′) shares minus inaccuracies introduced by privacy.

Theorem 6. *Suppose buyers and sellers update their bidding strategies according to Algorithm 4 (with any $\eta, \xi > 0$). Further, suppose the market allocation mechanism is Algorithm 1. There exists an*

integer $N(\alpha)$ such that for any $t \geq N(\alpha)$, the number of shares cleared at time t satisfies

$$\Pr\left[\Pi\left(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b\right) \geq OPT - \frac{2\ln(V/\alpha)}{\varepsilon} - \frac{2\ln(1/\alpha)}{\varepsilon} - \sqrt{6\left(OPT + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}\right] \geq 1 - 9\alpha.$$

where this probability is taken with respect to the randomness of both Algorithms 1 and 4.

The proof idea is the following: despite the price randomness due to privacy, the event in which all agents bid their value and an optimal price is picked happens infinitely often, as in the non-private case. In turn, (at least) *OPT* buyers (resp. sellers) eventually learn to bid above (resp. below) an optimal price $p^\star$. However, the mechanism will still pick sub-optimal prices to guarantee privacy, as per the bound of Theorem 1. We refer the reader to Appendix G.4 for a complete proof.

A similar statement holds, with respect to benchmark OPT′ (see Definition 8), when agents update according to Algorithm 3.

Theorem 7. *Suppose buyers and sellers use the Exponential Weights Algorithm 3 (with any $\eta > 0$) to update their bids. Further, suppose the market allocation mechanism is Algorithm 1. There exists an integer $N(\alpha) > 0$ such that for all $t \geq N(\alpha)$, the number of shares cleared at time t satisfies*

$$\Pr\left[\Pi\left(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b\right) \geq OPT' - \frac{2\ln(V/\alpha)}{\varepsilon} - \frac{2\ln(1/\alpha)}{\varepsilon} - \sqrt{6\left(OPT' + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}\right] \geq 1 - 9\alpha.$$

5 SIMULATIONS

In previous sections, we designed our mechanism and obtained theoretical guarantees of performance; these guarantees were given both in a *one-shot* setting and relative to the optimal result that could be reached given agents' bids and also in a *repeated* setting using no-regret learning. In this section, we conduct experiments on simulated data in both a one-shot and learning setting in order to explore how tightly these guarantees bind in practice.

We perform all simulations in MATLAB using a similar starting configuration. We have 5000 buyers and 5000 sellers, and valuations must be integer values between 1 and 100. Valuations are drawn from normal distributions centered at 45 for sellers and 55 for buyers, with standard deviations of 15 for both. The draws are rounded to the nearest integer, and draws below 1 or above 100 are replaced with 1 and 100 respectively.

The mechanism we implement is the first we defined, namely, that described in Algorithm 1, run once or repeatedly for the one-shot game and learning settings, respectively. We vary ε over a range from $\varepsilon = 0.01$ to $\varepsilon = 0.5$.

Single-shot game. For the single shot game, we perform 800 trials per value of ε with a fixed set of agent valuations. These valuations were drawn randomly according to the procedure described above. We assume agents bid truthfully (and all of our comparisons are to the truthful optimal).

In the first plot of Figure 1, we show the empirical 5% quantile (i.e. the value for which only 5% of draws saw lower values) of the *competitive ratio* defined as the shares cleared as a fraction of OPT, the optimal number of shares that can be cleared given the realized valuations. This competitive ratio quantile is plotted in blue. The appropriate guarantee to compare to is the lower bound on this quantile given in Theorem 1, with confidence parameter $\alpha = 0.05/8$; this bound is plotted in orange. While the realized competitive ratio indicates, unsurprisingly, that privacy is not costless for very small levels of ε, it remains far above the worst-case guarantee predicted, and rapidly increases to nearly 1 in the practical regime (i.e., even for $\varepsilon = 0.1$). This shows that, for a large enough

number of agents and valuations drawn from well-behaved distributions, reasonable privacy can be achieved in practice with very little loss in utility.

The second plot of Figure 1 shows the absolute value of the inventory taken on by the mechanism, again plotting this quantile as a ratio of the optimal number of shares cleared in blue (again, limited to the top 95% of runs) and the theoretical upper bound for $\alpha = 0.05/6$ (as per the inventory bound of Theorem 1) in orange. Notice that for very small ε, the theoretical guarantee can be extremely large; yet, again, the realized inventory is far below the guarantee and never exceeds 23% for even $\varepsilon = 0.01$ and is less than 5% for $\varepsilon \geq 0.05$.

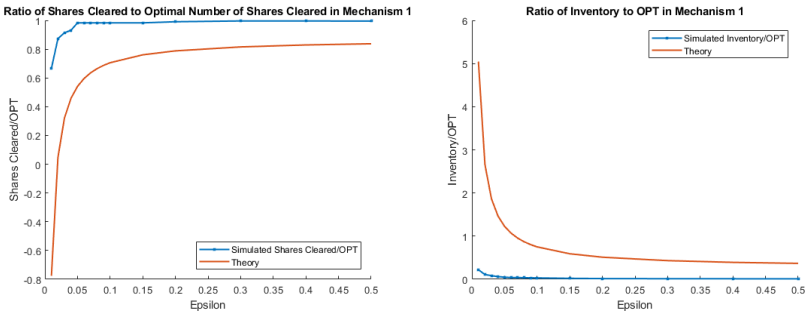

Fig. 1. Realized payoff and inventory figures relative to theoretical optimal in one-shot game for varying ε.

Learning Setting. In the *learning* setting, we plot the shares cleared *over time* as agents learn to bid given their valuations. We repeat the auction for 1000 rounds (1500 for the imbalance plot) with learning rate $\eta = 0.1$ and "fake" utility $\xi = 0.1$. Agents draw fixed valuations and then use the Social Exponential Weights described in Algorithm 4 to learn and bid each round. We repeat this process for several different values of ε.

The first three plots in Figure 2 tell similar stories: agents, and thus the system, learn to bid over time in such a way as to clear the optimal number of shares (were the mechanism privacy-free). The noisiness in the plots is due to privacy and depends on the choice of ε: the smaller the value of ε, the more likely the mechanism is to pick a sub-optimal price, even after agents learn to bid optimally. For $\varepsilon = 0.01$, the randomness of the mechanism induces enough noise as to occasionally forego a large portion of utility; at larger values of ε, the added randomness costs relatively little.

The fourth plot displays the imbalance between number of buyers bidding above vs. sellers bidding below the price chosen by the repeated *standard* (i.e., non-private) call auction when buyers use Social Exponential Weights. We highlight an interesting connection to real-world behavior: NYSE and NASDAQ perform pre-opening or pre-closing repeated "hypothetical" auctions aimed at price discovery. In these hypothetical auctions, the exchanges accept bids, announce the current price and imbalance, allow bidders to submit updated bids, and repeat. The pattern in imbalances documented by [3] agrees broadly with that of Figure 2: that is, the imbalance begins skewed to one side or another, but it repeatedly oscillates as bidders adjust before converging to a settled state.

ACKNOWLEDGMENTS

This work was partially funded by the Warren Center for Network and Data Sciences at the University of Pennsylvania.

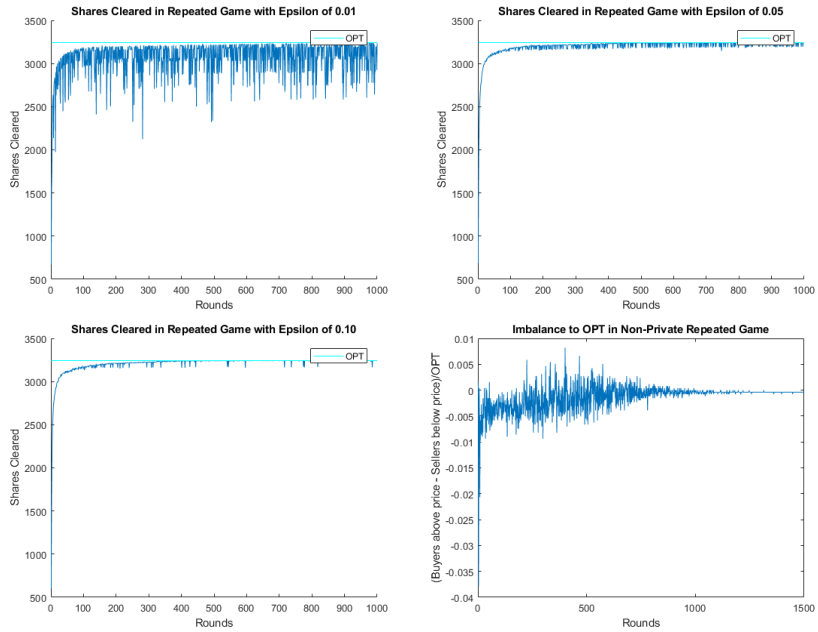

Fig. 2. The first three plots show the shares cleared over time, in the repeated setting of our mechanism, using Social Exponential Weights (Algorithm 4) for various choices of ε. The last plot shows the imbalance between buyers and sellers over time in a repeated (non-private) auction. The agents' updates use $\eta = 0.1$ and $\xi = 0.1$.

REFERENCES

[1] Eric Budish, Peter Cramton, and John Shim. 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics* 130, 4 (07 2015), 1547–1621. https://doi.org/10.1093/qje/qjv027 arXiv:https://academic.oup.com/qje/article-pdf/130/4/1547/30637414/qjv027.pdf

[2] Damien Challet. 2019. Strategic Behaviour and Indicative Price Diffusion in Paris Stock Exchange Auctions. In *New Perspectives and Challenges in Econophysics and Sociophysics*. Springer, 3–12.

[3] Damien Challet and Nikita Gourianov. 2018. Dynamical Regularities of US Equities Opening and Closing Auctions. *Market Microstructure and Liquidity* 4, 1 (2018).

[4] Zhili Chen, Tianjiao Ni, Hong Zhong, Shun Zhang, and Jie Cui. 2018. Differentially Private Double Spectrum Auction with Approximate Social Welfare Maximization. *arXiv preprint arXiv:1810.07873* (2018).

[5] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. 2015. Privacy and Truthful Equilibrium Selection for Aggregative Games. In *International Conference on Web and Internet Economics*. Springer, 286–299.

[6] Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. 2018. Strategic Classification from Revealed Preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 55–70.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.

[8] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407. https://doi.org/10.1561/0400000042

[9] Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. 2010. Boosting and Differential Privacy. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS '10)*. IEEE Computer Society, Washington, DC, USA, 51–60. https://doi.org/10.1109/FOCS.2010.12

[10] Eyal Even-Dar, Sham M. Kakade, Michael Kearns, and Yishay Mansour. 2006. (In)Stability Properties of Limit Order Dynamics. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC '06)*. Association for Computing Machinery, New York, NY, USA, 120–129. https://doi.org/10.1145/1134707.1134721

[11] Jim Gatheral. 2010. No-Dynamic-Arbitrage and Market Impact. *Quantitative finance* 10, 7 (2010), 749–759.

[12] Jim Gatheral. 2010. Three Models of Market Impact. In *Market Microstructure and High Frequency Data*. https://mfe.baruch.cuny.edu/wp-content/uploads/2017/05/Chicago2016OptimalExecution.pdf
[13] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. 2014. Private Matchings and Allocations. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing (STOC '14)*. Association for Computing Machinery, New York, NY, USA, 21–30. https://doi.org/10.1145/2591796.2591826
[14] Justin Hsu, Zhiyi Huang, Aaron Roth, and Zhiwei Steven Wu. 2016. Jointly Private Convex Programming. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 580–599.
[15] Sampath Kannan, Jamie Morgenstern, Aaron Roth, and Zhiwei Steven Wu. 2014. Approximately Stable, School Optimal, and Student-Truthful Many-to-One Matchings (via Differential Privacy). In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1890–1903.
[16] Michael Kearns, Mallesh M. Pai, Aaron Roth, and Jonathan Ullman. 2014. Mechanism Design in Large Games: Incentives and Privacy. *The American Economic Review* 104, 5 (2014), 431–435. http://www.jstor.org/stable/42920975
[17] M. Lewis. 2014. *Flash Boys: A Wall Street Revolt*. W. W. Norton. https://books.google.com/books?id=UcIkAwAAQBAJ
[18] Frank McSherry and Kunal Talwar. 2007. Mechanism Design via Differential Privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS '07)*. IEEE Computer Society, Washington, DC, USA, 94–103. https://doi.org/10.1109/FOCS.2007.41
[19] NASDAQ. 2020. NASDAQ Opening and Closing Crosses. https://www.nasdaqtrader.com/Trader.aspx?id=OpenClose
[20] NYSE. 2020. NYSE Opening and Closing Auctions Fact Sheet. https://www.nyse.com/publicdocs/nyse/markets/nyse/NYSE_Opening_and_Closing_Auctions_Fact_Sheet.pdf
[21] Simon Parsons, Marek Marcinkiewicz, Jinzhong Niu, and Steve Phelps. 2006. Everything you wanted to know about double auctions, but were afraid to (bid or) ask. (2006).
[22] Simon Parsons, Juan A Rodriguez-Aguilar, and Mark Klein. 2011. Auctions and bidding: A guide for computer scientists. *ACM Computing Surveys (CSUR)* 43, 2 (2011), 1–59.
[23] Ryan Rogers, Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. 2015. Inducing Approximately Optimal Flow Using Truthful Mediators. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. 471–488.
[24] Ryan M Rogers and Aaron Roth. 2014. Asymptotically Truthful Equilibrium Selection in Large Congestion Games. In *Proceedings of the Fifteenth ACM conference on Economics and Computation*. 771–782.
[25] Elaine Wah and Michael Wellman. 2013. Latency Arbitrage, Market Fragmentation, and Efficiency: A Two-Market Model. In *Proceedings of the Thirteenth ACM conference on Economics and Computation*.

