Incentivising Diffusion while Preserving Differential Privacy

Fengjuan Jia[1] **Mengxiao Zhang**[2] **Jiamou Liu**[3] **Bakh Khoussainov**[4]

[1,2,4]University of Electronic Science and Technology of China, China
[3]The University of Auckland, New Zealand

Abstract

Diffusion auction refers to an emerging paradigm of online marketplace where an auctioneer utilises a social network to attract potential buyers. Diffusion auction poses significant privacy risks. From the auction outcome, it is possible to infer hidden, and potentially sensitive, preferences of buyers. To mitigate such risks, we initiate the study of differential privacy (DP) in diffusion auction mechanisms. DP is a well-established notion of privacy that protects a system against inference attacks. Achieving DP in diffusion auctions is non-trivial as the well-designed auction rules are required to incentivise the buyers to truthfully report their neighbourhood. We study the single-unit case and design two differentially private diffusion mechanisms (DPDMs): recursive DPDM and layered DPDM. We prove that these mechanisms guarantee differential privacy, incentive compatibility and individual rationality for both valuations and neighbourhood. We then empirically compare their performance on real and synthetic datasets.

1 INTRODUCTION

New technological shift in AI and data science has given rise to an imminent need to address data privacy issues in online platforms. Indeed, a Gartner survey shows that 41% of the surveyed organisations have experienced a privacy breach or security incident[1]. Data privacy issues have been especially serious and impactful around the use of social commerce platforms such as Instagram and Facebook. As users of such a platform find, browse and buy products through the social network, they are also exposed to a significant risk of privacy leakage. A recent PCI Pal survey shows that fewer than 7% of users are confident about their data security on social commerce sites[2]. Thus designing new tools to facilitate safe and private use of social commerce platforms is of crucial importance.

Auction is important in facilitating online commerce. Auctions have been applied in many contexts, e.g., radio spectrum, sponsored search ads, virtual resource allocation. 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 auctions have recently been studied in [McSherry and Talwar, 2007, Jian et al., 2018, Ni et al., 2021, Zhang et al., 2020a, 2023]. To mitigate privacy risks, these studies employ the well-established notion of *differential privacy* (DP) [Dwork et al., 2006]. Here, DP is used to protect individual's bid information when the auction outcome is published. To achieve DP on bids, the work of McSherry and Talwar [2007] proposed *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. This mechanism has so far been a predominant method to protect privacy in auctions.

Diffusion auction is an emerging form of auction. 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. Thus, diffusion auction is especially suitable for facilitating online social commerce platforms where the social network [Liu and Wei, 2017] plays a prominent role. A challenge in diffusion auctions lies in resolving the

[1]https://blogs.gartner.com/avivah-litan/2022/08/05/ai-models-under-attack-conventional-controls-are-not-enough/

[2]https://www.pcipal.com/knowledge-centre/resource/fewer-than-10-of-people-are-confident-about-their-data-security-on-social-media-according-to-survey-from-pci-pal/

Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence (UAI 2023), PMLR 216:963–972.

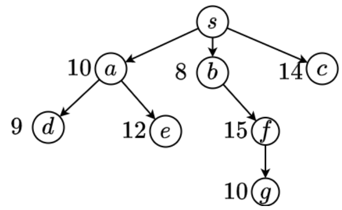

Figure 1: A social network with a seller s and seven buyers. The number beside each node is the valuations of the buyer. The seller s has an item to sell, and initially knows only a, b, c. The mechanism will construct a probability distribution over potential buyers which determines how likely a buyer is to win the item.

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., [Li et al., 2017, 2019, Zhang et al., 2020c,b], have proposed mechanisms for diffusion auction that achieve diffusion IC.

