Differentially Private Diffusion Auction: The Single-unit Case

Extended Abstract

Fengjuan Jia
University of Electronic Science and Technology of China
Chengdu, China
202221080937@std.uestc.edu.cn

Mengxiao Zhang
University of Electronic Science and Technology of China
Chengdu, China
mengxiao.zhang@uestc.edu.cn

Jiamou Liu
The University of Auckland
Auckland, New Zealand
jiamou.liu@auckland.ac.nz

Bakh Khoussainov
University of Electronic Science and Technology of China
Chengdu, China
bmk@uestc.edu.cn

ABSTRACT

Diffusion auction refers to an emerging paradigm where an auctioneer utilises a social network to attract potential buyers. We consider the risks of disclosing sensitive preferences of buyers from the published auction outcome and initiate the study of differential privacy in diffusion auction. We study the single-unit case and design two differentially private diffusion mechanisms (DPDMs): recursive DPDM and layered DPDM. We prove their incentive and privacy properties, and then empirically compare their performance on real and synthetic datasets.

KEYWORDS

Diffusion auction; differential privacy; mechanism design

ACM Reference Format:
Fengjuan Jia, Mengxiao Zhang, Jiamou Liu, and Bakh Khoussainov. 2023. Differentially Private Diffusion Auction: The Single-unit Case: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

Diffusion auction is an emerging business paradigm of online social network commerce. In this setting, a seller is able to harness the power of social network to diffuse auction information, inviting friends, friends-of-friends, etc., to join the auction, thereby attracting a large number of potential buyers. This differs from a standard auction (without social network) where the participants are fixed beforehand. A challenge in diffusion auctions lies in resolving the conflict between the seller who wants to attract more participants for better revenue and the buyers who are reluctant to invite their friends to avoid competition. Thus there is a need to extend *incentive compatibility* (IC) for hidden valuations in classical auctions, to *diffusion IC* for hidden valuation as well as social ties. Numerous studies, e.g., [5–7, 13–15], have proposed mechanisms for diffusion auction that achieve diffusion IC.

In an auction, buyers submit their (private) valuations in bids to the auctioneer. The bids often imply buyers' preferences and confidential business strategies, and competitors may exploit them to gain an advantage. Hence, there is a need to protect the privacy of bid information. The privacy issues in classical auctions have recently been studied in [1, 3, 4, 8–12, 16, 17]. These studies employ the well-established notion of *differential privacy* (DP) [2] to mitigate privacy risks. To achieve DP on bids, [8] proposes *exponential mechanism*. The mechanism randomises auction results so that a change in a buyer's bid does not significantly affect the auction outcome. In this way, the mechanism prevents the bid from being inferred from the auction outcome.

However, no study has focused on the privacy issues for diffusion auctions. Here we close this gap by investigating the following question: *How do we design a differentially private diffusion mechanism (DPDM) that guarantees desirable properties and preserves valuation privacy?*

2 PROBLEM FORMULATION

Consider the following setup: There is a seller, s and buyers $N = \{1, 2, \ldots, n\}$. Seller s has a single indivisible item to sell. Each buyer $i \in N$ is willing to buy the item and attaches a *valuation* v_i to the item. The seller and the buyers form a social network $G = (V, E)$, where $V = N \cup \{s\}$ and $E \subseteq V^2$. Each node $i \in V$ has a neighbour set $r_i := \{j \in V \mid (i, j) \in E\}$. The pair $\theta_i = (v_i, r_i)$ is called the *true profile of the buyer* i. Each buyer $i \in N$, once invited to the auction, is asked to report her profile $\theta'_i = (v'_i, r'_i)$, which might not be the true one. This forms the tuple $\theta' := (\theta'_1, \ldots, \theta'_n)$ called a *global profile of all buyers*. By Θ we denote the set of all such profiles. Given $\theta' \in \Theta$, we construct $G_{\theta'} = (V_{\theta'}, E_{\theta'})$ the directed graph: add a directed edge (i, j) if j is reported by i as a neighbour. We call such graph *profile digraph*. The *utility* of buyer i is $u_i(\theta') = v_i \pi_i(\theta') - p_i(\theta')$ when reported global profile is θ'. The *social welfare* of a mechanism M on θ', written $sw_M(\theta')$, is the sum of all utilities, i.e., $sw_M(\theta') = \sum_{i \in V} u_i(\theta')$.