A DIFFERENTIAL PRIVACY TOOLS

In this section, we remind the reader of mechanisms that are classically used to guarantee differential privacy. These mechanisms work by adding appropriately-chosen noise to the choices and outputs of a mechanism, so as to ensure that a change in a single individual's data cannot have a large distributional effect on the mechanism's output. The noise introduced by differentially private mechanisms depends not only on the level (ε, δ) of privacy one aims to guarantee, but also on the *sensitivity* of the query of interest. This sensitivity measures how much the real-valued function of interest is affected by a change in a single entry of an input data set, and will be formally defined in our introduced DP mechanisms.

A commonly used mechanism for releasing the answer to numerical queries while guaranteeing $(\varepsilon, 0)$-differential privacy is the Laplace mechanism. The Laplace mechanism takes a numerical query f as an input, and perturbs the value of f on the input data set with zero-mean Laplace noise that has scale proportional to $(\Delta f/\varepsilon)$ where Δf is the ℓ_1-sensitivity of f.

Definition 9 (Laplace Mechanism [7]). *Given a function $f : \mathcal{D}^n \to \mathbb{R}^k$ with ℓ_1-sensitivity Δf:*

$$\Delta f = \max_{\substack{D, D' \in \mathcal{D}^n \\ D \sim D'}} \|f(D) - f(D')\|_1,$$

a data set $D \in \mathcal{D}^n$, and a privacy parameter ε, the Laplace mechanism outputs:

$$f_\varepsilon(D) = f(D) + (W_1, \ldots, W_k)$$

where W_i's are i.i.d. random variables drawn from $Lap(\Delta f/\varepsilon)$.

We provide the privacy and accuracy guarantees of the Laplace mechanism below:

Theorem 8 (Privacy vs. Accuracy of the Laplace Mechanism [7]). *The Laplace Mechanism guarantees $(\varepsilon, 0)$-differential privacy and that with probability at least $1-\delta$,*

$$\|f_\varepsilon(D) - f(D)\|_\infty \leq \ln\left(\frac{k}{\delta}\right) \cdot \left(\frac{\Delta f}{\varepsilon}\right)$$

We remark that the Laplace mechanism can be used to privately output the answer to numerical queries. However, suppose we want to privately output the solution to a maximization problem defined on the input data. Then, directly adding noise to the optimal solution could completely destroy the objective value of the maximization problem in question (for example, in an auction, adding a small amount of noise on the price of an item could significantly reduce revenue). In such situations, the Laplace mechanism performs poorly, and a better choice of private mechanism is the Exponential Mechanism, defined below:

Definition 10 (Exponential Mechanism [18]). *Let $U : \mathcal{D}^n \times P \to \mathbb{R}$ be a utility function that takes a data set $D \in \mathcal{D}^n$ and a parameter $p \in P$ as inputs, and let ΔU be its sensitivity. In other words,*

$$\Delta U = \max_{p \in P} \max_{\substack{D, D' \in \mathcal{D}^n \\ D \sim D'}} |U(D,p) - U(D',p)|.$$

Given a data set $D \in \mathcal{D}^n$ and a privacy parameter ε, the exponential mechanism outputs $p \in P$ with probability proportional to $\exp\left(\frac{\varepsilon U(D,p)}{2\Delta U}\right)$ where $\exp(\cdot)$ is the exponential function.

Theorem 9 (Privacy vs. Accuracy of the Exponential Mechanism [18]). *The Exponential Mechanism guarantees $(\varepsilon, 0)$-differential privacy. Further, let $p_\varepsilon \in P$ be the output of the Exponential mechanism, we have that with probability at least $1-\delta$,*

$$\left| U(D, p_\varepsilon) - \max_{p \in P} U(D,p) \right| \leq \ln\left(\frac{|P|}{\delta}\right) \cdot \left(\frac{2\Delta U}{\varepsilon}\right)$$

An important property of differential privacy is that it is robust to *post-processing*. Applying any data-independent function to the output of an (ε, δ)-DP algorithm preserves (ε, δ)-differential privacy.

Lemma 2 (Post-Processing [7]). *Let $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}$ be an (ε, δ)-DP algorithm and let $f : \mathcal{R} \to \mathcal{R}'$ be any function. We have that the algorithm $f \circ \mathcal{M} : \mathcal{D}^n \to \mathcal{R}'$ is (ε, δ)-DP.*

Another important property of differential privacy is that DP algorithms can be composed adaptively with a graceful degradation in their privacy parameters.

Theorem 10 ((Simple) Composition [9]). *Let $\mathcal{M}_t$ be an $(\varepsilon_t, \delta_t)$-DP algorithm for $t \in [T]$. We have that the composition $\mathcal{M} = (\mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_T)$ is (ε, δ)-DP where $\varepsilon = \sum_t \varepsilon_t$ and $\delta = \sum_t \delta_t$.*

To prove that our mechanisms satisfy (ε, δ)-joint differential privacy, we will leverage the billboard lemma ([13]). The billboard lemma shows that for every individual i in the data set, restricting i's output to be a function of only the output of a differentially private mechanism (run on all agents' data) and his own input guarantees joint differential privacy.

Lemma 3 (Billboard Lemma [13]). *Suppose $\mathcal{M} : \mathcal{D}^n \to \mathcal{R}'$ is (ε, δ)-differentially private. Consider any set of functions $f_i : \mathcal{D}_i \times \mathcal{R}' \to \mathcal{R}$, where $\mathcal{D}_i$ is the portion of the data set containing i's data. The composition $\{f_i(\Pi_i(D), \mathcal{M}(D))\}$ is (ε, δ)-jointly differentially private, where $\Pi_i : \mathcal{D}^n \to \mathcal{D}_i$ is the projection to i's data.*

B A META-ALGORITHM FOR DIFFERENTIAL PRIVACY

We provide a meta algorithm that simply computes the difference of payoff bounds of Mechanisms 1 and 2

$$f \triangleq \frac{2\ln(1/\alpha)}{\varepsilon} + \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)} - \frac{4\ln(n/\alpha)}{\varepsilon}$$

in a differentially privately manner.[15] Then, based on the sign of f, the mechanism decides whether to run $\mathcal{M}_1$ or $\mathcal{M}_2$. The private computation of f will unavoidably add an extra term of order $O(1/\varepsilon)$ to the payoff bound.

ALGORITHM 5: Private Call Auction with Allocation: A Meta Algorithm ($\mathcal{M}_3$)

Input: Agents' valuations $(\mathbf{v}^s, \mathbf{v}^b)$, privacy level ε, confidence level α.
Output: Market price p, allocations $\mathbf{a} = (\mathbf{a}^s, \mathbf{a}^b)$.

$f \leftarrow \frac{2\ln(1/\alpha)}{\varepsilon} + \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)} - \frac{4\ln(n/\alpha)}{\varepsilon}$

$\widetilde{f} \leftarrow f + Lap\left(\frac{\sqrt{6\ln(1/\alpha)}}{\varepsilon}\right)$ ▷ Private estimate of f

if $\widetilde{f} < 0$ **then**

Run $\mathcal{M}_1\left(\mathbf{v}^s, \mathbf{v}^b, \varepsilon, \alpha\right)$ (Algorithm 1) and get $p, \mathbf{a}$.

else

Run $\mathcal{M}_2\left(\mathbf{v}^s, \mathbf{v}^b, \varepsilon\right)$ (Algorithm 2) and get $p, \mathbf{a}$.

end

Claim 5 (Joint DP). *The allocation mechanism described in Algorithm 5 is 7ε-joint DP.*

Theorem 11 (Payoff and Inventory of Mechanism). *Suppose $OPT \geq 5\ln(V/\alpha)/\varepsilon$.*

(1) *Payoff: with probability $1 - 18\alpha$,*

$$\Pi(\mathcal{M}_3) \geq OPT - \frac{2\ln(\frac{V}{\alpha})}{\varepsilon} - \min\left\{\frac{2\ln\left(\frac{1}{\alpha}\right)}{\varepsilon} + \sqrt{6\left(OPT + \frac{\ln(\frac{1}{\alpha})}{\varepsilon}\right)\ln\left(\frac{1}{\alpha}\right)}, \frac{4\ln(\frac{n}{\alpha})}{\varepsilon}\right\} - \frac{\sqrt{6}\ln^{1.5}(\frac{1}{\alpha})}{\varepsilon}$$

(2) *Inventory: with probability $1 - 14\alpha$,*

$$I(\mathcal{M}_3) \leq 4\min\left\{\frac{2\ln\left(\frac{1}{\alpha}\right)}{\varepsilon} + \sqrt{6\left(OPT + \frac{\ln\left(\frac{1}{\alpha}\right)}{\varepsilon}\right)\ln\left(\frac{1}{\alpha}\right)}, \frac{4\ln(\frac{n}{\alpha})}{\varepsilon}\right\} + \frac{4\sqrt{6}\ln^{1.5}(\frac{1}{\alpha})}{\varepsilon} + \frac{10\ln(\frac{1}{\alpha})}{\varepsilon} + \frac{4\ln(\frac{2}{\alpha})}{3}$$

This theorem follows from Theorems 1 and 2, as well as the accuracy guarantee of the Laplace mechanism used in Algorithm 5 to compute f. We defer the full proof to Appendix D.3.

[15] OPT is a function of the input data set, hence a direct comparison of the bounds without addition of noise may leak information about the reported bids.

C PROOFS OF PRIVACY GUARANTEES OF OUR MECHANISMS

Proof of Claim 1. We start the proof by noticing that the sensitivity of Π (as per Definition 10) is 1: indeed, changing one element in the data $(\mathbf{v}^s, \mathbf{v}^b)$, i.e. the valuation of a single agent, will change the number of shares cleared by at most one, for any fixed price p. We can therefore conclude that by Theorem 9 the mechanism that takes the data set as input and outputs a price $p \propto \exp\left(\varepsilon \Pi(p, \mathbf{v}^s, \mathbf{v}^b)/2\right)$ is ε-DP. One can similarly argue that given a fixed price p, quantities $\sum_{i \in \mathcal{S}} \mathbf{1}[p \geq \mathbf{v}_i^s]$ and $\sum_{j \in \mathcal{B}} \mathbf{1}[p \leq \mathbf{v}_j^b]$ have sensitivity 1 (see Definition 9), and therefore by Theorem 8, $\widehat{s}$ and $\widehat{b}$ both satisfy ε-DP. We can now invoke the Composition Theorem 10 to conclude that the triplet $(p, \widehat{s}, \widehat{b})$ computed in Algorithm 1 satisfies 3ε-differential privacy. The claim then follows by the Billboard Lemma 3 and noticing that each agent's allocation depends only on their own data and the triplet $(p, \widehat{s}, \widehat{b})$. □

Proof of Claim 2. Notice first that according to Definition 10, the sensitivity of Π is 1 because for any p, changing one agent's valuation can change Π by at most 1. Now fixing the price p output by the first exponential mechanism, it similarly follows that the sensitivity of loss functions L^s and L^b are 2. We therefore have that the exponential mechanisms outputting p, τ^s, and τ^b are all ε-DP by Theorem 9, and hence the triplet (p, τ^s, τ^b) satisfies 3ε-DP by the Composition Theorem 10. The claim then follows by the Billboard Lemma 3 and noticing that each agent's allocation depends only on their own data and the triplet (p, τ^s, τ^b). □

Proof of Claim 5. First, we note that the sensitivity of f is upper-bounded by $\sqrt{6\ln(1/\alpha)}$. Indeed, we remind the reader that OPT has sensitivity of 1, and note that for any $x \geq 0$

$$\sqrt{6\ln(1/\alpha)} \geq \sqrt{6\left(x + 1 + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)} - \sqrt{6\left(x + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)},$$

using the classical inequality $\sqrt{a} + \sqrt{b} \geq \sqrt{a+b}$. Therefore, by Theorem 8, the computation of $\widetilde{f}$ is ε-differential private. In each of mechanisms $\mathcal{M}_1$ and $\mathcal{M}_2$, p_1 the price output by $\mathcal{M}_1$, respectively p_2 the price output by $\mathcal{M}_2$, are computed in an ε-differentially private manner. Similarly, $\hat{s}$, $\hat{b}$ (resp. τ^s, τ^b), the private counts of the number of agents willing to trade in $\mathcal{M}_1$ at price p_1 (resp. the thresholds picked by mechanism $\mathcal{M}_2$ for price p_2) are each the result of an ε-differentially private computation (conditional on p_1, p_2). In turn, our mechanism can be seen as one that computes $(\widetilde{f}, p_1, p_2, \hat{s}, \hat{b}, \tau^s, \tau^b)$ in a 7ε-differentially private manner (by the composition guarantee of Theorem 10), then outputs an allocation $\mathbf{a}_i^s$ for each given seller i (resp. $\mathbf{a}_j^b$ for each buyer j) as a function of only $\mathbf{v}_i^s$ (resp. $\mathbf{v}_j^b$) and $(\widetilde{f}, p_1, p_2, \hat{s}, \hat{b}, \tau^s, \tau^b)$. Hence, by Lemma 3, $\mathcal{M}$ is 7ε-joint differentially private. □

D PROOFS OF PROFIT AND INVENTORY OF OUR MECHANISMS

D.1 Proof of Theorem 1

Proof. We will be using the following concentration inequalities in our proof.

