Optimal, Truthful, and Private Securities Lending

Emily Diana, Michael Kearns, Seth Neel, Aaron Roth

NeurIPS 2019 Workshop on Robust AI in Financial Services: Data, Fairness, Explainability, Trustworthiness, and Privacy

Penn

Motivation

Motivated by challenges associated with securities lending, the mechanism underlying short selling of stocks in financial markets

- Consider allocation of a scarce commodity in settings in which privacy concerns or demand uncertainty may be in conflict with truthful reporting
- Want to construct a privacy protecting allocation mechanism that motivates truthful reporting without sacrificing too much utility

Model

- Lender distributes up to V shares to n clients over time horizon T
- At each time t, client i draws from a joint distribution over usages and requests, $Q_{it}(u_{it}, r_{it})$, but only request is visible to lender
- Lender chooses share allocation $S_t = \{s_{it}\}$ s.t. $\sum_i s_{it} \leq V$
- Client's payoff is number of shares actually used, and lender's utility for allocation rule A is:

$$v(A) = \sum_i \mathbb{E}_{Q_{it},A}[\min(A(r_1, \ldots, r_n; Q_1, \ldots, Q_n)_i, u_{it})]$$

Table 1: Sample Truthful Distribution

u_{it} \ r_{it}	0	1	2
0	$\frac{1}{3}$	0	0
1	0	$\frac{1}{3}$	0
2	0	0	$\frac{1}{3}$

Table 2: Sample Untruthful Distribution

u_{it} \ r_{it}	0	1	2
0	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$
1	0	$\frac{1}{6}$	$\frac{1}{6}$
2	0	0	$\frac{1}{3}$

Optimal Allocation Rule

Given knowledge of Q_i, the lender can compute the posterior distribution $Q_i(u_i|r_i)$ on the true demand u_i given r_i, via Bayes' rule:

$$Q_i(u_i|r_i) = \frac{Q_i(r_i|u_i)U_i(u_i)}{\sum_{u'} Q(r_i|u')U_i(u')}$$

Algorithm 1 Greedy Allocation Rule

Input: $n, \{Q_i(u_i|r_i)\}_{i\in[n]}, V$
Output: feasible allocation $S = \{s_i\}$.
procedure GREEDY($n, \{Q_i(u_i|r_i)\}_{i\in[n]}, V$)
  Initialize $s_i = 0$, $\forall i$. ▷ number of shares allocated to client i
  for $t = 1 \ldots V$ **do**
    Let $i^* = \text{argmax}_i T_i(s_i + 1|r_i)$
    update $s_i \leftarrow s_i + 1$
  end for
end procedure

Theorem: The allocation returned by *Greedy* maximizes the expected payoff for the lender: For S the allocation output by greedy:

$$S \in \arg\max_{S:\sum_i s_i = V} v(S) = \sum_i \mathbb{E}_{Q_i(u|r_i)}[\min(s_i, u_i)]$$

Dominant-Strategy Truthfulness

Given that the lender is solving the allocation problem optimally for the reported Q distributions, truth telling is a dominant strategy

Theorem: Fix a set of choices Q_{-i} and reports r_{-i} for all clients other than i, and a realization of client i's usage $u_i \sim U_i$. Let Q_i^T denote the truthful strategy $Q_i^T(r_i|u_i) = \mathbf{1}_{r_i}$, and let $Q_i(r_i|u_i)$ denote any other strategy. Let A denote the lender's optimal allocation. Then:

$$v_A^i(Q_i) \leq v_A^i(Q_i^T)$$

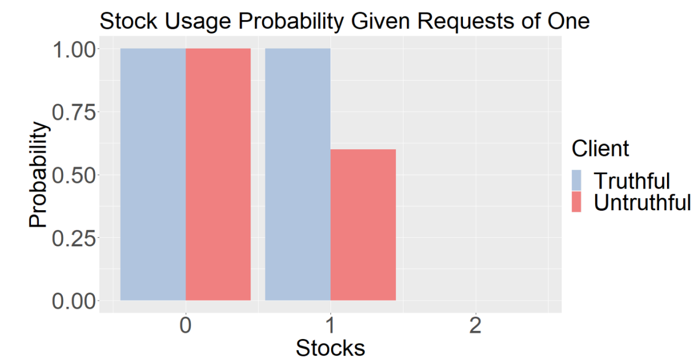

Private Auction Formulation

- Optimal allocation policy can be implemented as a virtual ascending auction among clients
- Bidders (clients) have decreasing marginal valuation functions for up to U units of each good (stock)
- We modify auction to guarantee joint differential privacy by
  1. Reporting number of bids placed so far with a differentially private estimator
  2. Allowing the algorithm to stop early
  3. Running the auction with $V - E$ shares, where E corresponds to error of differentially private bid counter
- Then, truthful reporting is still an approximately dominant strategy
- Finally, if clients are allowed to adapt strategies with time, joint differential privacy enforces truthfulness as an approximately dominant strategy and guarantees near optimality

Theorem: Let A be a private auction with appropriate values of U, V, ϵ and ρ such that A is $(\epsilon', \beta/T)$-JDP with $\epsilon' = \tilde{O}(\epsilon/\sqrt{T})$ and outputs S such that $E[V(S)] \geq (1-\rho)OPT_V - \rho$. Take β, ρ such that $\sqrt{\beta + (1-\beta)\rho} \leq \beta^2/T$. Then for a $(1-\beta)$ fraction of the n clients i, let L_{i*}^t denote the truthful strategies, and let L_i^t be any other set of strategies. Then a private greedy allocation rule for the private auction satisfies:

$$v_i(L_i^1, \ldots, L_i^n) \leq e^{2\epsilon} v_i(L_{i*}^1, \ldots, L_{i*}^n) + 2\beta UT + e^{\epsilon}\frac{\beta^2}{1-\beta^2/T}$$

$$v_A(L_{i*}^t) \geq (1-\rho)OPT_V - \rho T,$$

where OPT_V denotes the lender's optimal utility.

Selected References

Kuzman Ganchev, Michael Kearns, Yuriy Nevmyvaka, and Jennifer Wortman Vaughan (2010)
Censored Exploration and the Dark Pool Problem
Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence 185–194.

Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu (2014)
Private matchings and allocations
Proceedings of the Forty-sixth Annual ACM Symposium on Theory of Computing 21–30.

Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman (2014)
Mechanism design in large games: Incentives and privacy