Diffusion auctions are prone to the aforementioned privacy risks for auctions in general. 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?*

Answering this question is not a trivial task. As mentioned above, the exponential mechanism is the main approach to ensure DP for auctions. An exponential mechanism firstly creates a probability distribution over all possible auction results such that more preferable result is associated with a higher probability, and then outputs an auction result according to the distribution. However, this mechanism can not be directly extended to diffusion auctions as it fails to ensure diffusion IC property. For instance, run the exponential mechanism to the scenario in Figure 1 (See Example 4.1 for a detailed implementation). Assume that all buyers except buyer b reveal their neighbours truthfully. From b's perspective, revealing her neighbour f means getting a lower probability of winning the auction, as the exponential mechanism would distribute the winning probabilities over 7 buyers instead of 5. Therefore, the buyers are not incentivised to diffuse auction information to their friends.

Contribution. In this paper, we design DPDM for the case where a single seller sells an indivisible item to multiple potential buyers. The seller and the buyers are assumed to be nodes in a social network with their connections represented as edges. The seller initially only has access to her direct neighbours, and must incentivise the buyers to truthfully report their valuations of the item, and diffuse the auction information to their neighbours. At the same time, the DPDM should ensure the DP property for buyers' bids.

To this end, we design two DPDMs: **recursive DPDM** and **layered DPDM**. The idea for these two mechanisms is *market division* that partitions the buyers into sub-markets. The mechanism then associates a probability with each sub-market. To ensure diffusion IC, the probability should be monotonic on the size of the sub-markets:

- The recursive DPDM maps the network into a tree that captures information flow among buyers. Then it recursively divides the market such that each sub-tree is a sub-market and its probability is non-decreasing on the size of the sub-tree.
- The layered DPDM also relies on the tree above, except the market is not partitioned by sub-trees, but rather by buyers' distances from the seller. In this way, each layer is a sub-market and its probability is fixed.

These two mechanisms are proven to meet desirable incentive and privacy properties. The layered DPDM has a lower bound on expected social welfare. The recursive DPDM achieves a better social welfare empirically. We demonstrate this using a series of experiments that simulate diffusion auctions over three real-world social network datasets. Our experiments reveal that in most cases, the recursive DPDM reaches comparable social welfare as the theoretical upper bound. We now highlight our contributions:

1. We expand diffusion mechanisms adding the DP condition. This builds a bridge between diffusion auctions and privacy preservation. See Section 3.
2. Using the idea of market division, we design recursive DPDM (Section 4) and layered DPDM (Section 5). These mechanisms are IC and differentially private.
3. We empirically evaluate our two mechnaisms on real-world network datasets. See Section 6.

2 RELATED WORK

Differentially private mechanism. Differential privacy (DP) is proposed to protect individual data from inference attacks on aggregate queries over a database [Dwork et al., 2006]. The notion has since been extended to various domains such as statistical data inference [Dwork, 2008] and decision trees [Fletcher and Islam, 2019]. McSherry and Talwar [2007] extend DP to auctions and propose exponential mechanism. This mechanism ensures a weaker version of IC, approximate IC, which ensures any user can only gain a bounded extra utility from misreporting. This solution concept is adopted in subsequent studies [Zhu et al., 2014] and [Diana et al., 2020] on multi-item auctions and double auctions. As approximate IC allows bidders to have non-zero incentives to lie, these methods would not meet the requirements in our problem.

Many work design DP auctions that ensure the classical version of IC [Huang and Kannan, 2012, Xiao, 2013, Zhu and Shin, 2015, Jian et al., 2018]. Specifically, Xiao [2013], Huang and Kannan [2012] proposes general methods to transform a classical IC mechanism to a privacy preserving counterpart that is still IC. However, [Xiao, 2013]'s method works only when the valuation space is small and can not be applied to general problems, including ours. In contrast, [Huang and Kannan, 2012]'s method can be seen as a generalisation of Vickrey-Clarke-Groves (VCG) mechanism [Groves, 1973] paired with a carefully designed payment rule, and thus is applicable to general problems. Yet, when the mechanism is applied to auctions, it is only approximately IC. Later, Zhu and Shin [2015] and Xu et al. [2017], Jian et al. [2018] propose mechanisms that combine the exponential mechanism with the payment rule in [Archer and Tardos, 2001], applying to combinatorial auctions and reverse auctions, respectively.

No mechanism above can be applied as DPDM in our problem because they fail to ensure diffusion IC.