Fact 1 (Multiplicative Chernoff Bound). *Let $\{X_i\}_{i=1}^n$ be a collection of independent random variables where $X_i \in [0,1]$ for all i. Let $S = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[S]$. We have that for any $0 \leq t \leq 1$,*

$$\Pr\left[S < (1-t)\mu\right] \leq e^{-\frac{\mu t^2}{2}}$$

Fact 2 (Bernstein's Inequality). *Let $\{X_i\}_{i=1}^n$ be a collection of i.i.d random variables where for each i, $X_i \in [0,1]$, $E[X_i] = \mu$, and $Var(X_i) = \sigma^2$. Let $S = \sum_{i=1}^n X_i$. We have that for any $t \geq 0$,*

$$\Pr\left[|S - n\mu| > t\right] \leq 2e^{-\frac{t^2}{2n\sigma^2 + 2t/3}}$$

Let $s(p) = \sum_{i \in S} \mathbf{1}\left[p \geq \mathbf{v}_i^s\right]$ and $b(p) = \sum_{j \in B} \mathbf{1}\left[p \leq \mathbf{v}_j^b\right]$ be the number of sellers and buyers available at price p, where p is the price chosen by the exponential mechanism. Note that as p is a random variable, so are $s(p)$ and $b(p)$. From now on, for simplicity of notations, we omit the dependency of s and b in p. We start the proof by noting that by the accuracy guarantee of the Laplace mechanism (see Theorem 8) and a union bound, we have with probability at least $1 - 2\alpha$ that

$$\left|\widehat{b} - b\right| \leq \frac{\ln(1/\alpha)}{\varepsilon}, \quad |\widehat{s} - s| \leq \frac{\ln(1/\alpha)}{\varepsilon}, \tag{3}$$

and by the accuracy guarantee of the Exponential mechanism (see Theorem 9), we have with probability at least $1 - \alpha$ that,

$$\Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right) = \min\{s, b\} \geq \mathrm{OPT} - \frac{2\ln(V/\alpha)}{\varepsilon}. \tag{4}$$

By a union bound, Equations (3) and (4) hold simultaneously with probability at least $1 - 3\alpha$, and throughout this proof we condition on these events. Let $\widetilde{s} = \sum_{i \in \mathcal{S}} \mathbf{a}_i^s$ and $\widetilde{b} = \sum_{j \in \mathcal{B}} \mathbf{a}_j^b$ be the number of sellers and buyers who participate in a trade, output by the mechanism.

First, let's focus on $\widetilde{s}$. Observe that

$$\widetilde{s} \,|\, s, \widehat{s}, \widehat{b} \sim \mathit{Binomial}\left(s, \widehat{q} = \min\left\{1, \frac{\left(\widehat{b}\right)_+}{\left(\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right\}\right).$$

Note that we have

$$\begin{aligned}
\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon} &\geq s - \frac{2\ln(1/\alpha)}{\varepsilon} && \text{(by Equation (3))} \\
&\geq \mathrm{OPT} - \frac{2\ln(1/\alpha)}{\varepsilon} - \frac{2\ln(V/\alpha)}{\varepsilon} && \text{(by Equation (4))} \\
&\geq \mathrm{OPT} - \frac{4\ln(V/\alpha)}{\varepsilon} && (V \geq 1) \\
&\geq \frac{\ln(V/\alpha)}{\varepsilon} && \text{(by assumption, } OPT \geq \tfrac{5\ln(V/\alpha)}{\varepsilon}\text{)} \\
&> 0,
\end{aligned}$$

and also

$$\begin{aligned}
\widehat{b} &\geq b - \frac{\ln(1/\alpha)}{\varepsilon} && \text{(by Equation (3))} \\
&\geq \mathrm{OPT} - \frac{\ln(1/\alpha)}{\varepsilon} - \frac{2\ln(V/\alpha)}{\varepsilon} && \text{(by Equation (4))} \\
&\geq \mathrm{OPT} - \frac{3\ln(V/\alpha)}{\varepsilon} && (V \geq 1) \\
&\geq \frac{2\ln(V/\alpha)}{\varepsilon} && \text{(by assumption, } OPT \geq \tfrac{5\ln(V/\alpha)}{\varepsilon}\text{)} \\
&> 0.
\end{aligned}$$

As such, $\widehat{q}$ is well-defined, and we can rewrite it as

$$\widehat{q} = \min\left(1, \frac{\widehat{b}}{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}}\right).$$

We have by the multiplicative Chernoff Bound (Fact 1) with $t = \sqrt{\frac{2\ln(1/\alpha)}{s\widehat{q}}}$ that,

$$\widetilde{s} \geq s\widehat{q} - \sqrt{2s\widehat{q}\ln(1/\alpha)} \tag{5}$$

with probability at least $1 - \alpha$ when $t \leq 1$. Note that the bound applies when $t > 1$ too, noting that then $s\widehat{q} - \sqrt{2s\widehat{q}\ln(1/\alpha)} < 0$ but $\widetilde{s} \geq 0$. In what follows, we will provide an upper bound and a lower bound for the term $s\widehat{q}$ so that we can further lower bound $\widetilde{s}$ in Equation (5). Symmetrically, we can get a similar lower bound for $\widetilde{b}$ which completes the first part of the proof because $\Pi(\mathcal{M}) = \min\{\widetilde{s}, \widetilde{b}\}$.

On the one hand, note that

$$\begin{aligned}
s\widehat{q} &= s \cdot \frac{\min\left\{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}, \widehat{b}\right\}}{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}} \\
&\geq \frac{s}{s} \cdot \min\left\{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}, \widehat{b}\right\} \\
&= \min\left\{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}, \widehat{b}\right\} \\
&\geq \min\{s, b\} - \frac{2\ln(1/\alpha)}{\varepsilon} \\
&\geq \mathrm{OPT} - \frac{2\ln(1/\alpha)}{\varepsilon} - \frac{2\ln(V/\alpha)}{\varepsilon}.
\end{aligned} \tag{6}$$

The first inequality follows from $\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon} \leq s$ by Equation (3). The second inequality follows from

$$\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon} \geq s - \frac{2\ln(1/\alpha)}{\varepsilon}, \quad \widehat{b} \geq b - \frac{\ln(1/\alpha)}{\varepsilon}$$

by Equation (3). The third inequality is a direct application of Equation (4). On the other hand,

$$\begin{aligned}
s\widehat{q} &= s \cdot \frac{\min\left\{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}, \widehat{b}\right\}}{\widehat{s} - \frac{\ln(1/\alpha)}{\varepsilon}} \\
&\leq \frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \cdot \min\left\{s, b + \frac{\ln(1/\alpha)}{\varepsilon}\right\} \\
&\leq \frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \cdot \left(\min\{s, b\} + \frac{\ln(1/\alpha)}{\varepsilon}\right) \\
&\leq \frac{s}{s - \frac{2\ln(V/\alpha)}{\varepsilon}} \cdot \left(\min\{s, b\} + \frac{\ln(1/\alpha)}{\varepsilon}\right) \\
&\leq 3\left(\min\{s, b\} + \frac{\ln(1/\alpha)}{\varepsilon}\right) \\
&\leq 3\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right).
\end{aligned} \tag{7}$$

The first inequality follows from

$$s - \frac{\ln(1/\alpha)}{\varepsilon} \leq \widehat{s} \leq s + \frac{\ln(1/\alpha)}{\varepsilon}, \quad \widehat{b} \leq b + \frac{\ln(1/\alpha)}{\varepsilon}$$

by Equation (3). The second-to-last inequality follows from the fact that

$$f : (0, +\infty) \to \mathbb{R}, \quad f(x) = \frac{x}{x - \frac{2\ln(V/\alpha)}{\varepsilon}} = \frac{1}{1 - \frac{2\ln(V/\alpha)}{\varepsilon x}}$$

is a non-increasing function of x, and that by Equation (4),

$$s \geq \mathrm{OPT} - \frac{2\ln(V/\alpha)}{\varepsilon} \geq \frac{3\ln(V/\alpha)}{\varepsilon}.$$

Combining Equations (5), (6), and (7), we obtain that with probability $1 - 4\alpha$,

$$\widetilde{s} \geq \mathrm{OPT} - \frac{2\ln(1/\alpha)}{\varepsilon} - \frac{2\ln(V/\alpha)}{\varepsilon} - \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}. \tag{8}$$

Symmetrically, we can get the same bound for $\widetilde{b}$: with probability $1 - 4\alpha$,

$$\widetilde{b} \geq \mathrm{OPT} - \frac{2\ln(1/\alpha)}{\varepsilon} - \frac{2\ln(V/\alpha)}{\varepsilon} - \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)} \tag{9}$$

Combining Equations (8) and (9) and noting that $\Pi(\mathcal{M}_1) = \min\{\widetilde{s}, \widetilde{b}\}$ proves the first part of the theorem. We conclude the proof by noting that the statements hold with probability at least $1 - 8\alpha$ by union bound.

Let us now analyze the inventory of the mechanism. We have by the triangle inequality that

$$I(\mathcal{M}_1) = \left|\widetilde{s} - \widetilde{b}\right| \leq \left|\widetilde{s} - \min\{s, b\}\right| + \left|\widetilde{b} - \min\{s, b\}\right| \tag{10}$$

We will provide an upper bound for the first term, and by symmetry, an upper bound on the second will follow immediately. We have that by the triangle inequality

$$\left|\widetilde{s} - \min\{s, b\}\right| \leq \left|\widetilde{s} - s\widehat{q}\right| + \left|s\widehat{q} - \min\{s, b\}\right| \tag{11}$$

First

$$\begin{aligned}
\left|\widetilde{s} - s\widehat{q}\right| &\leq \sqrt{2s\widehat{q}(1-\widehat{q})\ln(2/\alpha)} + \frac{2\ln(2/\alpha)}{3} \\
&\leq \sqrt{2s\widehat{q}\ln(2/\alpha)} + \frac{2\ln(2/\alpha)}{3} \\
&\leq \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(2/\alpha)} + \frac{2\ln(2/\alpha)}{3}
\end{aligned} \tag{12}$$

where the first inequality holds with probability $1 - \alpha$ and follows from the Bernstein's inequality (Fact 2) with $\sigma^2 = \widehat{q}(1-\widehat{q})$ and taking $t = \frac{\ln(2/\alpha)}{3} + \sqrt{\frac{\ln^2(2/\alpha)}{9} + 2n\sigma^2\ln(2/\alpha)}$. Notice that $t \leq \frac{2\ln(2/\alpha)}{3} + \sqrt{2n\sigma^2\ln(2/\alpha)}$. The last inequality follows from the upper bound developed in Equation (7). Second,

$$\left|s\widehat{q} - \min\{s, b\}\right| \leq \frac{9\ln(1/\alpha)}{\varepsilon}. \tag{13}$$

This is because as we showed in Equation (6),

$$-\frac{2\ln(1/\alpha)}{\varepsilon} \leq s\widehat{q} - \min\{s, b\},$$

and as we showed in Equation (7),

$$\begin{aligned}
s\widehat{q} - \min\{s,b\} &\leq \frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \cdot \left(\min\{s,b\} + \frac{\ln(1/\alpha)}{\varepsilon}\right) - \min\{s,b\} \\
&= \left(\frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} - 1\right) \cdot \min\{s,b\} + \frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \cdot \frac{\ln(1/\alpha)}{\varepsilon} \\
&\leq \left(\frac{2\ln(1/\alpha)/\varepsilon}{s - \frac{2\ln(1/\alpha)}{\varepsilon}}\right) \cdot s + \frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \cdot \frac{\ln(1/\alpha)}{\varepsilon} \\
&= \frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \cdot \frac{3\ln(1/\alpha)}{\varepsilon} \\
&\leq \frac{9\ln(1/\alpha)}{\varepsilon}
\end{aligned}$$

Because we showed in Equation (7) that $\frac{s}{s - \frac{2\ln(1/\alpha)}{\varepsilon}} \leq 3$.