Designing a DPDM has mainly three challenges: **(1) Valuation asymmetry.** Buyers' true valuations are hidden from the seller. Thus buyers have an advantage over the seller as they can misreport their valuations. **(2) Neighbourhood asymmetry.** Buyers' neighbours are hidden, so buyer i may misreport the neighbour set $r'_i \subseteq r_i$ as disseminating the auction information may hinder their own chance of winning. **(3) Valuation privacy.** Once the auction result is announced, an attacker may infer the bid information from the published auction result. This disadvantages the buyer(s) whose private valuation is diclosed.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

To preserve privacy, we use randomisation. Let Ω be a probability space. A *randomised mechanism* M consists of two randomised functions $(\pi(\cdot), p(\cdot))$, where $\pi\colon \Theta\times\Omega \to \{0,1\}^n$, $p\colon \Theta\times\Omega \to \mathbb{R}^+$.

Definition 2.1. *Let M be a randomised mechanism. Call the mechanism M ϵ-*differentially private *(ϵ-DP) if for any two global profiles $\theta', \theta'' \in \Theta$ that differ on a single buyer's valuation, and for any possible outcome $o \in O$, $\Pr[M(\theta') = o] \le \exp(\epsilon)\Pr[M(\theta'') = o]$.*

In randomised mechanisms, we use $\mathbf{E}_M[u_i(\cdot)]$ to denote i's expected utility in M. The social welfare of M is also in expectation, i.e., $\mathbf{E}_M[sw_M(\theta)] = \sum_{i\in V}\mathbf{E}_M[u_i(\theta)]$.

Definition 2.2. *Let M be a randomised mechanism,*

- *The mechanism M is* IC *if for all $i \in N$, all $\theta_i, \theta_i' \in \Theta$ and for all $\theta'_{-i}, \theta''_{-i} \in \Theta^{n-1}$, we have the following, $\mathbf{E}_M[u_i((\theta_i, \theta'_{-i}))] \ge \mathbf{E}_M[u_i((\theta_i', \theta''_{-i}))]$.*
- *The mechanism M is* IR *if for all $i \in N$ and all $\theta'_{-i} \in \Theta^{n-1}$, we have $\mathbf{E}_M[u_i((\theta_i, \theta'_{-i}))] \ge 0$.*

We aim to design a randomised mechanism that is IC, IR, ϵ-DP (for reasonable ϵ) while maximising expected social welfare.

3 RECURSIVE DPDM

Preserving valuation privacy in diffusion auctions is not a trivial task. Existing diffusion auctions, including IDM [7], CMD [6], and FDM [14], are deterministic, and thus fail to preserve privacy. Existing DP mechanisms, including exponential mechanism, fail to incentivise truthful report of neighbours as inviting more participants means a lower probability of winning the auction.

To incentivise buyers to diffuse auction information, we need to ensure each buyer's utility of reporting her neighbours should be no less than that of non-reporting. We propose *recursive DPDM* REC to achieve this. The basic idea is "market division", i.e., treat the social network as a market, partition the market into multiple sub-markets and assign each sub-market a probability with which buyers in this sub-market win. Then each buyer would report as many neighbours as possible in order to maximise the probability of the sub-market she belongs to. The buyers in a sub-market share the probability of the sub-market in such a way that the winning probability of any buyer is independent from her children. Therefore, the buyers have no competition with their children and have no incentive to hide them.