Diffusion auction mechanisms. Diffusion auction is an emerging topic in mechanism design. Li et al. [2017] are the first to investigate diffusion auction and they propose information diffusion mechanism (IDM), a mechanism for single-unit auction in a social network. The basic idea is to give monetary reward to buyers who are critical to diffusion, and it ensures diffusion IC. Following this idea, Li et al. [2019], Zhang et al. [2020c,b] study single-unit diffusion auction from other various aspects. Later, Zhao et al. [2018], Kawasaki et al. [2020] extend single-unit diffusion auctions to multiple-unit cases and propose generalised IDM and DNA-MU, resp. All of these mechanisms are deterministic and none addresses privacy leakage risks.

3 PROBLEM FORMULATION

3.1 PRELIMINARIES

Consider the following setup: There is a seller, denoted by s, and n buyers, denoted by $N = \{1, 2, 3, ..., n\}$. 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. Valuation v_i is the maximum amount of money that i is willing to pay. This value is private to the buyer.

The seller and the buyers form a social network, represented by a graph $G = (V, E)$, where $V = N \cup \{s\}$ and $E \subseteq V^2$. Each node $i \in V$ has a neighbour set, denoted by $r_i := \{j \in V \mid (i, j) \in E\}$. This set is a private for buyer i. The pair (v_i, r_i) is called the *true profile of the buyer* i.

During an auction, the seller would like to attract more buyers to the auction and spread the auction information. Initially, only the seller's neighbours are invited to the auction. 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 a profile θ', we set $\theta'_{-i} := (\theta'_1, \ldots \theta'_{i-1}, \theta'_{i+1}, \ldots \theta'_n)$ to denote the profile of all buyers except i. 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*.

Diffusion auctions have two forms of information asymmetry: **(1) Valuation asymmetry.** The buyers' true valuations are private and hidden from the seller. Thus buyers have an advantage over the seller as they can misreport their valuations. The auction should prevent misreporting of valuation through appropriate strategies. **(2) Neighbourhood asymmetry.** By Bulow-Klemperer theorem, the revenue of an auction increases as the number of buyers grows [Bulow and Klemperer, 1996]. However, as buyers' neighbours are hidden, the seller would hope the buyers to diffuse the auction information to their neighbours to allow more participants to join. Being rational, the buyers are not necessarily willing to disseminate the auction information as this may hinder their own chance of winning. Hence buyer i can misreport the neighbour set $r'_i \subseteq r_i$.

Definition 3.1. *A* mechanism M *consists of two functions* $(\pi(\cdot), p(\cdot))$*, where* $\pi\colon \Theta \to \{0,1\}^n$ *is an* allocation function *and* $p\colon \Theta \to \mathbb{R}^n$ *is a* payment function*.*

Thus, a mechanism M takes the reported profile $\theta' \in \Theta$ as input, and outputs $(\pi(\cdot), p(\cdot))$. The function $\pi(\cdot)$ determines which buyer gets the item, and the function $p(\cdot)$ determines the amount that each buyers pays. We write the *allocation result* $\pi(\theta')$ as $(\pi_1(\theta'), \ldots, \pi_n(\theta'))$ and the *payment result* $p(\theta')$ as $(p_1(\theta'), \ldots, p_n(\theta'))$. 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 M on θ', written $sw_M(\theta')$, is the sum of all utilities, i.e., $sw_M(\theta') = \sum_{i \in V} u_i(\theta')$. We aim to maximise the social welfare.

3.2 PRIVACY-AWARE DIFFUSION AUCTION

On top of the two challenges (1) and (2) provided by information asymmetry, another important challenge is **(3) Valuation privacy.** Once the auction result is announced, an attacker may infer the bid information from the published auction result. This is known as the *inference attack* [Li et al., 2017]. This disadvantages the buyer(s) whose private valuation is disclosed. Therefore, the buyers require the guarantees that their private valuations are protected. So, for privacy preservation, we use randomisation.