Putting together Equations (11), (12), and (13) we get that with probability $1 - 3\alpha$,

$$|\widetilde{s} - \min\{s,b\}| \leq \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(2/\alpha)} + \frac{2\ln(2/\alpha)}{3} + \frac{9\ln(1/\alpha)}{\varepsilon}.$$

Swapping the roles of the buyers and a similar proof yields the same bound on $\left|\widetilde{b} - \min\{s,b\}\right|$.

A union bound completes the proof, by Equation (10). □

D.2 Proof of Theorem 2

Proof. Let us define $\widetilde{s} = \sum_{i \in \mathcal{S}} \mathbf{a}_i^s$ and $\widetilde{b} = \sum_{j \in \mathcal{B}} \mathbf{a}_j^b$ to be the number of sellers and buyers who participate in a trade under output allocation $\mathbf{a}$. We have that with probability at least $1 - 3\alpha$, by the accuracy guarantee of the Exponential mechanism (see Theorem 9),

$$\Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right) \geq \mathrm{OPT} - \frac{2\ln(V/\alpha)}{\varepsilon}, \tag{14}$$

and that since $\min_{\tau^s} L^s(\tau^s, p, \mathbf{v}^s, \mathbf{v}^b) = \min_{\tau_b} L^b(\tau^b, p, \mathbf{v}^s, \mathbf{v}^b) = 0$,

$$\left|\widetilde{s} - \Pi(p, \mathbf{v}^s, \mathbf{v}^b)\right| = L^s(\tau^s, p, \mathbf{v}^s, \mathbf{v}^b) \leq \frac{4\ln\left(n^s/\alpha\right)}{\varepsilon} \quad \Longrightarrow \quad \widetilde{s} \geq \Pi(p, \mathbf{v}^s, \mathbf{v}^b) - \frac{4\ln\left(n^s/\alpha\right)}{\varepsilon}, \tag{15}$$

$$\left|\widetilde{b} - \Pi(p, \mathbf{v}^s, \mathbf{v}^b)\right| = L^b(\tau^b, p, \mathbf{v}^s, \mathbf{v}^b) \leq \frac{4\ln\left(n^b/\alpha\right)}{\varepsilon} \quad \Longrightarrow \quad \widetilde{b} \geq \Pi(p, \mathbf{v}^s, \mathbf{v}^b) - \frac{4\ln\left(n^b/\alpha\right)}{\varepsilon} \tag{16}$$

We therefore have that

$$\Pi\left(\mathcal{M}_2\right) = \min\left\{\widetilde{s}, \widetilde{b}\right\} \geq \mathrm{OPT} - \frac{2\ln(V/\alpha)}{\varepsilon} - \frac{4\ln(n/\alpha)}{\varepsilon} \tag{17}$$

Let's now analyze the inventory introduced by the private mechanism. We have that

$$\begin{aligned} I(\mathcal{M}_2) &= \left|\widetilde{s} - \widetilde{b}\right| \\ &\leq \left|\widetilde{s} - \Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right)\right| + \left|\widetilde{b} - \Pi\left(p, \mathbf{v}^s, \mathbf{v}^b\right)\right| \\ &= L^s(\tau^s, p, \mathbf{v}^s, \mathbf{v}^b) + L^b(\tau^b, p, \mathbf{v}^s, \mathbf{v}^b) \\ &\leq \frac{4\ln\left(n^s/\alpha\right)}{\varepsilon} + \frac{4\ln\left(n^b/\alpha\right)}{\varepsilon} \\ &\leq \frac{8\ln(n/\alpha)}{\varepsilon} \end{aligned}$$

where the second inequality holds with probability $1 - 2\alpha$ by Equations (15) and (16). □

D.3 Proof of Theorem 11

Proof. This theorem follows from Theorems 1 and 2 and conditioning on the accuracy guarantee of the additional Laplace mechanism used in Algorithm 5:

$$\text{w.p. } 1-\alpha, \qquad \left|\widetilde{f} - f\right| \leq \frac{\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon} \tag{18}$$

Suppose $\widetilde{f} < 0$. Note that in this case,

$$\begin{aligned} \text{OPT} - \Pi(\mathcal{M}_3) &= \text{OPT} - \Pi(\mathcal{M}_1) \\ &\leq \frac{2\ln(V/\alpha)}{\varepsilon} + \frac{2\ln(1/\alpha)}{\varepsilon} + \sqrt{6\left(\text{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)} && (\star) \\ &= \frac{2\ln(V/\alpha)}{\varepsilon} + \frac{4\ln(n/\alpha)}{\varepsilon} + f \\ &\leq \frac{2\ln(V/\alpha)}{\varepsilon} + \frac{4\ln(n/\alpha)}{\varepsilon} + \widetilde{f} + \frac{\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon} \\ &\leq \frac{2\ln(V/\alpha)}{\varepsilon} + \frac{4\ln(n/\alpha)}{\varepsilon} + \frac{\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon} && (\star\star) \end{aligned}$$

where the first inequality follows from Theorem 1, with probability $1 - 8\alpha$. The second inequality follows from Equation 18, with probability $1 - \alpha$. Combining the bounds given by the second and the last inequalities (specified by $\star$ and $\star\star$), we get that with probability $1 - 9\alpha$,

$$\begin{aligned} \text{OPT} - \Pi(\mathcal{M}_3) \leq \min\left\{\frac{2\ln(1/\alpha)}{\varepsilon} + \sqrt{6\left(\text{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}, \frac{4\ln(n/\alpha)}{\varepsilon}\right\} \\ + \frac{2\ln(V/\alpha)}{\varepsilon} + \frac{\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon}. \end{aligned}$$

A similar analysis for $\widetilde{f} \geq 0$ which uses Theorem 2 gives us the same bound and proves the first part of the theorem.

Now let's look at the inventory. Suppose $\widetilde{f} \leq 0$. We have that

$$\begin{aligned}
I(\mathcal{M}_3) &= I(\mathcal{M}_1) \\
&\leq \frac{18\ln(1/\alpha)}{\varepsilon} + 2\sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(2/\alpha)} + \frac{4\ln(2/\alpha)}{3} \\
&\leq \frac{18\ln(1/\alpha)}{\varepsilon} + 4\sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)} + \frac{4\ln(2/\alpha)}{3} \qquad (\star) \\
&= 4\left(f + \frac{4\ln(n/\alpha)}{\varepsilon}\right) + \frac{10\ln(1/\alpha)}{\varepsilon} + \frac{4\ln(2/\alpha)}{3} \\
&\leq 4\left(\widetilde{f} + \frac{\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon} + \frac{4\ln(n/\alpha)}{\varepsilon}\right) + \frac{10\ln(1/\alpha)}{\varepsilon} + \frac{4\ln(2/\alpha)}{3} \\
&\leq 4\left(\frac{\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon} + \frac{4\ln(n/\alpha)}{\varepsilon}\right) + \frac{10\ln(1/\alpha)}{\varepsilon} + \frac{4\ln(2/\alpha)}{3} \qquad (\star\star)
\end{aligned}$$

where the first inequality follows from Theorem 1, with probability $1 - 6\alpha$. The second inequality follows because $\alpha < 1/2$ (note we need $\alpha < 1/18$ to give non-trivial guarantee for the payoff of the mechanism). The third inequality follows from Equation 18 with probability $1 - \alpha$. Looking at the bounds given by the third and the last lines of the above equation (specified by $\star$ and $\star\star$), we get that with probability $1 - 7\alpha$,

$$\begin{aligned}
I(\mathcal{M}_3) \leq 4\min\left\{\frac{2\ln(1/\alpha)}{\varepsilon} + \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}, \frac{4\ln(n/\alpha)}{\varepsilon}\right\} \\
+ \frac{4\sqrt{6}\ln^{1.5}(1/\alpha)}{\varepsilon} + \frac{10\ln(1/\alpha)}{\varepsilon} + \frac{4\ln(2/\alpha)}{3}.
\end{aligned}$$

A similar analysis for $\widetilde{f} \geq 0$ which uses Theorem 2 gives us the same bound and proves the second part of the theorem. □

E PROOF OF THEOREM 3

Consider the following family of data sets: first, we initialize D_0 as the data set that has n sellers with valuations $\{1, \ldots, n\}$, and n buyers with valuations $\{n, \ldots, 2n-1\}$. We then recursively construct D_l for all l. To construct D_{l+1} from D_l, we increase all valuations in D_l by 1, and assign buyers' (resp. sellers) identities in D_{l+1} such that all buyers (resp. sellers) except one have the same valuation as in D_l. Equivalently, our construction works as follows: for any $l \in \mathbb{N}$,

$$D_l = \begin{cases} \mathbf{v}^s = \{l+1, \ldots, l+n\} & \text{sellers' valuations} \\ \mathbf{v}^b = \{l+n, \ldots, l+2n-1\} & \text{buyers' valuations,} \end{cases}$$

up to re-ordering of the agents' identities. The result will follow from the fact that a differentially private algorithm should output similar distributions of prices on data sets D_0 and D_l, but that at the same time, for l large enough, D_0 and D_l are far enough from each other that no distribution of prices can perform well over both of them.

We first show the following lemma, which will be of use in the proof of Theorem 3:

Lemma 4. *Let $\{D_l\}$ be the family of data sets described above. If $\mathcal{A} : \mathcal{D}^n \to P$ is an (ε, δ)-DP algorithm, then for every price $p \in P$ and every $k, m \in \mathbb{N}$:*

$$\Pr\left[|\mathcal{A}(D_k) - p| < m\right] \geq e^{-2k\varepsilon}\Pr\left[|\mathcal{A}(D_0) - p| < m\right] - 2k\delta.$$

Proof. By the definition of (ε, δ)-DP, if D and D' are neighboring data sets, we must have that for any event E,

$$\Pr[\mathcal{A}(D) \in E] \leq e^{\varepsilon} \Pr[\mathcal{A}(D') \in E] + \delta,$$

or equivalently

$$\Pr[\mathcal{A}(D') \in E] \geq e^{-\varepsilon} \left(\Pr[\mathcal{A}(D) \in E] - \delta\right) \tag{19}$$

Notice that for every k, by construction, D_k and D_{k+1} differ by only two entries (one buyer's and one seller's valuation). This immediately implies that D_k and D_0 differ by at most $2k$ entries, hence we can apply inequality (19) recursively $2k$ times to obtain that for any event E,

$$\begin{aligned}
\Pr[\mathcal{A}(D_k) \in E] &\geq e^{-\varepsilon} \left(e^{-\varepsilon} \ldots \left(e^{-\varepsilon} \Pr[\mathcal{A}(D_0) \in E] - \delta\right) \ldots - \delta\right) - \delta \\
&= e^{-2k\varepsilon} \Pr[\mathcal{A}(D_0) \in E] - \delta(e^{-(2k-1)\varepsilon} + e^{-(2k-2)\varepsilon} + \ldots. + e^{-\varepsilon} + 1) \\
&\geq e^{-2k\varepsilon} \Pr[\mathcal{A}(D_0) \in E] - 2k\delta
\end{aligned}$$

where the last inequality follows from the fact that $e^x \leq 1$ for $x \leq 0$. Fixing the price p and k, m, and taking E to be the ball of radius m around p, i.e.

$$E = \{p' : |p' - p| < m\}$$

concludes the proof. □

We are now ready to prove Theorem 3.

Proof of Theorem 3. In this proof, for any given data set $D = (\mathbf{v}^s, \mathbf{v}^b)$, we let

$$u(\mathcal{A}, D) \triangleq \min \left\{ \sum_{i \in \mathcal{S}} \mathbb{1}\left[\mathbf{v}_i^s \leq p\right], \sum_{j \in S} \mathbb{1}\left[\mathbf{v}_j^b \geq p\right] \right\}$$

where p is drawn according to $\mathcal{A}(D)$.

