Differentially Private Call Auctions and Market Impact

Michael Kearns
University of Pennsylvania

IAQF/Thalesians Seminar
June 28, 2022

Joint work with Emily Diana, Hadi Elzayn, Aaron Roth, Saeed Sharifi-Malvajerdi, Juba Ziani

Ad-hoc defenses preserve privacy and limit market impact…

- Front-Running
- Fill-or-Kill
- Iceberg/Hidden/MinQty Orders
- Heartbeat Detection/Avoidance
- Dark Pools
- Speed Bumps
- “Whack-a-Mole”

Ad-hoc defenses preserve privacy and limit market impact…

- Front-Running
- Fill-or-Kill
- Iceberg/Hidden/MinQty Orders
- Heartbeat Detection/Avoidance
- Dark Pools
- Speed Bumps
- “Whack-a-Mole”

Can interpret all of these as attempts to preserve/compromise *privacy* of information

Can we achieve privacy in a market setting without sacrificing simplicity?

Call Auctions

- A call auction is a mechanism for allocating multiple identical items from sellers to buyers at a uniform price
- Large markets - NYSE, NASDAQ - run call auctions at opening and closing every day
- Some exchanges (IEX) run frequent call auctions (a la Budish, Cramton and Shin 2015)
- Two main purposes:
  - Allocate auctioned items to participants who value them
  - Price discovery

Private Call Auctions

- Want a mechanism with *provable, gracefully degrading* guarantees of privacy
- Still clear as many shares as possible and permit efficient price discovery
- Want cost of mechanism to be low. In our case, small net position (inventory)
- Want the mechanism to be simple for participants, and ideally incentive compatible

Call Auction

Valuation	Participant Type
1	Buyer
2	Buyer
2	Buyer
5	Buyer
2	Seller
3	Seller
3	Seller
7	Seller

Call Auction

Valuation	Participant Type
1	Buyer
2	Buyer
2	Buyer
5	Buyer
2	Seller
3	Seller
3	Seller
7	Seller