Definition 3.2. *A* randomised mechanism M *is one that,* *given $\theta' \in \Theta$, outputs (π, p) such that π and p are randomised allocation and payment functions, respectively.*

Given $\theta' \in \Theta$, $\pi(\theta')$ is a random variable with values $\{0,1\}^n$, and $p(\theta')$ is a random variable with values in $(\mathbb{R}^+)^n$. We use the concept of differential privacy to define the privacy protection of M. Differential privacy requires that the distributions over the outcomes are nearly identical when global profiles are nearly identical. The privacy protection level is measured by a privacy parameter $\epsilon \in \mathbb{R}^+$.

Definition 3.3. *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 $o \in O$,*

$$\Pr[M(\theta') = o] \leq \exp(\epsilon)\Pr[M(\theta'') = o] \qquad (1)$$

Eqn. (1) shows if any buyer i changes her reported profile from $\theta'_i = (v'_i, r'_i)$ to $\theta''_i = (v''_i, r'_i)$, the auction outcome does not change much. Therefore, no one could infer the valuation of any buyer from the randomised outcome.

Exponential mechanism [McSherry and Talwar, 2007] ensures ϵ-DP for valuation privacy. Given a global profile, an exponential mechanism creates a distribution over all possible auction outcomes, and outputs an outcome according to the distribution. Intuitively, the higher a reported valuation is, the more likely the corresponding buyer gets the item. Specially, given a global profile θ', define a *score function* $\sigma\colon \Theta \times O \to \mathbb{R}$ that assigns a real valued score to each pair (θ', o) from $\Theta \times O$. The more preferable an outcome o is, the higher the score of o is. An exponential mechanism $M(\theta')$ outputs a result $o^* \in O$ with probability

$$\frac{\exp(\epsilon\sigma(\theta', o^*))}{\sum_{o\in O}\exp(\epsilon\sigma(\theta', o))}$$

In our problem, we use o_i to denote the outcome where buyer $i \in N$ gets the item.

In randomised mechanisms, we assume that the buyers are risk-neutral and care about their utilities in expectation. We use $\mathbf{E}_M[u_i(\cdot)]$ to denote i's expected utility in M and redefine the IC and IR properties by expected utility.

Definition 3.4. *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, $\mathbf{E}_M[u_i((\theta_i, \theta'_{-i}))] \geq \mathbf{E}_M[u_i((\theta'_i, \theta''_{-i}))]$.*
- *The mechanism M is* IR *if for all $i \in N$ and all $\theta'_{-i} \in \Theta^{n-1}$, we have $\mathbf{E}_M[u_i((\theta_i, \theta'_{-i}))] \geq 0$.*

The IR and IC properties ensure that buyers participate in the auction and reveal their true profiles as they are rational and doing so leads to the best expected utilities. Hence, information asymmetry issues can be addressed.

The social welfare of M is also in expectation, i.e.,

$$\mathbf{E}_M[sw_M(\theta)] = \sum_{i\in V}\mathbf{E}_M[u_i(\theta)].$$

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

4 RECURSIVE DPDM

Preserving valuation privacy in diffusion auctions is not a trivial task. On one hand, existing diffusion auctions, including IDM [Li et al., 2017], CMD [Li et al., 2019], and FDM [Zhang et al., 2020c], are deterministic, and thus fail to preserve privacy. On the other hand, existing DP mechanisms, including exponential mechanism, fail to incentivise truthful report of neighbours, as illustrated in Example 4.1.

Example 4.1. *We apply the exponential mechanism paired with score function $\sigma(\theta', o_i) = v'_i$ to the scenario in Fig. 1. Assume that the buyers truthfully report their valuations. Then buyer i wins with probability $\exp(\epsilon v_i)/\sum_{\kappa\in N}\exp(\epsilon v_\kappa)$. If buyer b reports her neighbour f, b wins with probability $\exp(8\epsilon)/\sum_{\kappa\in N}\exp(\epsilon v_\kappa)$, whereas she wins with probability $\exp(8\epsilon)/\sum_{\kappa\in N\setminus\{f,j\}}\exp(\epsilon v_\kappa)$ had she chose not to report f. In the latter case, the winning probability is even higher, and thus b has incentive to hide her neighbours.*

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, as shown in Eqn. (2). 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, as shown in Eqn. (3). Therefore, the buyers have no competition with their children and have no incentive to hide them.