Specifically, fix a score function $\sigma(\cdot)$ non-decreasing in v_i'. Given $\theta' \in \Theta$, a privacy parameter ϵ and the function $\sigma(\cdot)$ as input, REC works as follows: **(1)** From the profile digraph $G_{\theta'}$, REC constructs a *diffusion critical tree* $T_{\theta'}$ [15]. **(2)** REC determines winning probabilities. The process is recursive and starts with $T_{\theta'}$. Given a (sub-)tree rooted by $i \in V$, REC assigns a probability to each subtree rooted by $j \in r_i$, and a winning probability to each $j \in r_i$. This operation is repeated for j's children, children of j's children and so on until there is no more children. **(3)** REC randomly selects a buyer w as a winner according to the constructed distribution in Step (2). REC sets w's allocation $\pi_w = 1$, and payment as $p_w = v'_w - \int_0^{v'_w} \Pr_w((x, r'_w))dx / \Pr_w(\theta'_w)$.

Let $d_{\max}$ be the maximum depth of the diffusion critical tree and $\Delta\sigma$ be the largest possible difference in score function σ when applied to two global profiles that differ only on a single valuation, for all possible outcome $o \in O$. We have the following theorem.

Theorem 3.1. *Recursive DPDM* REC *is IC, IR and $\epsilon d_{\max}\Delta\sigma$-DP.*

4 LAYERED DPDM

Following the idea of market division, we propose layered DPDM LAY in this section. Different from REC, LAY divides the market by buyers' distances to seller, i.e., LAY allocates a probability to each layer of the tree, which is shared by the buyers on this layer. For any buyer, once she is invited, her layer is fixed. Also, the buyer(s) whom she invites is on the next layer, and thus has no competition with her.

Specifically, LAY executes the same operations as in REC, where the only difference is in Step (2). In Step (2), given a critical diffusion tree $T_{\theta'}$ and an infinite decreasing sequence $\gamma = (\gamma_1, \gamma_2, \ldots)$, where $\sum\gamma_i = 1$, LAY assigns a probability $\Pr_{L_\ell} = \gamma_\ell$ to each layer $L_\ell, 1 \le \ell \le d_{\max}$, of the tree and then assigns a winning probability to buyers on layer L_ℓ. The following theorem shows the incentive and privacy properties of LAY.

Theorem 4.1. *Layered DPDM* LAY *is IC, IR and $\epsilon\Delta\sigma$-DP.*

Next we analyse the expected social welfare of LAY. We consider a hypothetical scenario where the exponential mechanism is applied to the whole social network where the seller knows all buyers and the auction information is diffused to all buyers without any incentive. We call such a mechanism as *exponential mechanism with diffusion (EMD).* EMD has the optimal expected social welfare among all DPDMs and thus is used as the benchmark.

Theorem 4.2. *Given a global profile θ, layered DPDM* LAY *has $\mathbf{E}_{\text{LAY}}[sw_{\text{LAY}}(\theta)] \ge \gamma_{d_{\max}}\mathbf{E}_{\text{EMD}}[sw_{\text{EMD}}(\theta)]$.*

Corollary 4.3. *For $\gamma = (\frac{a-1}{a}, \frac{a-1}{a^2}, \ldots)$, where $a > 1$, the expected social welfare of layered DPDM* LAY *is $\ge \frac{a-1}{a^{d_{\max}}}\mathbf{E}_{\text{EMD}}[sw_{\text{EMD}}(\theta)]$.*

5 EXPERIMENT

We evaluate the performances of REC and LAY, in terms of social welfare under different privacy levels and valuations on three real world social network datasets. We also analyse the effect of sequence $\gamma = (\frac{a-1}{a}, \frac{a-1}{a^2}, \ldots)$ on the performance of LAY. We use three benchmarks: IDM [7], EMD and EMWD. EMWD: Apply the exponential mechanism only to the seller's neighbours. The expected social welfare of EMWD can be seen as the lower bound.