First, we note that in data set D_0, at most n trades (where every trading agent gets non-negative utility) can occur, setting a price of n. Further, n is the unique price that makes n trades possible, noting that decreasing (resp. increasing) the price leads to strictly less than n sellers (resp. buyers) willing to trade at that price. We let $p_0^* = n$ be this (unique) optimal price that clears n shares on data set D_0. For a given (ε, δ)-DP algorithm $\mathcal{A} : \mathcal{D}^n \to P$ that outputs a price p given an input data set $D = (\mathbf{v}^s, \mathbf{v}^b)$, let us define, for any $k, m \in \mathbb{N}$ (we will choose these values later on),

$$q_m^0 := \Pr\left[|\mathcal{A}(D_0) - p_0^*| < m\right], \; q_m^k := \Pr\left[|\mathcal{A}(D_k) - p_0^*| < m\right].$$

Notice by Lemma 4 that

$$q_m^k \geq e^{-2k\varepsilon} q_m^0 - 2k\delta. \tag{20}$$

Now, fix $m = \lceil \frac{1}{\varepsilon} \rceil$, $k = 2\lceil \frac{1}{\varepsilon} \rceil$, and take $n \geq m$. We have that the expected loss of $\mathcal{A}$ on D_0 is

$$\begin{aligned}
\mathbb{E}_{\mathcal{A}}\left[L(\mathcal{A}, D_0)\right] &= \mathrm{OPT}(D_0) - \mathbb{E}_{\mathcal{A}}\left[u(\mathcal{A}, D_0)\right] \\
&= n - \mathbb{E}_{\mathcal{A}}\left[u(\mathcal{A}, D_0)\right] \\
&\geq n - \left(q_m^0 \cdot n + (1 - q_m^0) \cdot (n - m)\right) \\
&= (1 - q_m^0) \cdot m \\
&\geq (1 - q_m^0) \cdot \left(\frac{1}{\varepsilon}\right).
\end{aligned} \tag{21}$$

The first inequality follows from a simple application of the law of total expectation on event $E = \{p : |p - p_0^*| < \frac{m}{2n}\}$ and its complement: with probability $1 - q_m^0$ the outputted price is outside E, which implies that it can only clear at most $n - m \geq 0$ shares (picking a price that is m away from n necessarily leads to either m fewer buyers or m fewer sellers willing to trade); the rest of the time, with probability q_m^0, algorithm $\mathcal{A}$ clears at most n shares. The second inequality is an immediate consequence of the choice of m. Similarly, on data set D_k,

$$\begin{aligned}
\mathbb{E}_{\mathcal{A}}\left[L(\mathcal{A}, D_k)\right] &= \mathrm{OPT}_k - u(\mathcal{A}, D_k) \\
&= n - \mathbb{E}_{\mathcal{A}}\left[u(\mathcal{A}, D_k)\right] \\
&\geq n - \left(q_m^k \cdot (n - (k - m)) + (1 - q_m^k) \cdot n\right) \\
&= q_m^k \cdot (k - m) \qquad (22) \\
&\geq \left(e^{-2k\varepsilon} q_m^0 - 2k\delta\right) \cdot (k - m) \\
&\geq \left(e^{-8} q_m^0 - 8(\delta/\varepsilon)\right) \cdot \left(\frac{1}{\varepsilon}\right)
\end{aligned}$$

where the first inequality follows from another use of the law of total expectation on the event E and its complement (notice we choose our parameters so that $k > m$ and $n \geq k - m$): with probability q_m^k, the price is at most $n + m$, and there are $k - m$ sellers that are willing to trade at price $n + m$ but not at price $n + k$, implying that such a price clears at most $n - (k - m)$ shares; the rest of the time, the number of shares cleared is at most n always. The second follows from Equation (20) and the last one follows from the choice of k and m and the fact that $\varepsilon\lceil\frac{1}{\varepsilon}\rceil \leq 1 + \varepsilon \leq 2$ for $0 \leq \varepsilon \leq 1$. Now let $L(\mathcal{A})$ be the worst-case expected loss of $\mathcal{A}$. We have that

$$\begin{aligned}
L(\mathcal{A}) &\geq \max\left\{(1 - q_m^0), \left(e^{-8} q_m^0 - 8(\delta/\varepsilon)\right)\right\} \cdot \left(\frac{1}{\varepsilon}\right) \\
&\geq \left(\frac{e^{-8} - 8(\delta/\varepsilon)}{1 + e^{-8}}\right) \cdot \left(\frac{1}{\varepsilon}\right)
\end{aligned}$$

where the first inequality follows from Equations (21) and (22) and the second is a simple observation that $f(q_m^0) := \max\left\{(1 - q_m^0), \left(e^{-8} q_m^0 - 8(\delta/\varepsilon)\right)\right\}$ is minimized at $q_m^0 = \frac{1 + 8(\delta/\varepsilon)}{1 + e^{-8}}$. Notice the lower bound is valid only when $\delta < \frac{e^{-8}}{8}\varepsilon = O(\varepsilon)$. This proves our claim that $L(\mathcal{A}) = \Omega\left(\frac{1}{\varepsilon}\right)$. □

F PROOFS OF APPROXIMATE TRUTHFULNESS

Our proof of truthfulness for Mechanism 1 will leverage the following lemma, which shows the output of an $(\varepsilon, 0)$-DP mechanism does not change by much in expectation when the input data set is changed by at most one element.

Lemma 5. *Let $Y = \mathcal{M}(D)$ where $\mathcal{M} : D \to \mathcal{Y}$ is an $(\varepsilon, 0)$-DP mechanism, and let $\max_{y \in \mathcal{Y}} |y| \leq K$. Then for any neighboring data sets $D \sim D'$,*

$$\left|\,\mathbb{E}\left[Y(D)\right] - \mathbb{E}\left[Y(D')\right]\,\right| \leq (e^{\varepsilon} - 1)K$$

Proof. $Y(D)$ and $Y(D')$ are random variables; we represent the possible values they may take on as $y \in \mathcal{Y}$, and represent the probability distribution of Y under D, D' as $\mathcal{P}$, $\mathcal{P}'$, respectively. It follows that

$$\mathbb{E}\left[Y(D)\right] - \mathbb{E}\left[Y(D')\right] = \mathbb{E}_{Y \sim \mathcal{P}}\, Y - \mathbb{E}_{Y \sim \mathcal{P}'}\, Y = \sum_{y \in \mathcal{Y}} \left(\Pr_{\mathcal{P}}[Y = y] - \Pr_{\mathcal{P}'}[Y = y]\right) y$$

Therefore,

$$\begin{aligned}|\mathbb{E}[Y(D)] - \mathbb{E}[Y(D')]| &\leq \sum_{y\in\mathcal{Y}} \left|\Pr_{\mathcal{P}}[Y=y] - \Pr_{\mathcal{P}'}[Y=y]\right| |y| \\ &\leq \sum_{y\in\mathcal{Y}} (e^{\varepsilon}-1)\max\left\{\Pr_{\mathcal{P}}[Y=y],\ \Pr_{\mathcal{P}'}[Y=y]\right\}|y| \\ &\leq (e^{\varepsilon}-1)K,\end{aligned}$$

where the second inequality follows from the definition of $(\varepsilon, 0)$-differential privacy. □

Proof of Claim 4. We prove the claim for any seller. A similar proof holds for buyers. Fix an index i, and any reports/bid vector $\left(\mathbf{r}^s_{-i}, \mathbf{r}^b\right)$ for the remaining buyers and sellers. For simplicity of notation, let us denote $(\mathbf{v}^s_i, \mathbf{r}^s_{-i}, \mathbf{r}^b)$ where i submits his bid truthfully as data set D, and $(\mathbf{r}^s_i, \mathbf{r}^s_{-i}, \mathbf{r}^b)$ for some (other) report $\mathbf{r}^s_i$ as data set D'. Notice D and D' are neighboring data sets. Writing $\mathbb{E}_{\mathcal{M}}$ for the expectation with respect to the mechanism $\mathcal{M}$, we have that:

$$\begin{aligned}\mathbb{E}_{\mathcal{M}}\left[\mathbf{u}^s_i(\mathcal{M}(D'))\right] &= \mathbb{E}_{\mathcal{M}}[\mathbf{a}^s_i \cdot (p - \mathbf{v}^s_i)|D'] \\ &= \mathbb{E}_{\mathcal{M}}\left[\mathbf{1}\left[p \geq \mathbf{r}^s_i\right] Bern(q^s)(p - \mathbf{v}^s_i)|D'\right] \\ &= \mathbb{E}_{\mathcal{M}}\left[\mathbf{1}\left[p \geq \mathbf{v}^s_i\right] \cdot \mathbf{1}\left[p \geq \mathbf{r}^s_i\right] Bern(q^s)(p - \mathbf{v}^s_i)|D'\right] \\ &+ \mathbb{E}_{\mathcal{M}}\left[\mathbf{1}\left[p < \mathbf{v}^s_i\right] \cdot \mathbf{1}\left[p \geq \mathbf{r}^s_i\right] Bern(q^s)(p - \mathbf{v}^s_i)|D'\right] \\ &\leq \mathbb{E}_{\mathcal{M}}\left[\mathbf{1}\left[p \geq \mathbf{v}^s_i\right] Bern(q^s)(p - \mathbf{v}^s_i)|D'\right] \\ &\leq \mathbb{E}_{\mathcal{M}}\left[\mathbf{1}\left[p \geq \mathbf{v}^s_i\right] Bern(q^s)(p - \mathbf{v}^s_i)|D\right] + \left(e^{3\varepsilon} - 1\right) V \\ &= \mathbb{E}_{\mathcal{M}}\left[\mathbf{u}^s_i(\mathcal{M}(D))\right] + \left(e^{3\varepsilon} - 1\right) V\end{aligned}$$

where the first inequality follows because the second term appearing in the sum is nonpositive and that $\mathbf{1}\left[p \geq \mathbf{v}^s_i\right] \cdot \mathbf{1}\left[p \geq \mathbf{r}^s_i\right] \leq \mathbf{1}\left[p \geq \mathbf{v}^s_i\right]$. The second inequality follows from Lemma 5 and the fact that the computation of the pair of random variables (p, q^s) combined with any post-processing of the pair (p, q^s) that is independent of the reported data $D' = (\mathbf{r}^s, \mathbf{r}^b)$ satisfies $(3\varepsilon, 0)$-differential privacy by the Post-processing Lemma 2 and the Composition Theorem 10. Also note that the price/bids range is $\{1, 2, \ldots, V\}$, so we can take $K = V$ in Lemma 5. □

The proofs of approximate truthfulness of Mechanisms 2 and 5 follow the exact same argument that leverages the stability properties of differential privacy. The only difference comes in the choice of tie-breaking rule and the level of differential privacy of Mechanisms 2 and 5. Rewriting the above proofs with the corresponding tie-breaking rules yields the argument.

G PROOFS FOR LEARNING DYNAMICS

G.1 Proof of No-Regret Lemma 1

We first show the claim below:

Claim 6. *Let $R_{j,t}$ be the random variable representing the reward of buyer j in Algorithm 4 at round t, and let $R^*_j(T)$ be the total reward of buyer j's best fixed action in hindsight, over T rounds. Moreover, let $\xi \leq V$ and $\eta \leq \frac{1}{V}$. Then, the regret of buyer j over T rounds is bounded as follows:*

$$R^*_j(T) - \mathbb{E}\left[\sum_{t=1}^{T} R_{j,t}\right] \leq \xi T + \eta V^2 T + \frac{\ln V}{\eta} \tag{23}$$

Proof. We can think of Algorithm 4 as Exponential Weights with a modified utility function:

$$\text{buyer } j\text{'s modified utility at time } t \text{ for bid } k: \quad \mu^b_{j,t}(k) = \begin{cases} \xi \cdot q^b_t & k = \mathbf{v}^b_j \text{ and } p_t = \mathbf{v}^b_j \\ u_{j,t}(k) & \text{otherwise} \end{cases}$$

where $u_{j,t}(k)$ is the actual utility of buyer j at time t if he were to bid k. Importantly, we show that using this modified utility function we can still achieve vanishing regret (with respect to the *original* reward $R_{j,t}$ which is the agent's *true/realized* utility).

First, notice that $u_{j,t}$ is always upper-bounded by $\mu^b_{j,t}$: $u_{j,t} \le \mu^b_{j,t}$; but also that $\mu^b_{j,t} \le u_{j,t} + \xi$. Recall $R^*_j(T)$ is the reward of the best fixed action in hindsight, with respect to the sequence of prices $p_1, \ldots, p_T$ and probabilities $q^b_1, \ldots, q^b_T$ as chosen by an adversary. Let r^*_j be the report that leads to achieving R^*_j, i.e.

$$R^*_j(T) \triangleq \max_{k \in \{1,\ldots,V\}} \sum_{t=1}^{T} u_{j,t}(k)$$

$$r^*_j \triangleq \operatorname*{argmax}_{k \in \{1,\ldots,V\}} \sum_{t=1}^{T} u_{j,t}(k)$$

Our goal will be to show that Equation (23) holds. Our proof technique will mostly follow standard arguments. In this proof – and this proof only – we let w denote the unnormalized weights that may not sum to 1, and note they induce probability distributions ρ by normalizing each weight by the sum of the weights.