We now describe REC in detail: 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) Construction of diffusion critical tree. From the profile digraph $G_{\theta'}$, REC constructs a *diffusion critical tree* $T_{\theta'}$. When the context is clear, we write the tree as T. The diffusion critical trees are introduced in [Zhao et al., 2018]. For buyers i, j, we say that i is *θ'-critical* to j, written $i \preceq_{\theta'} j$, if all paths from s to j in $G_{\theta'}$ go through i. The root of the tree $T_{\theta'}$ is s, the nodes $V_{\theta'}$ are the buyers, and for each

$j \in V_{\theta'}$, her parent is the node $i \preceq_{\theta'} j$ who has the closest distance to j. When there are more than one parents, only one node is randomly selected as the parent. The *depth* of buyer i, denoted by d_i, is the distance from s to i.

(2) Assignment of 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 sub-tree 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.

(a) Assignment of probabilities to sub-trees. Let $T[i]$ denote the sub-tree rooted by i. $T[i]$ consists of node i and all of i's descendants. Let $T(i)$ denote $T[i]$ with i removed, i.e., $T(i) := T[i] \setminus \{i\}$. Given a sub-tree $T[i]$, REC divides the market in $T[i]$ to $|r_i| + 1$ sub-markets, one for i and each of the other for a sub-tree $T[j]$, where $j \in r_i$. Then REC assigns a probability $\mathrm{Pr}_i^{\theta'}(\theta'_i)$ to i with θ'_i and $\mathrm{Pr}_{T[j]}^{\theta'}$ to each $T[j]$, where $j \in r_i$. When the context is clear, we write Pr_i and $\mathrm{Pr}_{T[j]}$ for $\mathrm{Pr}_i^{\theta'}(\theta'_i)$ and $\mathrm{Pr}_{T[j]}^{\theta'}$, respectively. We define Pr_i later in Step (2).b. For notational convenience, given a set of nodes $S \subseteq T$, we let $\mathrm{Exp}(S)$ be the sum

$$\mathrm{Exp}(S) = \sum_{\kappa \in S} \exp(\epsilon \sigma(\theta', o_\kappa)).$$

Now we define $\mathrm{Pr}_{T[j]}$ for each $j \in r_i$ as

$$\mathrm{Pr}_{T[j]} = \left(\mathrm{Pr}_{T[i]} - \mathrm{Pr}_i\right) \times \frac{\mathrm{Exp}(T[j])}{\mathrm{Exp}(T(i))} \tag{2}$$

(b) Assignment of winning probabilities to buyers within a sub-market. In a sub-tree $T[i]$, REC assigns the winning probability Pr_j to each $j \in r_i$ as

$$\mathrm{Pr}_j = \left(\mathrm{Pr}_{T[i]} - \mathrm{Pr}_i\right) \times \frac{\mathrm{Exp}(j)}{\mathrm{Exp}(T(i) \setminus T(j))} \tag{3}$$

At the very beginning, REC starts with the tree T rooted by s. We label s as node 0 and set $\mathrm{Pr}_{T[0]} = 1$ and $\mathrm{Pr}_0 = 0$. REC ends with the leaves. For a sub-tree $T[i]$ where each $j \in r_i$ are leaves, REC assigns the winning probability to each j as $\mathrm{Pr}_j = \left(\mathrm{Pr}_{T[i]} - \mathrm{Pr}_i\right) \times \frac{\mathrm{Exp}(j)}{\mathrm{Exp}(T(i))}$.

(3) Allocation and payment. Randomly select a buyer w as a winner according to the constructed distribution in Step (2). Set w's allocation $\pi_w = 1$, and payment as

$$p_w = v'_w - \int_0^{v'_w} \mathrm{Pr}_w((x, r'_w))dx / \mathrm{Pr}_w(\theta'_w) \tag{4}$$

We present the details of REC in Algorithm 1 and give a running example of Step (2) in Example 4.2.