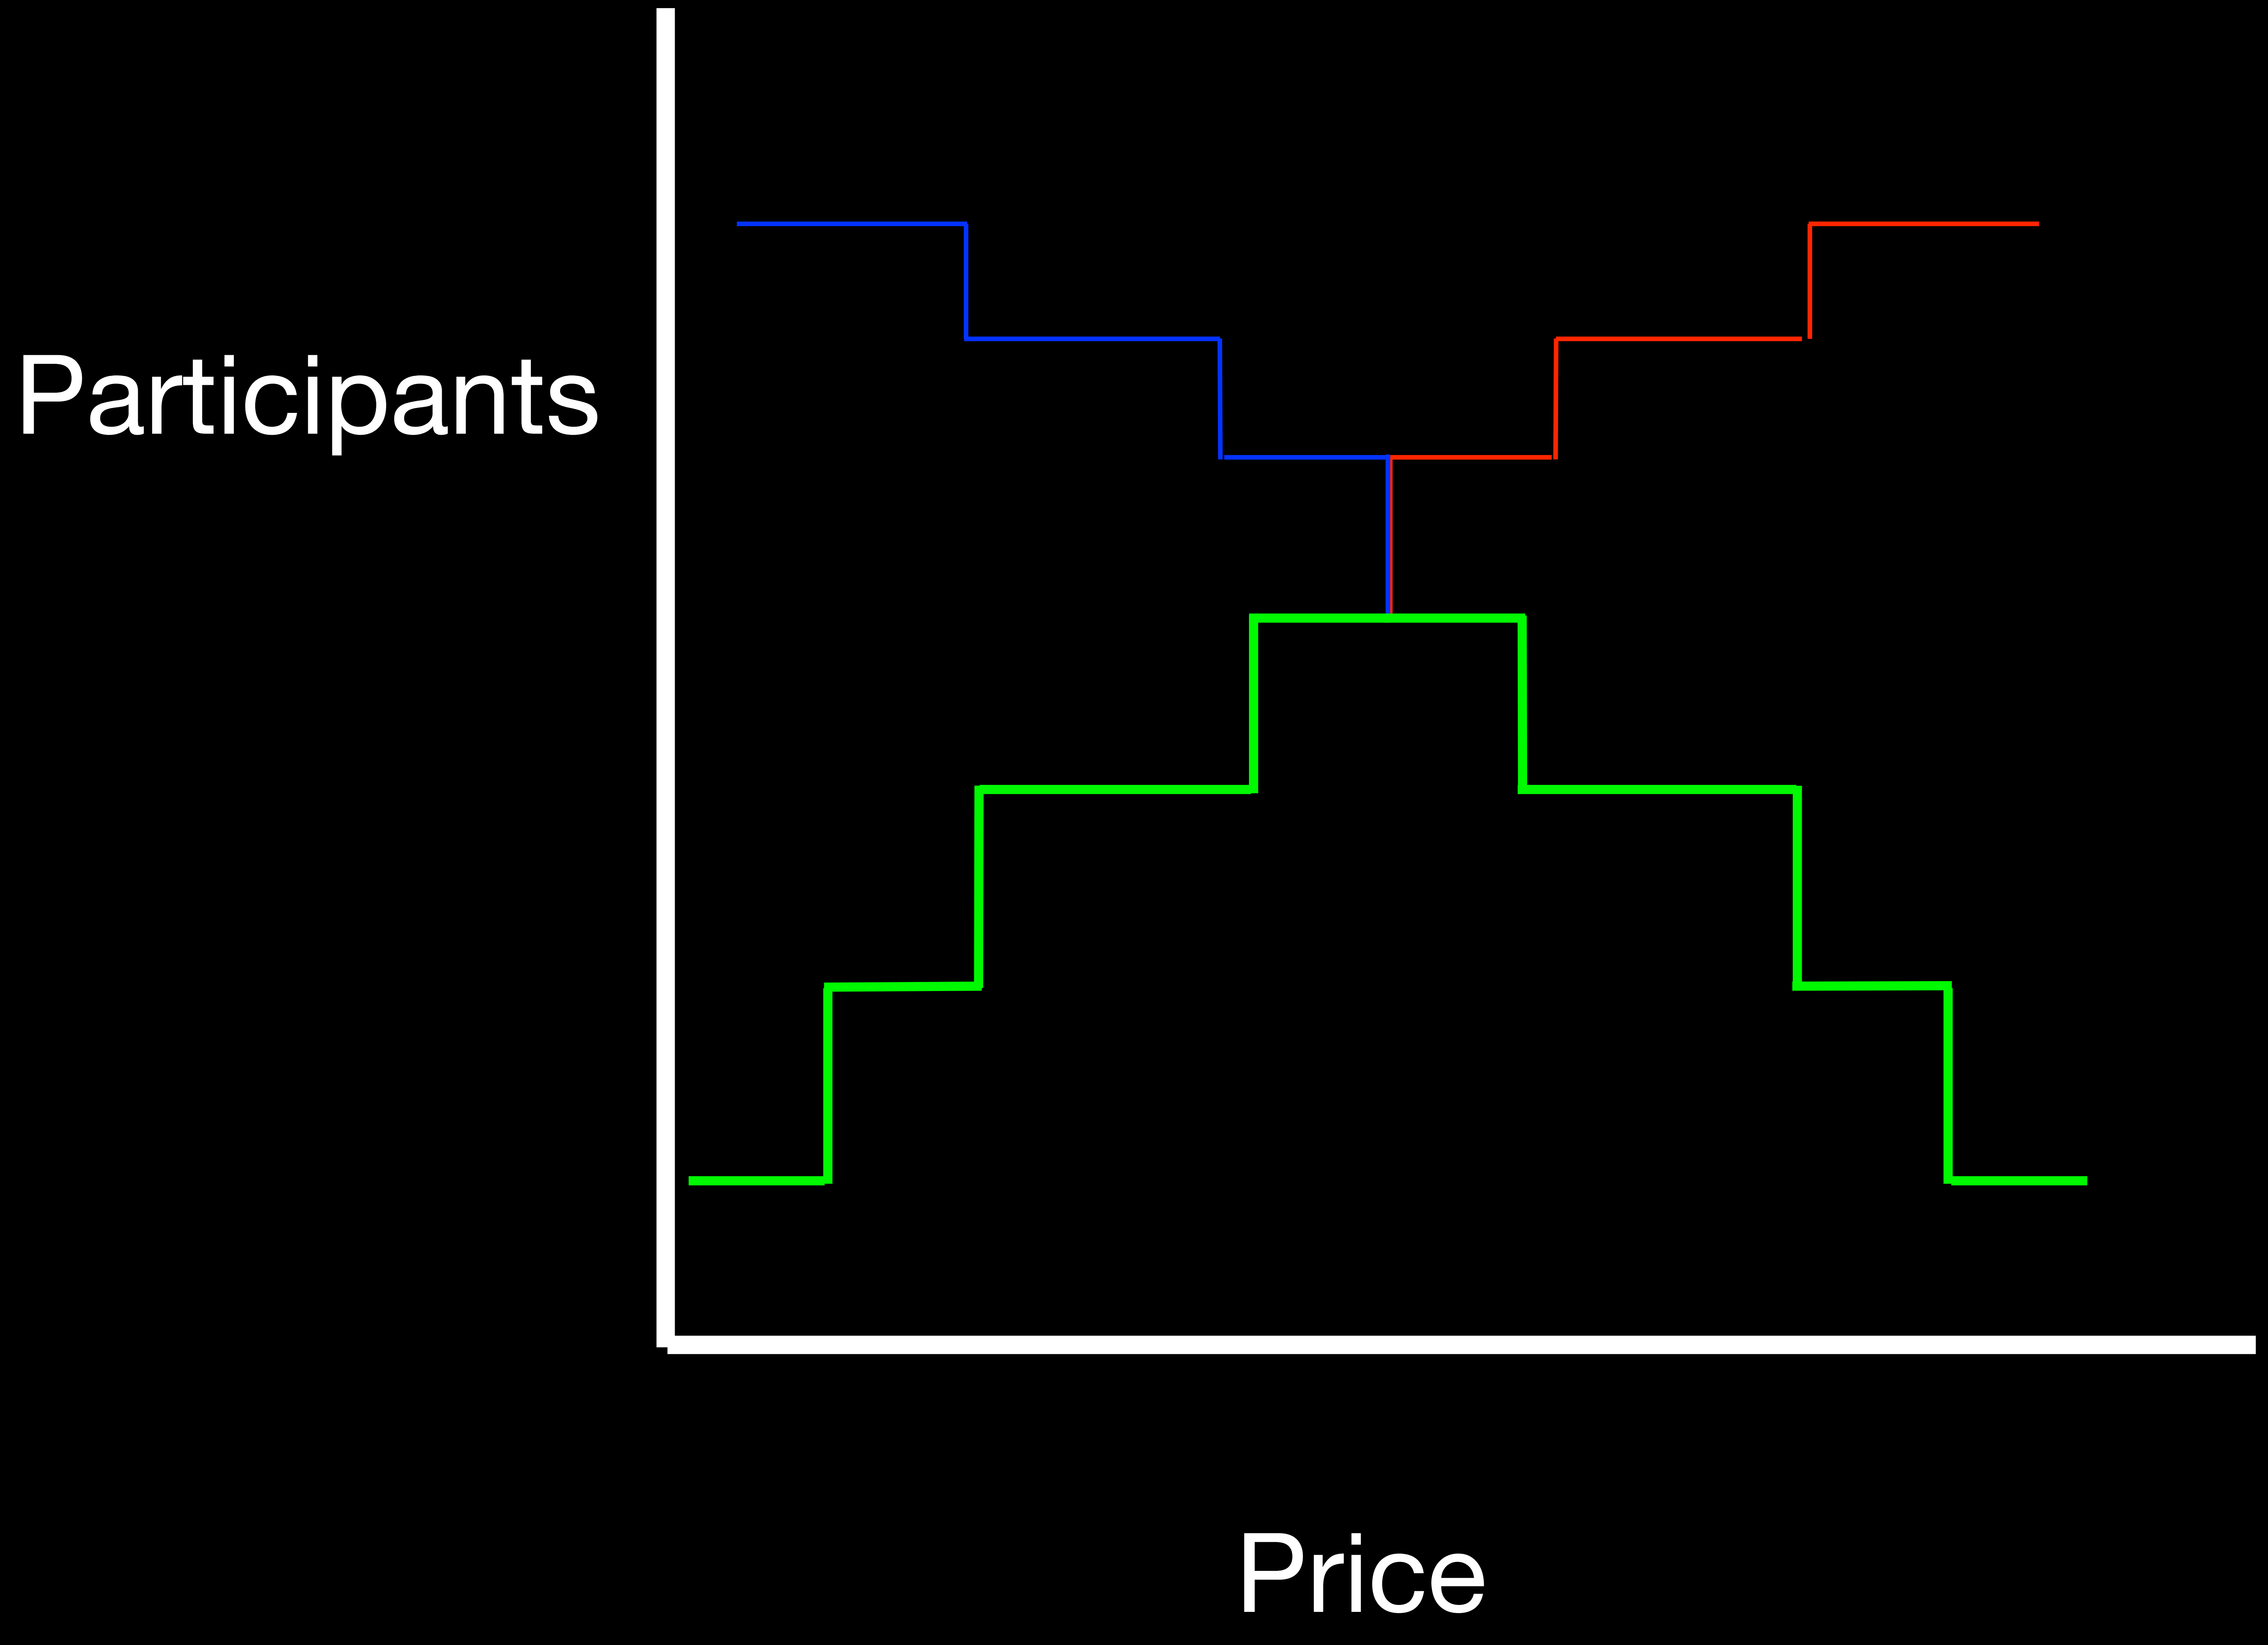

Call Auction

Valuation	Participant Type
1	Buyer
2	Buyer
2	Buyer
5	Buyer
2	Seller
3	Seller
3	Seller
7	Seller