The experiment results show that: Overall, when comparing to IDM, the difference in social welfare of the DPDMs decreases with ϵ increases. Then, among DPDMs, EMD performs best in most cases, followed by REC and LAY. Particularly, REC performs very well. The deviation of REC from EMD is at most 2.62%. Figure 1 shows the average social welfare of REC, LAY and three benchmarks under normally distributed valuation and γ with $a = 2$.

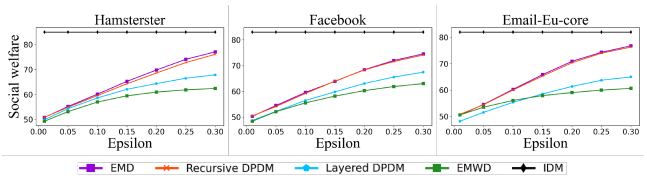

Figure 1: Average social welfare of LAY, REC, **EMD, EMWD and IDM**

REFERENCES

[1] 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*. 541–583.

[2] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.

[3] Zhiyi Huang and Sampath Kannan. 2012. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. IEEE, 140–149.

[4] Lin Jian, Yang Dejun, Li Ming, Xu Jia, and Xue Guoliang. 2018. Frameworks for Privacy-Preserving Mobile Crowdsensing Incentive Mechanisms. In *IEEE Trans. Mob. Comput.* 1851–1864.

[5] Takehiro Kawasaki, Nathanaël Barrot, Seiji Takanashi, Taiki Todo, and Makoto Yokoo. 2020. Strategy-proof and non-wasteful multi-unit auction via social network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 2062–2069.

[6] Bin Li, Dong Hao, Dengji Zhao, and Makoto Yokoo. 2019. Diffusion and auction on graphs. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. 435–441.

[7] Bin Li, Dong Hao, Dengji Zhao, and Tao Zhou. 2017. Mechanism design in social networks. In *Thirty-First AAAI Conference on Artificial Intelligence*.

[8] Frank McSherry and Kunal Talwar. 2007. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*. IEEE, 94–103.

[9] Tianjiao Ni, Zhili Chen, Lin Chen, Shun Zhang, Yan Xu, and Hong Zhong. 2021. Differentially private combinatorial cloud auction. *IEEE Transactions on Cloud Computing* (2021).

[10] David Xiao. 2013. Is privacy compatible with truthfulness? *IACR Cryptol. ePrint Arch.* 2011 (2013), 5.

[11] Jinlai Xu, Balaji Palanisamy, Yuzhe Tang, and S.D. Madhu Kumar. 2017. PADS: Privacy-Preserving Auction Design for Allocating Dynamically Priced Cloud Resources. In *2017 IEEE 3rd International Conference on Collaboration and Internet Computing (CIC)*. 87–96. https://doi.org/10.1109/CIC.2017.00023

[12] Mengxiao Zhang, Fernando Beltran, and Jiamou Liu. 2020. Selling data at an auction under privacy constraints. In *Conference on Uncertainty in Artificial Intelligence*. PMLR, 669–678.

[13] Wen Zhang, Dengji Zhao, and Hanyu Chen. 2020. Redistribution Mechanism on Networks. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 1620–1628.

[14] Wen Zhang, Dengji Zhao, and Yao Zhang. 2020. Incentivize Diffusion with Fair Rewards. In *ECAI 2020*. IOS Press, 251–258.

[15] Dengji Zhao, Bin Li, Junping Xu, Dong Hao, and Nicholas R Jennings. 2018. Selling Multiple Items via Social Networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. 68–76.

[16] Ruihao Zhu, Zhijing Li, Fan Wu, Kang Shin, and Guihai Chen. 2014. Differentially private spectrum auction with approximate revenue maximization. In *Proceedings of the 15th ACM international symposium on mobile ad hoc networking and computing*. 185–194.

[17] Ruihao Zhu and Kang G Shin. 2015. Differentially private and strategy-proof spectrum auction with approximate revenue maximization. In *2015 IEEE conference on computer communications (INFOCOM)*. IEEE, 918–926.