Algorithm 1 Recursive DPDM REC

Input: Reported global profile θ', privacy parameter ϵ and score function σ
Output: Allocation result $\pi(\theta')$ and payment result $p(\theta')$
1: Initialise $\pi(\theta') = \mathbf{0}, p(\theta') = \mathbf{0}$
2: Construct a profile digraph $G_{\theta'} = (V_{\theta'}, E_{\theta'})$
3: Construct a critical diffusion tree $T_{\theta'}$
4: Run GetPro($T_{\theta'}[0], 1, 0$)
5: Randomly select a buyer w with the distribution
6: Set $\pi_w = 1$ and p_w by Equation (4)

Algorithm 2 GetPro

Input: (Sub-)Tree $T[i]$, probabilities $\mathrm{Pr}_{T[i]}$ and Pr_i
Output: Probabilities $\mathrm{Pr}_{T[j]}$ and Pr_j, $j \in r_i$
1: **for** $j \in r_i$ **do**
2: Calculate $\mathrm{Pr}_{T[j]}$ of sub-tree $T[j]$ by Equation (2)
3: Calculate Pr_j of buyer j by Equation (3)
4: Run GetPro($T[j], \mathrm{Pr}_{T[j]}, \mathrm{Pr}_j$)
5: **end for**

Example 4.2. *Apply* REC *to scenario in Fig. 1, with* $\sigma(\theta, o_i) = v'_i$*. So,* $\mathrm{Pr}_T = 1$ *and* $\mathrm{Pr}_s = 0$*. Calculate the probabilities of* s*'s children. The probability for* $T[a]$ *is* $\mathrm{Pr}_{T[a]} = (\exp(10\epsilon) + \exp(9\epsilon) + \exp(12\epsilon)) / \mathrm{Exp}(T)$*. Buyer* a *wins with probability* $\mathrm{Pr}_a(10) = \exp(10\epsilon) / (\mathrm{Exp}(T) - (\exp(9\epsilon) + \exp(12\epsilon)))$*. Similarly, we get the probabilities for* $T[b], T[c]$ *and* b, c*. Buyer* d *wins with probability* $\mathrm{Pr}_d(9) = (\mathrm{Pr}(T[a]) - \mathrm{Pr}_a) \times \exp(9\epsilon) / (\exp(9\epsilon) + \exp(12\epsilon))$*. We can also get the probabilities for* e, f, g*.*

Lemma 4.3. *Recursive DPDM* REC *is individually rational in terms of both valuations and neighbours.*

Proof. Given a global profile θ, for each buyer i with (v_i, r_i), $\mathbf{E}_{\mathrm{REC}}[u_i(\theta)] = (v_i - p_i(\theta))\mathrm{Pr}_i(\theta_i) = \int_0^{v_i} \mathrm{Pr}_i((x, r_i))dx \geq 0$. Therefore, the lemma holds. ☐

To show that REC satisfies IC we need the following:

Theorem 4.4 ([Archer and Tardos, 2001])**.** *Let* $\mathrm{Pr}_i(v'_i)$ *be the winning probability assigned by mechanism* M *when buyer* i *reports* v'_i*. Then* M *is IC in terms of valuations iff for all* $i \in N$*: (1)* $\mathrm{Pr}_i(v'_i)$ *is monotonically non-decreasing in* v'_i*, and (2)* $\mathbf{E}[p_i] = v_i \mathrm{Pr}_i(v'_i) - \int_0^{v'_i} \mathrm{Pr}_i(x)dx$*.*

Lemma 4.5. *Recursive DPDM* REC *is incentive compatible in terms of both valuations and neighbours.*

Proof. We first show REC is IC in terms of valuations. By Equation (2), the probability for any sub-tree $T[i]$ is proportional to the score, which is non-decreasing in v'_i. Hence, $\mathrm{Pr}_{T[i]}$ in non-decreasing in v'_i. Similarly, by Equation (3), given a sub-tree $T[i]$, the winning probability Pr_i is non-decreasing in v'_i, which meets the condition (1) in Thm. 4.4. Also, by Equation (4), the expected payment

$$\mathbf{E}[p_i] = p_i \times \mathrm{Pr}_i = v'_i \mathrm{Pr}_i(\theta'_i) - \int_0^{v'_i} \mathrm{Pr}_i((x, r'_i))dx,$$

which meets the condition (2) in Theorem 4.4 when r'_i is fixed. Therefore, REC is IC in terms of valuations.