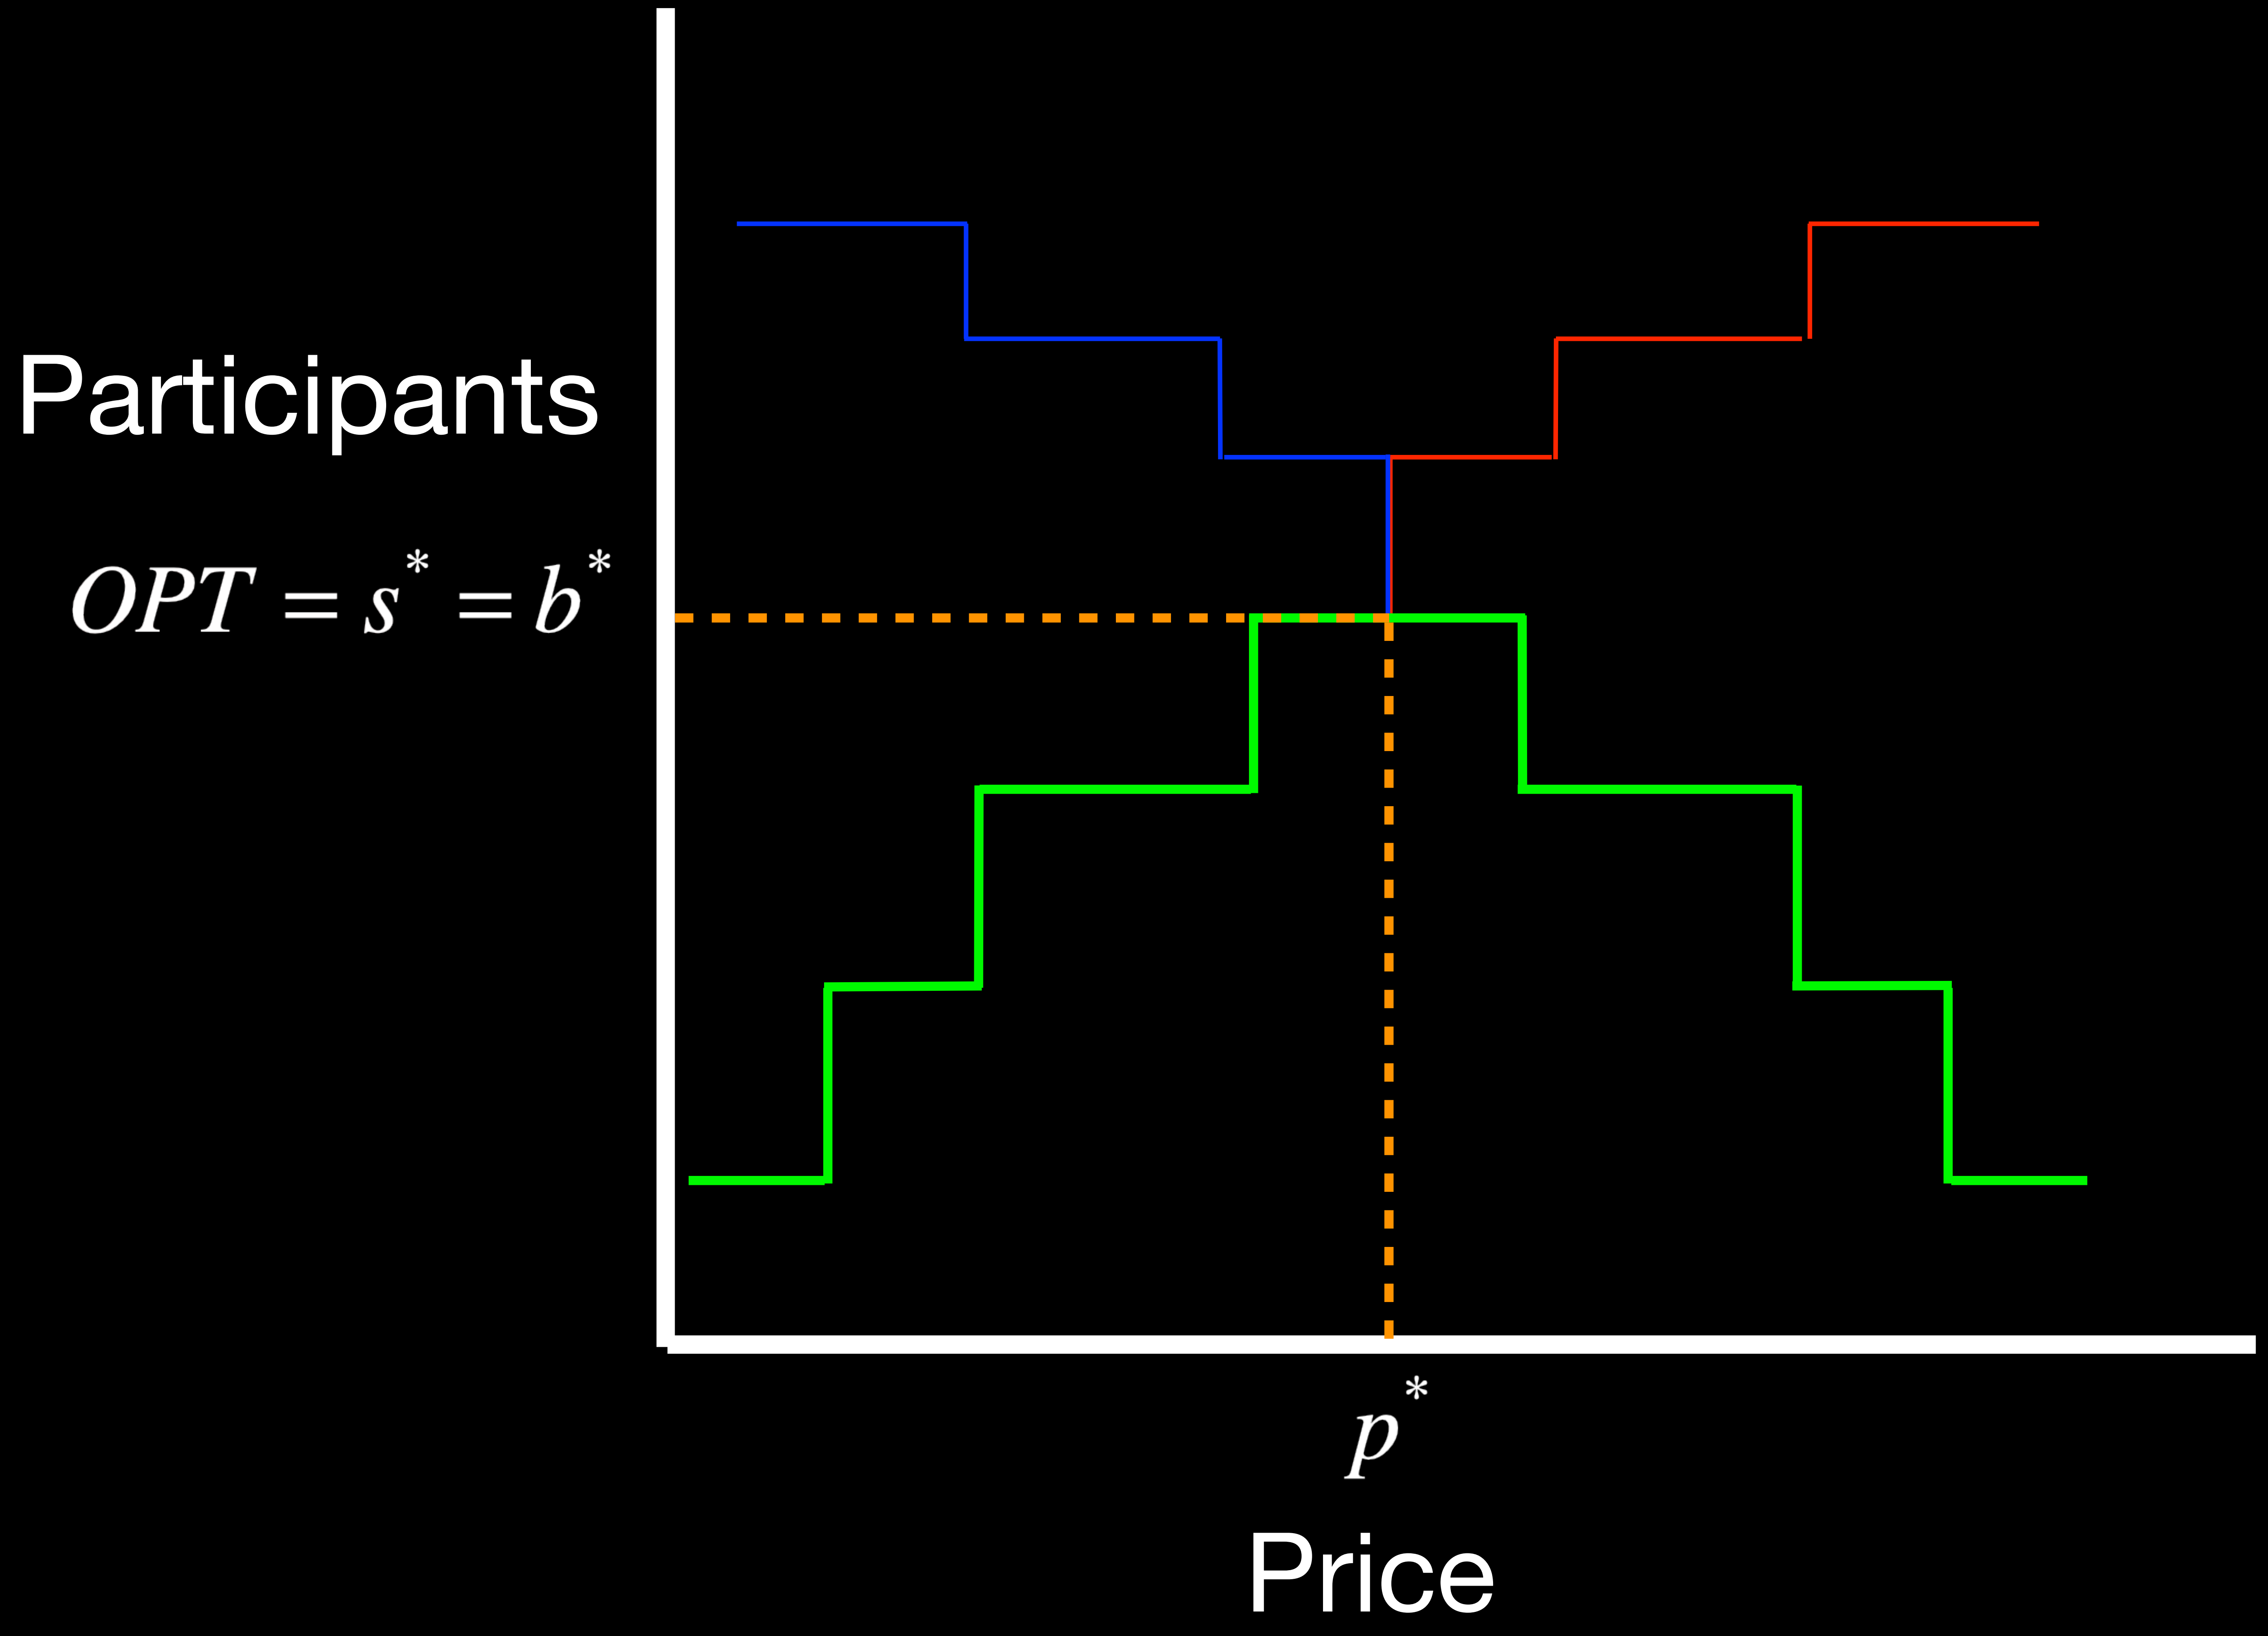

What do we mean by privacy?