First, let $W_{j,t} = \sum_{k=1}^{V} w_{j,t}(k)$. By definition:

$$\frac{W_{j,t+1}}{W_{j,t}} = \frac{\sum_{k=1}^{V} w_{j,t+1}(k)}{\sum_{k=1}^{V} w_{j,t}(k)} = \sum_{k=1}^{V} \frac{w_{j,t}(k) e^{\eta \mu^b_{j,t}(k)}}{\sum_{k=1}^{V} w_{j,t}(k)}$$

We will write $\rho_{j,t}(k) \triangleq w_{j,t}(k) / \sum_{k=1}^{V} w_{j,t}(k)$ as the probability distribution induced by weights $w_{j,t}(k)$, for all k. We can rewrite the above as

$$\frac{W_{j,t+1}}{W_{j,t}} = \sum_{k=1}^{V} \rho_{j,t}(k) e^{\eta \mu^b_{j,t}(k)}.$$

For $\eta \le \frac{1}{V}$ and $\xi \le 1$, we have $\eta \mu_{j,t}(k) \le 1$ for all k. Using the upper bound that $e^x \le 1 + x + x^2$ for all $x \in [0, 1]$, we obtain that

$$\frac{W_{j,t+1}}{W_{j,t}} \le 1 + \sum_{k=1}^{V} \rho_{j,t}(k) \cdot \eta \mu^b_{j,t}(k) + \sum_{k=1}^{V} \rho_{j,t}(k) \cdot \eta^2 \mu^b_{j,t}(k)^2$$

Then

$$\begin{aligned} \ln \frac{W_{j,t+1}}{W_{j,t}} &\le \ln \left(1 + \eta \sum_{k=1}^{V} \rho_{j,t}(k) \mu^b_{j,t}(k) + \eta^2 \sum_{k=1}^{V} \rho_{j,t}(k) \mu^b_{j,t}(k)^2 \right) \\ &\le \eta \sum_{k=1}^{V} \rho_{j,t}(k) \mu^b_{j,t}(k) + \eta^2 \sum_{k=1}^{V} \rho_{j,t}(k) \mu^b_{j,t}(k)^2, \end{aligned} \tag{24}$$

where we have used the fact that $\ln(1 + x) \le x$ for $x > -1$ (which holds in this case because payoffs are nonnegative given buyers (sellers) never bid above (below) their valuations). Now noting that

$\frac{W_{j,t+1}}{W_{j,1}} = \frac{W_{j,t+1}}{W_{j,t}} \frac{W_{j,t}}{W_{t-1}} \cdots \frac{W_{j,2}}{W_{j,1}}$, we can express

$$\ln \frac{W_{j,t+1}}{W_{j,1}} = \ln \frac{W_{j,t+1}}{W_{j,t}} \cdots \frac{W_{j,2}}{W_{j,1}} = \sum_{\tau=1}^{t} \ln \frac{W_{\tau+1,j}}{W_{\tau,j}}$$

And applying Inequality (24), we have that

$$\ln \frac{W_{j,t+1}}{W_{j,1}} \leq \eta \sum_{\tau=1}^{t} \sum_{k=1}^{V} \rho_{j,\tau}(k) \mu^b_{j,\tau}(k) + \eta^2 \sum_{\tau=1}^{t} \sum_{k=1}^{V} \rho_{j,\tau}(k) \mu^b_{j,\tau}(k)^2 \tag{25}$$

On the other hand, since $W_{j,t+1} \geq w_{j,t}(k)$ for all k, *including* for the best action in hindsight $k = r^*_j$, we have that

$$\begin{aligned}
\ln \frac{W_{j,t+1}}{W_{j,1}} \geq \ln \frac{w_{j,t+1}(r^*_j)}{W_{j,1}} &= \ln(e^{\eta \mu^b_{j,t}(r^*_j)} w_{j,t}(r^*_j)/W_{j,1}) \\
&= \ln(e^{\eta \mu^b_{j,t}(r^*_j)} e^{\eta \mu^b_{j,t-1}(r^*_j)} w_{j,t-1}(r^*_j)) - \ln W_{j,1} \\
&= \ldots. \\
&= \ln \left(\prod_{\tau=1}^{t} e^{\eta \mu^b_{j,\tau}(r^*_j)} w_{j,1}(r^*_j) \right) - \ln W_{j,1}.
\end{aligned}$$

Now, using the fact that the weights can be initialized with $w_{j,1}(k) = 1 \; \forall k$ and $W_{j,1} = V$, this gives

$$\ln \frac{W_{j,t+1}}{W_{j,1}} \geq \eta \sum_{\tau=1}^{t} \mu^b_{j,\tau}(r^*_j) - \ln V \tag{26}$$

But now combining Inequalities (25) and (26) gives:

$$\eta \sum_{\tau=1}^{t} \mu^b_{j,\tau}(r^*_j) - \ln V \leq \eta \sum_{\tau=1}^{t} \sum_{k=1}^{V} \rho_{j,\tau}(k) \mu^b_{j,\tau}(k) + \eta^2 \sum_{\tau=1}^{t} \sum_{k=1}^{V} \rho_{j,\tau}(k) \mu^b_{j,\tau}(k)^2$$

Now notice that $\sum_{k=1}^{V} \rho_{j,\tau}(k) \mu^b_{j,\tau}(k) = \mathbb{E}_k[\mu^b_{j,\tau}(k)]$. So rearranging and letting $t = T$, we have that

$$\sum_{\tau=1}^{T} \mu^b_{j,\tau}(r^*_j) - \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)] \leq \frac{\ln V}{\eta} + \eta \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)^2] \leq \frac{\ln V}{\eta} + \eta T V^2$$

where the inequality follows from the fact that $\mu^b_{j,t}$ is bounded by $\max(V, \xi) = V$ (remembering that $u_{j,t} \leq V$). But since $\mu^b_{j,\tau}(k) \geq u_{j,\tau}(k)$, we have that

$$\begin{aligned}
R^*_j(T) - \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)] &= \sum_{\tau=1}^{T} u_{j,\tau}(r^*_j) - \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)] \\
&\leq \sum_{\tau=1}^{T} \mu^b_{j,\tau}(r^*_j) - \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)] \\
&\leq \frac{\ln V}{\eta} + \eta T V^2
\end{aligned}$$

Further, since $\mu^b_{j,\tau}(k) \le u^b_{j,t}(k) + \xi$, we also have that

$$\begin{aligned} R^*_j(T) - \sum_{\tau=1}^{T} \mathbb{E}[\mu^b_{j,\tau}(k_{j,\tau})] &= \sum_{\tau=1}^{T} u_{j,\tau}(r^*_j) - \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)] \\ &\ge \sum_{\tau=1}^{T} \mu^b_{j,\tau}(r^*_j) - \sum_{\tau=1}^{T} \mathbb{E}_k[\mu^b_{j,\tau}(k)] - \xi T \\ &\ge \sum_{\tau=1}^{T} u_{j,\tau}(r^*_j) - \sum_{\tau=1}^{T} \mathbb{E}_k[u_{j,\tau}(k)] - \xi T \\ &= R^*_j(T) - \sum_{t=1}^{T} \mathbb{E}[R_{j,t}] - \xi T. \end{aligned}$$

Combining the last two inequalities, we get

$$R^*_j(T) - \sum_{t=1}^{T} \mathbb{E}\left[R_{j,t}\right] \le \xi T + \frac{\ln V}{\eta} + \eta V^2 T,$$

as desired. □

We can now conclude the proof, noting that Lemma 6 gives that the total regret of Algorithm 4 over T rounds for agent j is bounded by:

$$\text{Regret} \le \xi T + \frac{\ln V}{\eta} + \eta V^2 T$$

Choose $\eta = \frac{1}{V\sqrt{T}}$ and $\xi = \frac{1}{\sqrt{T}}$. Then we have that

$$\text{Regret} \le \sqrt{T} + V \ln V\sqrt{T} + \frac{1}{V\sqrt{T}} V^2 T = \sqrt{T} + V \ln V\sqrt{T} + V\sqrt{T}.$$

Then average regret can be bounded as:

$$\frac{1}{T}\,\text{Regret} \le \frac{1}{\sqrt{T}} + \frac{V \ln V}{\sqrt{T}} + \frac{V}{\sqrt{T}} = O\left(\frac{1}{\sqrt{T}}\right).$$

That is, average regret vanishes as $T \to \infty$.

G.2 Proof of Theorem 4

To prove Theorem 4, we will examine how the OPT′ sellers with the lowest values and the OPT′ buyers with the highest values update their weights. To do so, we will need the following definition:

Definition 11 (Highest (resp. lowest) value buyers (resp. sellers)). *Let $n^b(v) = \sum_{i=1}^{n^b} \mathbf{1}[\mathbf{v}^b_j \ge v]$ be the number of buyers with value bigger than or equal to v, and let $v^b = \max\{v : n^b(v) \ge OPT'\}$. Similarly, let $n^s(v) = \sum_{i=1}^{n^s} \mathbf{1}[\mathbf{v}^b_j \le v]$ be the number of sellers with value smaller than or equal to v, and let $v^s = \min\{v : n^s(v) \ge OPT'\}$.*

We note the following property of v^b, v^s:

Claim 7. *Suppose $OPT' > 0$. Then,*

$$v^b \ge v^s + 2.$$

Proof. By definition of OPT′, there exists a price $p^\star$ such that at least OPT′ buyers have value above or equal to $p^\star + 1$ and OPT′ sellers below or equal to $p^\star - 1$. But then, $v^s \le p^\star - 1$ and $v^b \ge p^\star + 1$, which concludes the proof. □

First of all, we show that if a given price p is picked infinitely many times, every agent j with $\mathrm{v}_j^b > p$ sees their probability of bidding more than p converge to 1. This is the object of Corollary 1, whose proof relies on Lemmas 6 and 7 below. We state the Lemmas for a buyer j and note that similar results hold for a seller i as well.

Lemma 6. *For all t, for all $p \in [V]$, for all $j \in [n^b]$,*

$$\sum_{k=p}^{\mathrm{v}_j^b} w_{j,t+1}^b(k) \geq \sum_{k=p}^{\mathrm{v}_j^b} w_{j,t}^b(k).$$

Proof. If $\sum_{k<p} w_{j,t}^b(k) = 0$, the result is immediate: it must be that for all $k < p$, $w_{j,t}^b(k) = 0$, so by exponential update, $w_{j,t+1}^b(k) = 0$, leading to $\sum_{k<p} w_{i,t+1}^b(k) = 0$. In turn,

$$\sum_{k\geq p} w_{j,t+1}^b(k) = \sum_{k\geq p} w_{j,t}^b(k) = 1.$$

We now focus on the case when $\sum_{k<p} w_{j,t}^b(k) > 0$. Remember that p_t is the optimal price at time t. If $p \leq p_t$,

$$\frac{\sum_{k\geq p} w_{j,t+1}^b(k)}{\sum_{k<p} w_{j,t+1}^b(k)} = \frac{\sum_{k=p}^{p_t-1} w_{j,t}^b(k) + \sum_{k\geq p_t} w_{j,t}^b(k)\exp(\eta q_t^b(\mathrm{v}_j^b - p_t))}{\sum_{k<p} w_{j,t}^b(k)} \geq \frac{\sum_{k\geq p} w_{j,t}^b(k)}{\sum_{k<p} w_{j,t}^b(k)}.$$

When $p > p_t$,

$$\begin{aligned}
\frac{\sum_{k\geq p} w_{j,t+1}^b(k)}{\sum_{k<p} w_{j,t+1}^b(k)} &= \frac{\sum_{k\geq p} w_{j,t}^b(k)\exp(\eta q_t^b(\mathrm{v}_j^b - p_t))}{\sum_{k<p_t} w_{j,t}^b(k) + \sum_{k=p_t}^{p-1} w_{j,t}^b(k)\exp(\eta q_t^b(\mathrm{v}_j^b - p_t))} \\
&\geq \frac{\sum_{k\geq p} w_{j,t}^b(k)\exp(\eta q_t^b(\mathrm{v}_j^b - p_t))}{\left(\sum_{k<p_t} w_{j,t}^b(k) + \sum_{k=p_t}^{p-1} w_{j,t}^b(k)\right)\exp(\eta q_t^b(\mathrm{v}_j^b - p_t))} \\
&= \frac{\sum_{k\geq p} w_{j,t}^b(k)}{\sum_{k<p} w_{j,t}^b(k)}.
\end{aligned}$$

Since

$$\sum_{k\geq p} w_{j,t+1}^b(k) + \sum_{k<p} w_{j,t+1}^b(k) = 1, \sum_{k\geq p} w_{j,t}^b(k) + \sum_{k<p} w_{j,t}^b(k) = 1,$$

we have that for all p,

$$\frac{\sum_{k\geq p} w_{j,t+1}^b(k)}{1 - \sum_{k\geq p} w_{j,t+1}^b(k)} \geq \frac{\sum_{k\geq p} w_{j,t}^b(k)}{1 - \sum_{k\geq p} w_{j,t}^b(k)}.$$

This in particular implies that for all p,

$$\sum_{k\geq p} w_{j,t+1}^b(k)\left(1 - \sum_{k\geq p} w_{j,t}^b(k)\right) \geq \sum_{k\geq p} w_{j,t}^b(k)\left(1 - \sum_{k\geq p} w_{j,t+1}^b(k)\right),$$

hence

$$\sum_{k\geq p} w_{j,t+1}^b(k) \geq \sum_{k\geq p} w_{j,t}^b(k).$$

□

Lemma 7 (Update moves mass up by a constant amount). *Suppose at time t, at least one buyer and one seller can trade.*

There exists a constant $C(\varepsilon) > 1$ such that for any buyer j with $\mathbf{v}_j^b > p_t$ and $\sum_{k=p_t}^{\mathbf{v}_j^b} w_{j,t}^b(k) \leq 1 - \varepsilon$, we have that

$$\frac{\sum_{k=p_t}^{\mathbf{v}_j^b} w_{j,t+1}^b(k)}{\sum_{k=p_t}^{\mathbf{v}_j^b} w_{j,t}^b(k)} \geq C(\varepsilon).$$

Proof. Let $X_t(p)$ be the probability that buyer j bids at least p on round t. For simplicity of notations, we omit the j subscripts in the proof.