Next we show REC is **IC** in terms of neighbours. By the definitions of expected utility and payment function (4), we know that i's expected utility is only determined by the winning probability Pr_i. Let a^ℓ be an ancestor of i with distance ℓ. When i reports truthfully as θ_i and the reported global profile is θ'_{-i}, then i's winning probability is

$$\begin{aligned}
\mathrm{Pr}_i =& \frac{\mathrm{Exp}(i)}{\mathrm{Exp}(T(a^1) \setminus T(i))} \times (\mathrm{Pr}_{T[a^1]} - \mathrm{Pr}_{a^1}) \\
=& \frac{\mathrm{Exp}(i)}{\mathrm{Exp}(T(a^1) \setminus T(i))} \times \left(\mathrm{Pr}_{T[a^2]} - \mathrm{Pr}_{a^2}\right) \\
&\times \left(\frac{\mathrm{Exp}(T[a^1])}{\mathrm{Exp}(T(a^2))} - \frac{\mathrm{Exp}(a^1)}{\mathrm{Exp}(T(a^2) \setminus T(a^1))}\right) \\
=& \frac{\mathrm{Exp}(i)}{\mathrm{Exp}(T(a^1) \setminus T(i))} \times (\mathrm{Pr}_T - \mathrm{Pr}_s) \\
&\times \prod_{\ell=1}^{d_i-1} \left(\frac{\mathrm{Exp}(T[a^\ell])}{\mathrm{Exp}(T(a^{\ell+1}))} - \frac{\mathrm{Exp}(a^\ell)}{\mathrm{Exp}(T(a^{\ell+1}) \setminus T(a^\ell))}\right)
\end{aligned} \tag{5}$$

If i hides some of her neighbours and reports any θ'_i where $r'_i \subseteq r_i$, instead, and the others report θ'_{-i}. Then in Eqn. (5), Pr_T, Pr_s and $\frac{\mathrm{Exp}(i)}{\mathrm{Exp}(T(a^1)\setminus T(i))}$ does not change. Also, for each ℓ, $\frac{\mathrm{Exp}(a^\ell)}{\mathrm{Exp}(T(a^\ell)\setminus T(a^{\ell+1}))}$ remains intact, but $\frac{\mathrm{Exp}(T[a^\ell])}{\mathrm{Exp}(T(a^{\ell+1}))}$ decreases. So we can know that Pr_i decreases when i misreports her neighbourhood. Therefore, we have $\mathbf{E}_{\mathrm{REC}}[u_i(((v_i, r_i), \theta'_{-i}))] \geq \mathbf{E}_{\mathrm{REC}}[u_i(((v_i, r'_i), \theta''_{-i}))]$. □

Lemmas 4.3 and 4.5 show the recursive mechanism REC incentivises buyers to reveal their true profiles and thus addresses valuation asymmetry and neighbourhood asymmetry. Next we show that REC also addresses the valuation privacy issue. In following lemma, we use the following notations.

- $d_{\max}$: the maximum depth of the diffusion critical tree,
- $\Delta\sigma$: 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 $o_i \in O$.

Lemma 4.6. *Given a reported global profile θ', recursive DPDM* REC *is $\epsilon d_{\max}\Delta\sigma$-differentially private, where ϵ is the DP parameter of* REC.