One definition is standard (ϵ, δ)-Differential Privacy

- Let $\mathcal{M}$ be a mechanism/algorithm
- Let *D, D'* be any pair of *neighboring* inputs/datasets - differ in single element
- Let *S* be any subset of possible outputs of mechanism
- We say $\mathcal{M}$ is (ϵ, δ)-Differentially Private if for any *S, D, D':*

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

Examples and Properties of DP

- Examples: computing the mean; randomized response
- Undetectability by third-party observer
- Immunity to post-processing
- Composition and graceful degradation
- Deployments: Apple, Google, 2020 U.S. Census

Joint Differential Privacy

- Similar semantics as differential privacy, but applies when algorithm has output for each participant as well

- Informally, asks that nothing can be learned about your input even if other participants collude

Joint Differential Privacy

- *D, D'* neighboring datasets that differ at single element *i*
- $\mathcal{M}$ a mechanism that outputs a vector whose dimension is the size of the databases, $\mathcal{M}_{-i}$ is output on all other than *i*

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

Mechanism*

Get Valuations

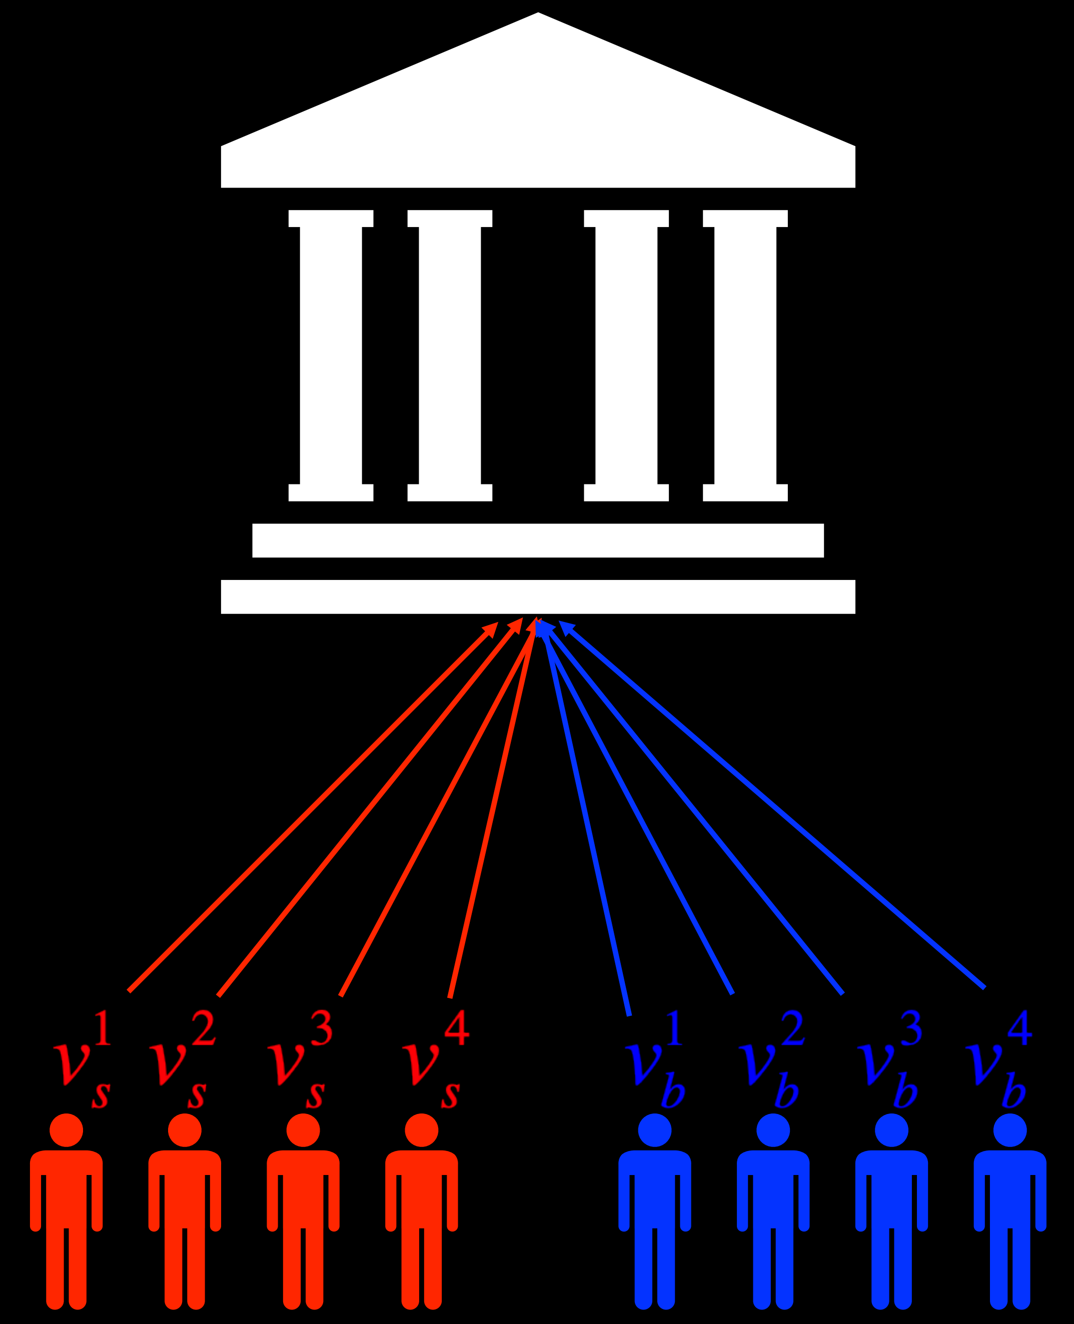

Select Price

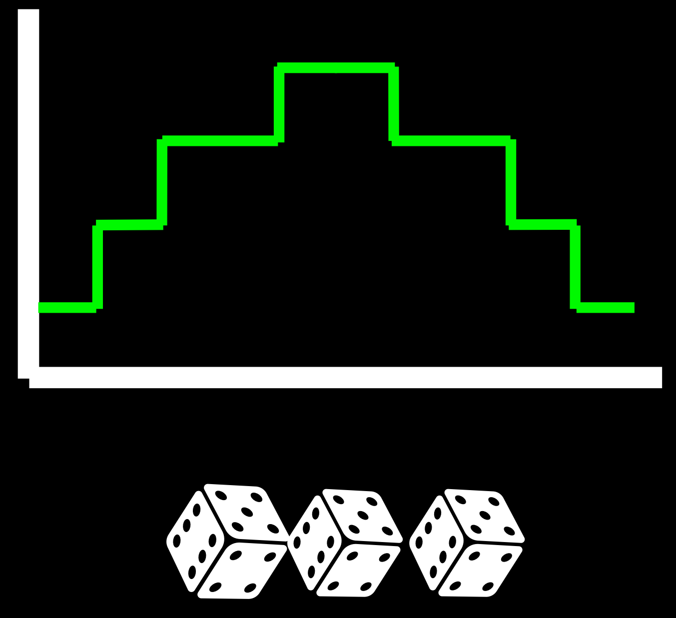