Trivially:

$$X_t(p_t) = \sum_{k=p_t}^{\mathbf{v}_j^b} w_{j,t}^b(k).$$

Now, by the definition of exponential weights, we have that

$$X^{t+1}(p_t) = \frac{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} X_t(p_t)}{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} X_t(p_t) + (1 - X_t(p_t))}$$

since the buyer updates $w_{j,t}^b(k)$ with $e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}$ for all bids k above p_t up to $\mathbf{v}_j^b$, and updates weights on bids $k \leq p_t$ with $e^{\eta q_t^b \cdot 0} = 1$. It immediately follows that

$$\frac{X^{t+1}(p_t)}{X_t(p_t)} = \frac{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}}{X_t(p_t)(e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} - 1) + 1}$$

Now by assumption, $X_t(p_t) < 1 - \varepsilon$, so

$$\begin{aligned}
\frac{X^{t+1}(p_t)}{X_t(p_t)} &> \frac{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}}{(1-\varepsilon)(e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} - 1) + 1} \\
&= \frac{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}}{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} - \varepsilon e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} + \varepsilon} \\
&= \frac{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}}{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)} + \varepsilon(1 - e^{\eta q_t^b(\mathbf{v}_j^b - p_t)})} \\
&= \frac{1}{1 - \varepsilon(1 - \frac{1}{e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}})}
\end{aligned}$$

Using the fact that $q_t^b \geq \frac{1}{n^b}$, as there are at most n^b buyers and at least one possible seller to trade with, and the fact that $\mathbf{v}_j^b - p_t \geq 1$, we get that

$$e^{\eta/n^b} \leq e^{\eta q_t^b(\mathbf{v}_j^b - p_t)}.$$

In turn,

$$\frac{X^{t+1}(p_t)}{X_t(p_t)} \geq \frac{1}{1 - \varepsilon(1 - e^{-\eta/n^b})} > 1.$$

Letting $C(\varepsilon) = \frac{1}{1-\varepsilon(1-e^{-\eta/n^b})}$ is enough to conclude the proof. □

Corollary 1. *Pick any buyer j, and let $p < \mathbf{v}_j^b$. Let $N_t(p)$ be the number of times price p is picked by the mechanism so that at least one trade is possible at p, up until time t. In other words,*

$$N_t(p) = \sum_{t' \le t} \mathbf{1}\left[\Pi_{t'}\left(p, \mathbf{r}_{t'}^s, \mathbf{r}_{t'}^b\right) \ge 1\right]$$

If $\lim_{t\to\infty} N_t(p) = +\infty$, then

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{j,t}^b \ge p] = 1$$

Proof. Fix $\varepsilon > 0$. At time t, by applying Lemma 7 and Lemma 6 repeatedly, we have that

$$\begin{aligned}\Pr[\mathbf{r}_{j,t}^b \ge p] &\ge \min\{1-\varepsilon, C(\varepsilon)^{N_t(p)} \Pr[\mathbf{r}_{j,0}^b \ge p]\} \\ &\ge \min\{1-\varepsilon, C(\varepsilon)^{N_t(p)} \frac{1}{V}\}\end{aligned}$$

where the last inequality follows because the initial weights are uniform over all bids. By assumption, there exists T such that for all $t \ge T$: $N_t(p) \ge \frac{\log((1-\varepsilon)(V))}{\log C(\varepsilon)}$, and consequently,

$$\Pr[\mathbf{r}_{j,t}^b \ge p] \ge 1 - \varepsilon.$$

Since this holds for every $\varepsilon > 0$, the limit statement follows. □

We note that a similar Corollary exists for sellers as well. Now, we need to show that there is a price that clears benchmark OPT′ and is chosen by the mechanism infinitely often. This is the object of Lemma 8, whose proof relies on Claim 8. Once again, we state the Claim only for buyers and note that a similar result for sellers as well.

Claim 8. *For every buyer j, for all t, $\frac{1}{V} \le w_{j,t}^b(\mathbf{v}_j^b) \le \frac{1}{2}$.*

Proof. At time step t, if $p_t > \mathbf{v}_j^b$, agent j does not update any weight. If $p_t \le \mathbf{v}_j^b$, it is easy to see that the weight on $\mathbf{v}_j^b$ cannot decrease in the next round. Indeed, for any k such that $p_t \le k \le \mathbf{v}_j^b$, we have that

$$\begin{aligned}w_{j,t+1}^b(k) &= w_{j,t}^b(k) \cdot \frac{\exp\left(\eta q_t^b(\mathbf{v}_j^b - p_t)\right)}{\sum_{k<p_t} w_{j,t}^b(k) + \sum_{k\ge p_t} w_{j,t}^b(k)\exp\left(\eta q_t^b(\mathbf{v}_j^b - p_t)\right)} \\ &= w_{j,t}^b(k) \cdot \frac{1}{\exp\left(-\eta q_t^b(\mathbf{v}_j^b - p_t)\right)\sum_{k<p_t} w_{j,t}^b(k) + \sum_{k\ge p_t} w_{j,t}^b(k)} \\ &= w_{j,t}^b(k) \cdot \frac{1}{\exp\left(-\eta q_t^b(\mathbf{v}_j^b - p_t)\right)\sum_{k<p_t} w_{j,t}^b(k) + 1 - \sum_{k<p_t} w_{j,t}^b(k)} \\ &= w_{j,t}^b(k) \cdot \frac{1}{1 - \left(1 - \exp\left(-\eta q_t^b(\mathbf{v}_j^b - p_t)\right)\right)\sum_{k<p_t} w_{j,t}^b(k)} \\ &\ge w_{j,t}^b(k),\end{aligned}$$

where the last step follows from noting that both $1 - \exp\left(-\eta q_t^b(\mathbf{v}_j^b - p_t)\right)$, $\sum_{k<p_t} w_{j,t}^b(k) \le 1$. As such, $w_{j,t}^b(\mathbf{v}_j^b)$ is non-decreasing in t, so $w_{j,t}^b(\mathbf{v}_j^b) \ge w_{j,0}^b(\mathbf{v}_j^b) = \frac{1}{V}$.

Let us now prove the second inequality. Note that at any time step t, let p_t be the price chosen by the mechanism. When $p_t > \mathbf{v}_j^b$, j does not update his weight. Similarly, when $p_t = \mathbf{v}_j^b$, the exponential update rule is the same for $w_{j,t}^b(\mathbf{v}_j^b)$ and $w_{j,t}^b(\mathbf{v}_j^b - 1)$ and given by $\exp\left(\eta q_t^b(\mathbf{v}_j^b - p_t)\right) =$

$\exp(0) = 1$. When $p_t < \mathbf{v}_j^b$, both $w_{j,t}^b(\mathbf{v}_j^b)$ and $w_{j,t}^b(\mathbf{v}_j^b - 1)$ are multiplied by the same amount $\exp\left(\eta q_t^b(\mathbf{v}_j^b - p_t)\right)$. Therefore, it immediately follows by induction that $w_{j,t}^b(\mathbf{v}_j^b) = w_{j,t}^b(\mathbf{v}_j^b - 1)$ for all t. In particular, this implies $w_{j,t}^b(\mathbf{v}_j^b) \leq 1/2$, as $w_{j,t}^b(\mathbf{v}_j^b) + w_{j,t}^b(\mathbf{v}_j^b - 1) \leq 1$. □

Lemma 8 (Good event). *Suppose* $OPT' > 0$, *and let*

$$\gamma \triangleq \left(\frac{1}{V}\right)^{1+|n^b(v^s+1)||n^s(v^b-1)|} \cdot \left(\frac{1}{2}\right)^{\left(n^b-|n^b(v^s+1)|\right)\left(n^s-|n^s(v^b-1)|\right)} > 0.$$

At any time t, $v^s < p_t < v^b$ *and at least one trade is possible with probability at least* γ.

Proof. By Claim 8, we have that with probability at least

$$\left(\frac{1}{V}\right)^{|n^b(v^s+1)||n^s(v^b-1)|} \cdot \left(\frac{1}{2}\right)^{\left(n^b-|n^b(v^s+1)|\right)\left(n^s-|n^s(v^b-1)|\right)} = V\gamma,$$

all buyers with value $\mathbf{v}_j^b > v^s$ bid their value, all buyers with value $\mathbf{v}_j^b \leq v^s$ bid strictly below their value, all sellers with value $\mathbf{v}_i^s < v^b$ bid their value, and all sellers with $\mathbf{v}_i^s \geq v^b$ bid strictly more than their value. In particular, since $v^s < v^b$, all buyers with value $\mathbf{v}_j^b \geq v^b$ bid their value and all sellers with value $\mathbf{v}_i^s \leq v^s$ bid their value. By definition of v^b and v^s, there are at least OPT′ such buyers and sellers, so setting any price p satisfying $v^s \leq p \leq v^b$ clears OPT′ shares at least. On the other hand, any price $p > v^b$ and any price $p < v^s$ cannot clear OPT′ shares. Therefore, $v^s \leq p_t \leq v^b$. Further, since all buyers with value $\mathbf{v}_j^b \geq v^b$ and all sellers with value $\mathbf{v}_i^s \leq v^s$ bid their values, and $v^b \geq v^s$, at least $\text{OPT}' \geq 1$ trades happen at price p_t.

When $v^s < p < v^b$ for all optimal prices, this is enough to conclude the proof. Now, suppose $p = v^b$ is an optimal price at time t. By construction, no seller bids v^b. As such, the number of sellers with bids under p and the number of sellers with bids under $p - 1$ are the same, and $p - 1 = v^b - 1$ clears at least as many shares as p, hence is optimal at time t. Because p_t is chosen uniformely at random among the set of optimal prices, and there are at most V optimal prices, $p - 1$ is picked with probability at least $\frac{1}{V}$, and satisfies $v^s < p - 1 < v^b$ by Claim 7. Similarly, if $p = v^s$ is optimal, then so is $p + 1 < v^b$, and it is picked by the mechanism with probability at least $\frac{1}{V}$. This concludes the proof. □

We are now ready to put everything together, and show Theorem 4.

Proof of Theorem 4. The case when $\text{OPT}' = 0$ is immediate. So let us assume $\text{OPT}' > 0$. Lemma 8 shows that at any given round, there is a constant probability $\gamma > 0$ to pick $p_t \in (v^s, v^b)$ and realize at least one trade at that price. As such, as $t \to +\infty$, the number of times the mechanism picks a price in (v^s, v^b) such that a trade is realized also tends to infinity. In particular, by the pigeonhole principle, there exists a price $p^\star \in (v^s, v^b)$ such that

$$\lim_{t\to\infty} N_t(p^\star) = +\infty.$$

By Corollary 1, for every buyer $j \in n^b(v^b)$,

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{j,t}^b \geq p^\star] = 1,$$

and similarly, for every seller $i \in n^s(v^s)$,

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{i,t}^s \leq p^\star] = 1.$$

Since there are at least OPT′ buyers in $n^b(v^b)$ and OPT′ sellers in $n^s(v^s)$, we have that

$$1 \geq \Pr\left[\Pi_t\left(p_t, \mathbf{r}^s, \mathbf{r}^b\right) \geq \mathrm{OPT}'\right] \geq \prod_{j \in n^b(v^b)} \Pr[\mathbf{r}^b_{j,t} \geq p^\star] \cdot \prod_{i \in n^s(v^s)} \Pr[\mathbf{r}^s_{i,t} \leq p^\star],$$

which concludes the proof. □

G.3 Proof of Theorem 5

The proof is similar to that of Theorem 4, and is given below. We start by showing in Corollary 2 that if a price p is picked by the mechanism infinitely many times, every buyer with value at least p learns to bid higher than p with probability going to 1.