$$\Pr[p] \propto \exp\left(\frac{\epsilon}{2}\text{shares}(p)\right)$$

$$\tilde{p}$$

Estimate Sellers and Buyers

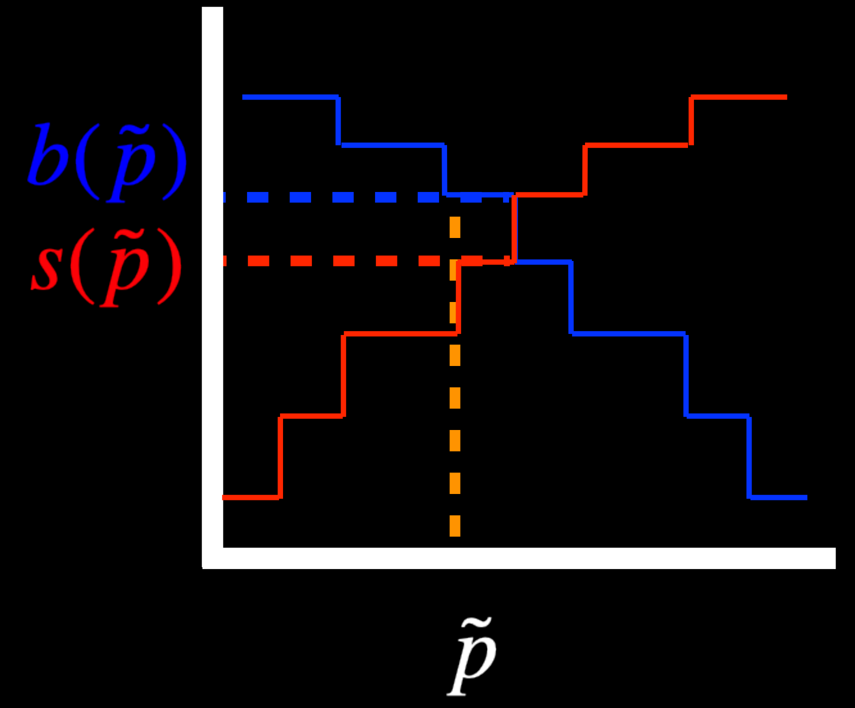

$$\hat{s} = s(\tilde{p}) + \text{Lap}\left(\frac{1}{\epsilon}\right)$$

$$\hat{b} = b(\tilde{p}) + \text{Lap}\left(\frac{1}{\epsilon}\right)$$

Select participants

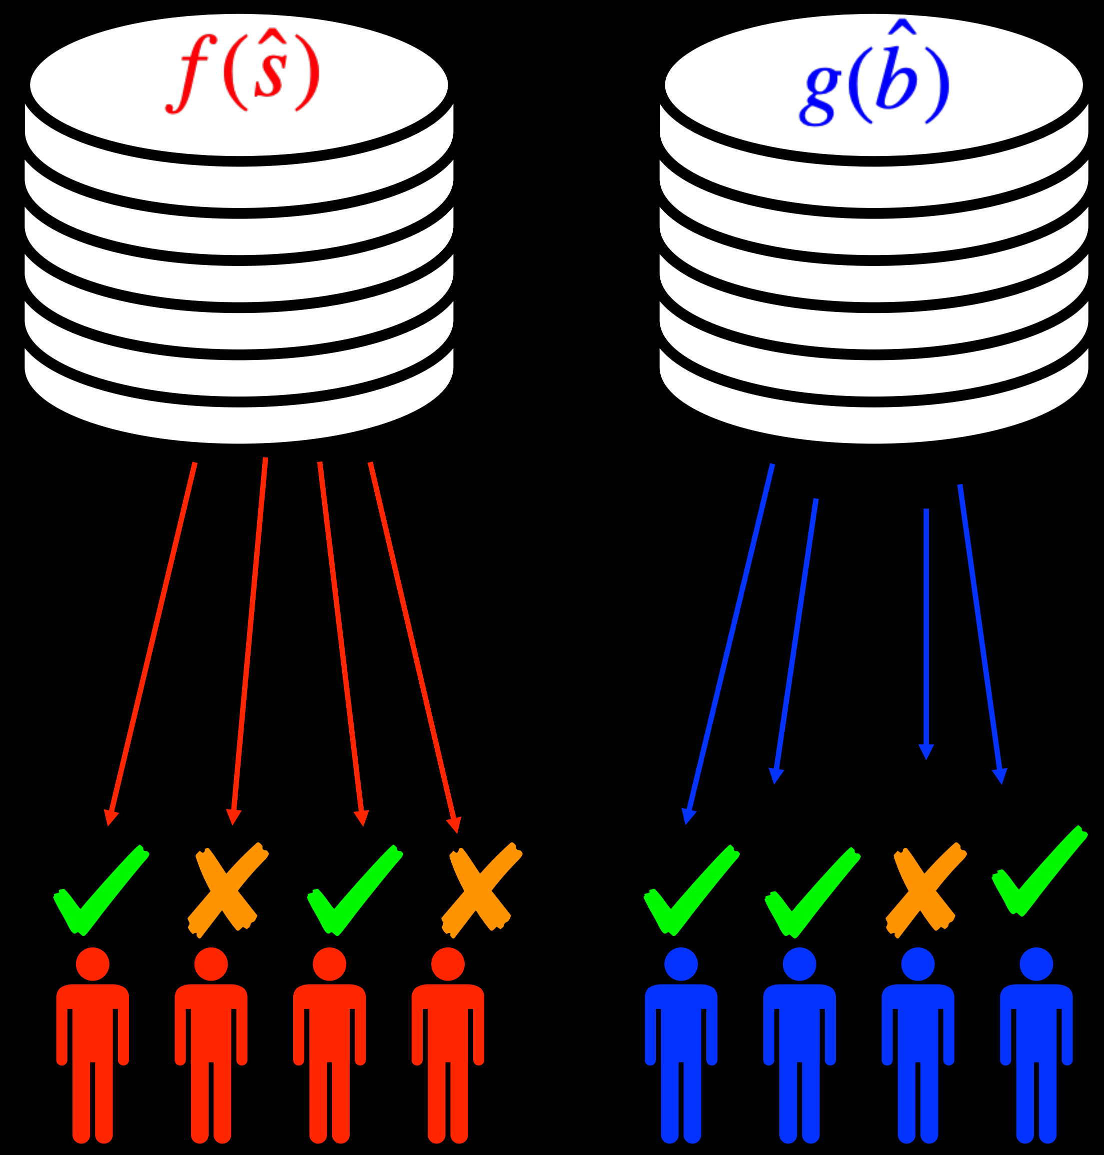

*Actually have two, with different guarantees, and privately select best one

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

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

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

Guarantees

- The mechanism satisfies $(\epsilon, 0)$*-joint* differential privacy

- With high probability $(1-\alpha)$, clear shares at least

$$\mathrm{OPT} - \mathcal{O}\left(\frac{\ln(1/\alpha)}{\epsilon} + \sqrt{\mathrm{OPT}}\right)$$

- With high probability, inventory taken on is at most

$$\mathcal{O}\left(\frac{\ln\ 1/\alpha}{\epsilon} + \sqrt{\mathrm{OPT}}\right)$$

A Lower Bound on Inventory

Theorem 4. *[Lower bound on the loss of private algorithms] Pick any* ε, δ *such that* $0 \leq \varepsilon \leq 1$ *and* $\delta = \mathcal{O}(\varepsilon)$. *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* $\Omega(1/\varepsilon)$.

Connections to Market Impact Literature

- E.g. widely studied *square root law*: change in price ~ $\sqrt{k/\mathcal{V}}$
  - [Gatheral 2010; Bouchaud et al.]
- *V* = interval total volume trade; *k* = participation rate
- Our results: change in price ~ $(e^{k\varepsilon} - 1)$
- Matching the two theories yields $\varepsilon \approx 1/\sqrt{kn}$ and $\mathrm{OPT}(1 - \sqrt{kn}/\mathrm{OPT})$ shares cleared
- So in regime where OPT ~ *n* and *k* = *o(n)*, two major advantages:
  - *No* assumptions on orders/participants
  - Bounding information leakage of *any* kind, not just price impact

Mechanism Incentive Properties

- Mechanism is *individually rational* - participants only trade if the price chosen is one they are willing to trade at

- Mechanism is *approximately incentive compatible* - agents can't gain by more than a small amount by misreporting their valuation

    - This is *not* true in standard call auction because optimal price is not stable - even a single person misreporting could change price a lot

Simulation: One-shot Game

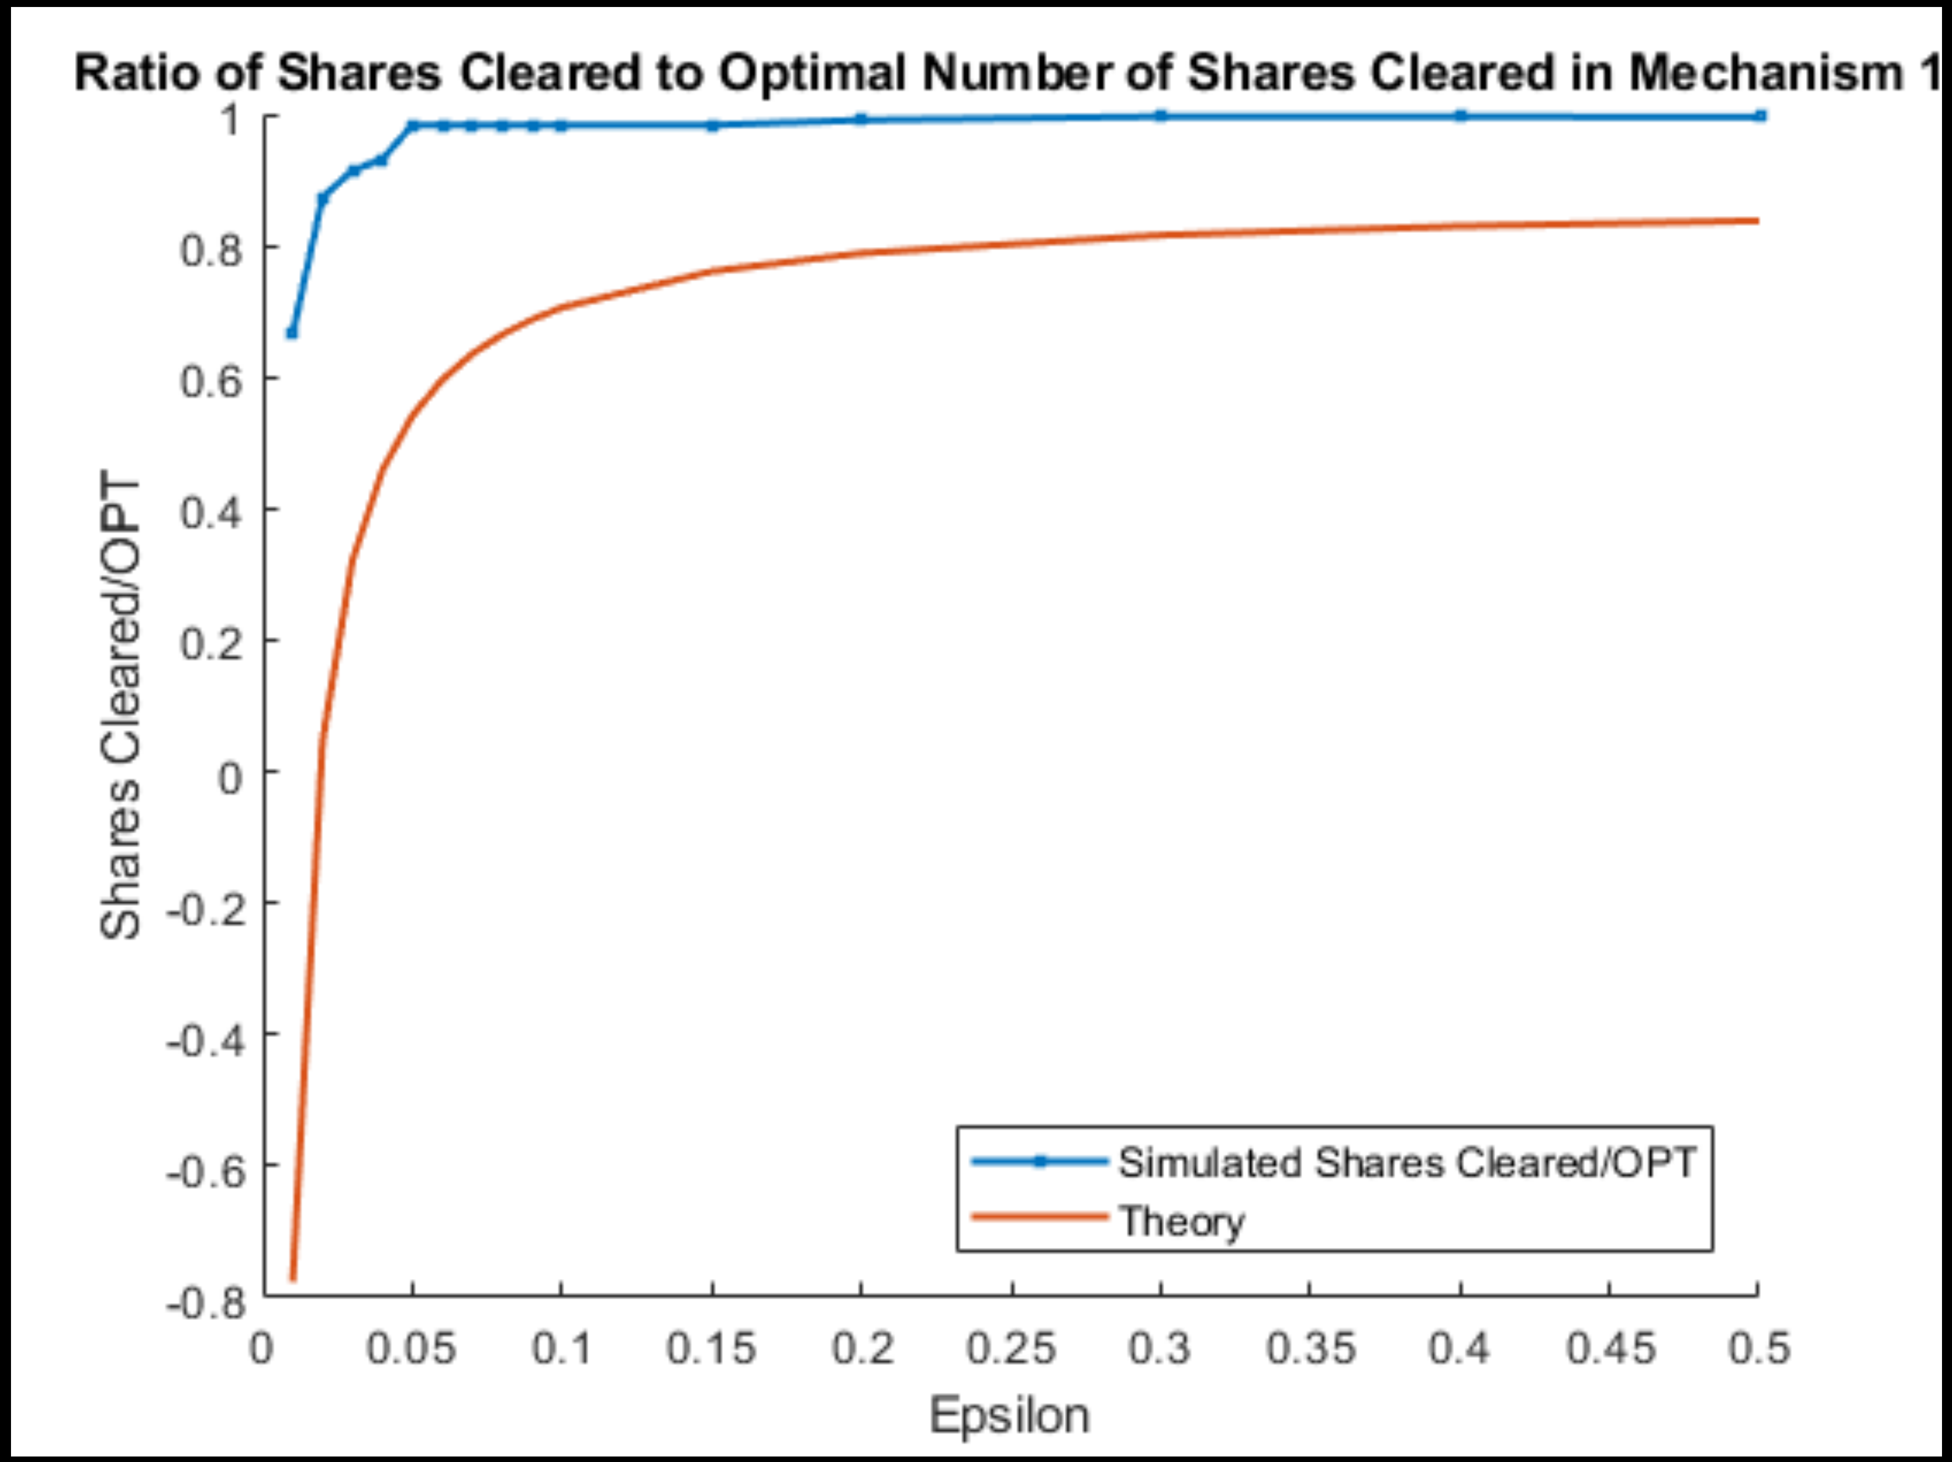

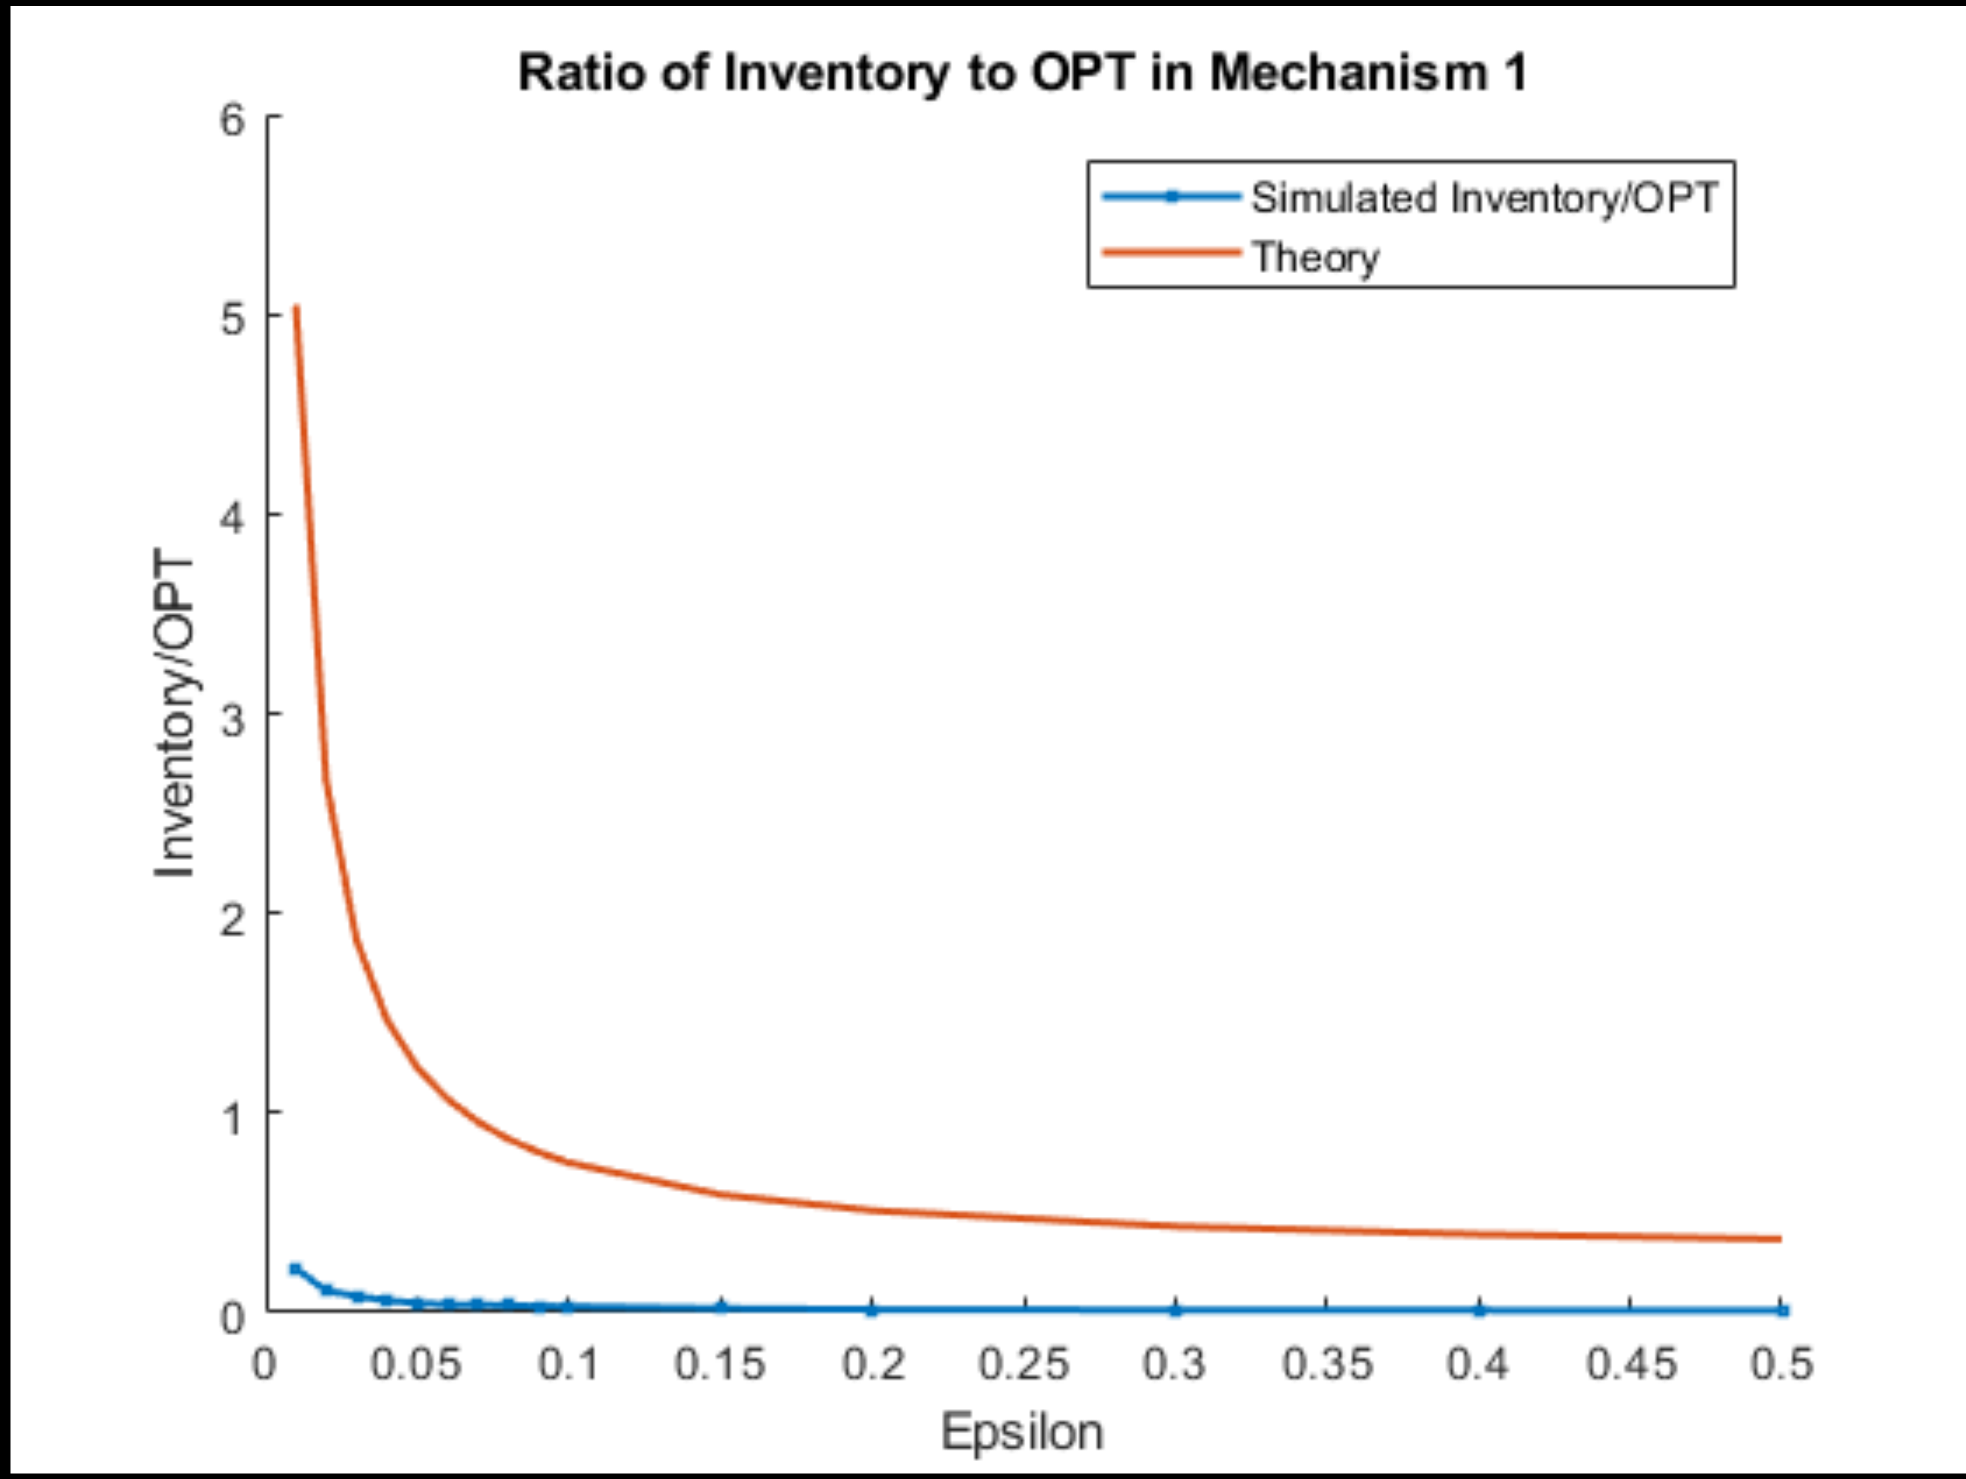

Learning in a repeated setting

- We study a setting where agents valuations are drawn once and the mechanism repeated for many rounds. Assume agents are not fully strategic but learn to bid over time via no-regret dynamics

- Captures pre-auction ‘hypothetical’ auctions NYSE/NASDAQ; also captures agents that may not initially bid truthfully, or may not trust that the mechanism is really IC

- Rather than try to mechanistically model traders, want robust algorithmic setting to capture reasonable behavior

- How does our mechanism perform in such a setting?

Learning in a repeated setting