Lemma 9. *For all t, for all $p \in [V]$, for all buyers j,*

$$\sum_{k=p}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t+1}(k) \geq \sum_{k=p}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t}(k).$$

Proof. The proof is identical to that of Lemma 6. □

We then characterize by how much the weight allocated to bids above the chosen price p_t increase for a buyer j, at every time step t:

Lemma 10 (Update moves mass up by a constant amount). *Suppose at time t, at least one buyer and one seller can trade. There exists a constant $C(\varepsilon) > 1$ such that for any buyer j with $\mathbf{v}^b_j \geq p_t$ and $\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t}(k) \leq 1 - \varepsilon$, we have that*

$$\frac{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t+1}(k)}{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t}(k)} \geq C(\varepsilon).$$

Proof. Note that when $p_t < \mathbf{v}^b_j$, we have by Lemma 7 that

$$\frac{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t+1}(k)}{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t}(k)} \geq \frac{1}{1 - \varepsilon(1 - e^{-\eta/n^b})} > 1.$$

Now, when $p_t = \mathbf{v}^b_j$ note that

$$\frac{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t+1}(k)}{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t}(k)} = \frac{\mathbf{w}^b_{j,t+1}(\mathbf{v}^b_j)}{\mathbf{w}^b_{j,t}(\mathbf{v}^b_j)} = \exp\left(\eta q^b_t \xi\right).$$

In particular, as there is at least one possible trade, we have that $q^b_t \geq 1/n^b$, hence

$$\frac{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t+1}(k)}{\sum_{k=p_t}^{\mathbf{v}^b_j} \mathbf{w}^b_{j,t}(k)} \geq \exp\left(\frac{\eta\xi}{n^b}\right).$$

Letting $C(\varepsilon) = \min\left(\frac{1}{1-\varepsilon(1-e^{-\eta/n^b})}, \exp\left(\frac{\eta\xi}{n^b}\right)\right)$ is enough to conclude the proof. □

Corollary 2. *Pick any buyer j, and let $p \le \mathbf{v}_j^b$. Let $N_t(p)$ be the number of times price p is picked and at least one trade is possible at price p, up until time t. If $\lim_{t\to\infty} N_t(p) = +\infty$, then*

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{j,t}^b \ge p] = 1$$

Proof. This is identical to the proof of Corollary 1. □

Second, we need to show that there is a price that clears benchmark OPT′ and is chosen by the mechanism infinitely often.

Lemma 11 (Good event). *With probability at least $\left(\frac{1}{V}\right)^{n^b+n^s}$, all agents bid their valuation.*

Proof. By the same proof as Corollary 8, for every agent j and for all t, $\frac{1}{V} \le \mathbf{w}_{j,t}^b(\mathbf{v}_j^b)$. This is enough to prove the lemma. □

We are now ready to put everything together, and show Theorem 5.

Proof of Theorem 5. Suppose OPT > 0 (otherwise the result is immediate). When all agents bid their values, the mechanism selects a price that executes OPT ≥ 1 trades. Lemma 11 shows this happens with constant probability at any given round, and as such happens infinitely often when the number of rounds goes to infinity. By the pigeonhole principle, there exists a price $p^\star$ such that there are at least OPT buyers (resp. sellers) with value at least (resp. at most) $p^\star$, and such that

$$\lim_{t\to\infty} N_t(p^\star) = +\infty.$$

By Corollary 2, for any buyer with $\mathbf{v}_j^b \ge p^\star$,

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{j,t}^b \ge p^\star] = 1,$$

and similarly, for every seller i with $\mathbf{v}_i^s \le p^\star$,

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{i,t}^s \le p^\star] = 1.$$

In turn, since

$$1 \ge \Pr\left[\Pi_t\left(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b\right) \ge \mathrm{OPT}\right] \ge \prod_{j\in[n^b]:\ \mathbf{v}_j^b \ge p^\star} \Pr[\mathbf{r}_{j,t}^b \ge p^\star] \cdot \prod_{i\in[n^s]:\ \mathbf{v}_i^s \le p^\star} \Pr[\mathbf{r}_{i,t}^s \le p^\star],$$

we have

$$\lim_{t\to+\infty} \Pr\left[\Pi_t\left(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b\right) \ge \mathrm{OPT}\right] = 1.$$

□

G.4 Proof of Theorem 6

We start by noting that in the private case, the weights are still non-decreasing over time.

Lemma 12. *For all t, for all $p \in [V]$, for all buyers j,*

$$\sum_{k=p}^{\mathbf{v}_j^b} \mathbf{w}_{j,t+1}^b(k) \ge \sum_{k=p}^{\mathbf{v}_j^b} \mathbf{w}_{j,t}^b(k).$$

One distinction compared to the non-private case arises with respect to the amount by which the weights above p_t are updated. This amount depends on q_t^b, which is a random variable over the randomness of private computation of the selection probability. We note that *conditionally on* $q_t^b \ge 1/n^b$, Lemma 7 carries through, as formalized below:

Lemma 13 (Update moves mass up by a constant amount). *Suppose at time t, at least one buyer and one seller can trade and that $q_t^b > 1/n^b$. There exists a constant $C(\varepsilon) > 1$ such that for any buyer j with $\mathbf{v}_j^b \geq p_t$ and $\sum_{k=p_t}^{\mathbf{v}_j^b} \mathbf{w}_{j,t}^b(k) \leq 1 - \varepsilon$,*

$$\frac{\sum_{k=p_t}^{\mathbf{v}_j^b} \mathbf{w}_{j,t+1}^b(k)}{\sum_{k=p_t}^{\mathbf{v}_j^b} \mathbf{w}_{j,t}^b(k)} \geq C(\varepsilon).$$

We now fix a price p. We show that for one such p, if the event where p is the price picked by the mechanism *and* $q_t^b \geq 1/n$ happens infinitely often, then all bidders with valuation equal to or larger than p learn to bid higher than p with probability that goes to 1.

Lemma 14. *Pick any buyer j, and let $p \leq \mathbf{v}_j^b$. Let $N_t(p)$ be the number of times price p is picked by the mechanism so that at least one trade is possible at p and $q^b > 1/n^b$, up until time t. In other words,*

$$N_t(p) = \sum_{t' \leq t} \mathbf{1}\left[\Pi_{t'}\left(p, \mathbf{r}_{t'}^s, \mathbf{r}_{t'}^b\right) \geq 1, q_{t'}^b > \frac{1}{n^b}\right]$$

If $\lim_{t\to\infty} N_t(p) = +\infty$, then $\lim_{t\to\infty} \Pr[\mathbf{r}_{j,t}^b \geq p] = 1$.

We note that the event in which all agents bid their valuation and the mechanism (despite the randomness due to privacy) picks an optimal price p and releases $q^b \geq 1/n^b$, $q^s \geq 1/n^s$ happens with at least constant probability (independent of the time dimension of the problem), hence infinitely many times when the time horizon goes to infinity:

Lemma 15 (Good event). *Suppose OPT ≥ 1. At any round t, with probability at least $C\left(\frac{1}{V}\right)^{n^b+n^s+1}$ for some constant $C > 0$: all buyers bid their valuations, $q_t^b > 1/n^b$, $q_t^s > 1/n^s$, and the chosen price p_t is an optimal price that clears OPT shares.*

Proof. We have shown before in the proof of Lemma 11 that with probability at least $V^{-(n^b+n^s)}$ every agent bids their valuation.

In the rest of the proof, we condition on all agents bidding their valuations in the current round t. Conditional on this, we show that with constant probability, simultaneously: $q_t^b > 1/n^b$, and $q_t^s > 1/n^s$. Recall from Algorithm 1 that in each round t, given the selected price p_t, we have

$$q_t^b = \min\left(1, \frac{(\widehat{s}_t)_+}{\left(\widehat{b}_t - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right),\ q_t^s = \min\left(1, \frac{\left(\widehat{b}_t\right)_+}{\left(\widehat{s}_t - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right).$$

By the accuracy guarantees of the Laplace mechanism and the fact that Laplace noise has positive value with probability $1/2$, we have that with constant probability C (for some C that only depends on α and ε but not on t), the 4 following events simultaneously hold:

(1) $\left(\widehat{b}_t - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+ \leq \sum_{j\in\mathcal{B}} \mathbb{1}\left[\mathbf{v}_j^b \geq p_t\right] \leq n^b$,
(2) $\widehat{b}_t \geq \sum_{j\in B} \mathbb{1}\left[\mathbf{v}_j^b \leq p_t\right] \geq \text{OPT} \geq 1$,
(3) $\left(\widehat{s}_t - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+ \leq \sum_{i\in\mathcal{S}} \mathbb{1}\left[\mathbf{v}_i^s \leq p_t\right] \leq n^s$,
(4) $\widehat{s}_t \geq \sum_{i\in\mathcal{S}} \mathbb{1}\left[\mathbf{v}_i^s \leq p_t\right] \geq \text{OPT} \geq 1$, noting that at least one trade is possible at price p_t.

Using the above inequalities, we obtain that with probability C,

$$q_t^b = \min\left(1, \frac{(\widehat{s}_t)_+}{\left(\widehat{b}_t - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right) \geq \min\left(1, \frac{1}{n^b}\right) \geq \frac{1}{n^b}$$

$$q_t^s = \min\left(1, \frac{\left(\widehat{b}_t\right)_+}{\left(\widehat{s}_t - \frac{\ln(1/\alpha)}{\varepsilon}\right)_+}\right) \geq \min\left(1, \frac{1}{n^s}\right) \geq \frac{1}{n^s}.$$

To finish the proof, we just need to show that, conditional on all agents bidding their valuations in the current round t, with probability at least $1/V$, p_t – the price selected when every agent bids their valuation – is an optimal price. Note there exists a price $p_t^\star$ that is optimal for round t, i.e. such that $\Pi_t(p_t^\star, \mathbf{v}^s, \mathbf{v}^b) \geq \Pi_t(p, \mathbf{v}^s, \mathbf{v}^b)$ for all p. By the exponential mechanism, this price $p_t^\star$ is selected with probability

$$\frac{\exp\left(\varepsilon\Pi_t(p_t^\star, \mathbf{v}^s, \mathbf{v}^b)/2\right)}{\sum_{p=1}^V \exp\left(\varepsilon\Pi_t(p, \mathbf{v}^s, \mathbf{v}^b)/2\right)} \geq \frac{\exp\left(\varepsilon\Pi_t(p_t^\star, \mathbf{v}^s, \mathbf{v}^b)/2\right)}{\sum_{p=1}^V \exp\left(\varepsilon\Pi_t(p_t^\star, \mathbf{v}^s, \mathbf{v}^b)/2\right)} = \frac{1}{V}.$$

□

We are now ready to put everything together, and show Theorem 6.

Proof of Theorem 6. Let us for simplicity call:

$$f(\varepsilon, \alpha) \triangleq \frac{2\ln(V/\alpha)}{\varepsilon} - \frac{2\ln(1/\alpha)}{\varepsilon} - \sqrt{6\left(\mathrm{OPT} + \frac{\ln(1/\alpha)}{\varepsilon}\right)\ln(1/\alpha)}$$

Suppose OPT > 0 (otherwise the result is immediate). By Lemma 15 that shows that with constant probability (independent of time) in every round, the mechanism picks an optimal price and $q^b, q^s \geq \frac{1}{n}$, this event must happen infinitely many times. By the pigeonhole principle, there exists an optimal price $p^\star$ such that infinitely many times, $p^\star$ is picked by the mechanism with $q^b, q^s \geq \frac{1}{n}$. In turn, all buyers j with $\mathbf{v}_j^b \geq p^\star$ and all sellers i with $\mathbf{v}_i^s \leq p^\star$ (there are at least OPT of them, since $p^\star$ is optimal) learn to bid above, respectively below price $p^\star$ with probability that tends to 1 as t goes to infinity, by Lemma 14. Formally, for every buyer j with $\mathbf{v}_j^b \geq p^\star$,

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{j,t}^b \geq p^\star] = 1,$$

and similarly, for every seller i with $\mathbf{v}_i^s \leq p^\star$, we have that

$$\lim_{t\to\infty} \Pr[\mathbf{r}_{i,t}^s \leq p^\star] = 1.$$

Hence

$$\lim_{t\to\infty} \prod_{j\in[n^b]:\ \mathbf{v}_j^b \geq p^\star} \Pr[\mathbf{r}_{j,t}^b \geq p^\star] \cdot \prod_{i\in[n^s]:\ \mathbf{v}_i^s \leq p^\star} \Pr[\mathbf{r}_{i,t}^s \leq p^\star] = 1$$

and consequently, there exists $N(\alpha)$ large enough such that for all $t \geq N(\alpha)$,

$$\prod_{j\in[n^b]:\ \mathbf{v}_j^b \geq p^\star} \Pr[\mathbf{r}_{j,t}^b \geq p^\star] \cdot \prod_{i\in[n^s]:\ \mathbf{v}_i^s \leq p^\star} \Pr[\mathbf{r}_{i,t}^s \leq p^\star] \geq 1 - \alpha.$$

When all buyers with value at least the price and sellers with value at most the price bid between their valuation and $p^\star$, the optimal number of shares that can be cleared is OPT. By the accuracy guarantee of Mechanism 1, it must then be the case that for all $t \geq N(\alpha)$,

$$\begin{aligned}
\Pr\left[\Pi_t(p_t, \mathbf{r}_t^s, \mathbf{r}_t^b) \geq \mathrm{OPT} - f(\varepsilon, \alpha)\right] &\geq (1 - 8\alpha) \prod_{j \in [n^b]:\, \mathbf{v}_j^b \geq p^\star} \Pr[\mathbf{r}_{j,t}^b \geq p^\star] \cdot \prod_{i \in [n^s]:\, \mathbf{v}_i^s \leq p^\star} \Pr[\mathbf{r}_{i,t}^s \leq p^\star] \\
&\geq (1 - 8\alpha)(1 - \alpha) \\
&\geq 1 - 9\alpha.
\end{aligned}$$

This concludes the proof. □

The proof for benchmark OPT' follows the same argument, and is omitted for simplicity of exposition.