Theorem 7. *Suppose buyers and sellers update their bidding strategies according to Algorithm 5 (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*

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

where this probability is taken with respect to the randomness of both Algorithms 1 and 5.

Simulation: Repeated Game

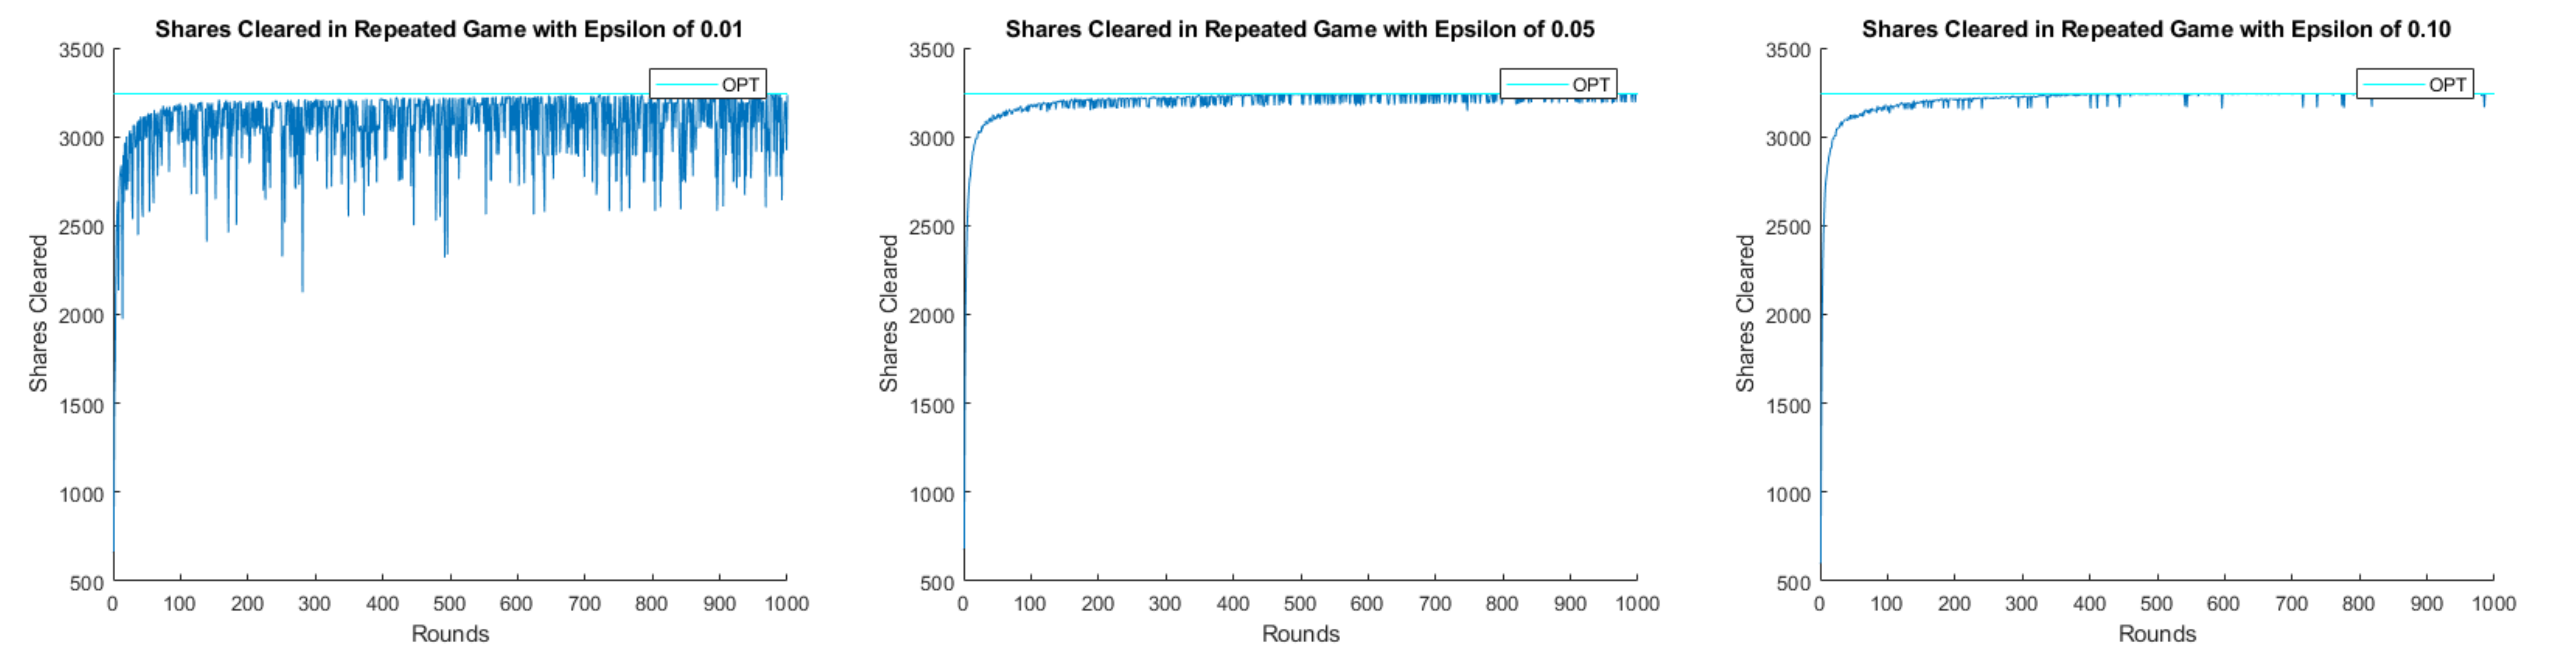

Recap

In this work, we:

- Design a joint-differentially private call auction
  - Guaranteed to be (ϵ, δ)-joint DP
  - Achieve good performance in terms of shares cleared
  - Take on only small amount inventory
- Simulate empirical performance and show that we do well relative to theory
- Prove that mechanism will converge to the optimal shares cleared when agents use some no-regret learning algorithms
- Demonstrate that this happens reasonably quick in simulation

Thanks!

https://arxiv.org/abs/2002.05699