Proof. Let θ and θ' be two profiles where a buyer i's reports i reports v_i in θ and v'_i in θ' such that $v_i \neq v'_i$. Consider the probabilities that $\mathrm{REC}(\theta)$ and $\mathrm{REC}(\theta')$ return a winner w. In a critical diffusion tree T_θ, let d_w denote the depth of w, a_w^ℓ be an ancestor of w with distance ℓ. Also, let $\mathrm{Exp}^\theta(T(a_w^1) - T(w))$ and $\mathrm{Exp}^{\theta'}(T(a_w^1) - T(w))$ denote the value derived from θ and θ', respectively. Then by Equation (3), we have

$$\frac{\mathrm{Pr}[\mathrm{REC}(\theta) = o_w]}{\mathrm{Pr}[\mathrm{REC}(\theta') = o_w]} = \frac{\frac{\mathrm{Exp}(w)}{\mathrm{Exp}^\theta(T(a_w^1)-T(w))}}{\frac{\mathrm{Exp}^{\theta'}(w)}{\mathrm{Exp}^{\theta'}(T(a_w^1)-T(w))}} \times \frac{\mathrm{Pr}^\theta_{T[a_w^1]} - \mathrm{Pr}^\theta_{a_w^1}}{\mathrm{Pr}^{\theta'}_{T[a_w^1]} - \mathrm{Pr}^{\theta'}_{a_w^1}}$$

We repeatedly replace $\mathrm{Pr}^\theta_{T[a_w^\ell]}$, $\mathrm{Pr}^\theta_{a_w^\ell}$, $\mathrm{Pr}^{\theta'}_{T[a_w^\ell]}$, $\mathrm{Pr}^{\theta'}_{a_w^\ell}$ by expressions of $a_w^{\ell+1}$ until we get an expression of s. For each distance $0 \leq \ell < d_w$, we denote $\frac{\mathrm{Exp}(T[a_w^\ell])}{\mathrm{Exp}(T(a_w^{\ell+1}))}$ as A_ℓ^θ, $\frac{\mathrm{Exp}(a_w^\ell)}{\mathrm{Exp}(T(a_w^{\ell+1})\setminus T(a_w^\ell))}$ as B_ℓ^θ. For θ', we have similar notations as $A_\ell^{\theta'}$ and $B_\ell^{\theta'}$. Then the above ratio can be written as

$$\frac{\mathrm{Pr}[\mathrm{REC}(\theta) = o_w]}{\mathrm{Pr}[\mathrm{REC}(\theta') = o_w]} = \frac{B_0^\theta}{B_0^{\theta'}} \times \prod_{\ell=1}^{d_w-1} \frac{A_\ell^\theta - B_\ell^\theta}{A_\ell^{\theta'} - B_\ell^{\theta'}}$$

Next we prove the lemma through that for each $0 \leq \ell < d_w$, $\frac{A_\ell^\theta - B_\ell^\theta}{A_\ell^{\theta'} - B_\ell^{\theta'}}$ is bounded by $\exp(\epsilon\Delta\sigma)$. Here we skip this due to space limit. See details in **App. B**. Then we have

$$\begin{aligned}
\frac{\mathrm{Pr}[\mathrm{REC}(\theta) = o_w]}{\mathrm{Pr}[\mathrm{REC}(\theta') = o_w]} &\leq \exp(\epsilon\Delta\sigma) \times \prod_{\ell=1}^{d_w-1} \exp(\epsilon\Delta\sigma) \\
&\leq \exp(\epsilon d_w \Delta\sigma) \leq \exp(\epsilon d_{\max}\Delta\sigma)
\end{aligned}$$

□

Next theorem easily follows from Lemmas 4.5, 4.3 & 4.6.

Theorem 4.7. *Recursive DPDM* REC *is IC, IR and $\epsilon d_{\max}\Delta\sigma$-DP.*

5 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. Specifically, given a critical diffusion tree, 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.

LAY executes the same operations as in REC, where the only difference is in Step (2) "Assignment of winning probabilities". Below we describe Step (2) of LAY in detail.

(2) Assignment of winning probabilities. In this step, given a critical diffusion tree $T_{\theta'}$, LAY assigns a probability to each layer of the tree and then assigns a winning probability to buyers on each layer.

(a) Assignment of probability to layers. Now we give the definition of layer. Given a tree, the buyers with the same distance d_i form a *layer* of a tree. The distance $d_i \in \{1, \ldots, d_{\max}\}$. We use L_ℓ to denote the set of buyers with distance ℓ, i.e., $L_\ell \coloneqq \{i \mid d_i = \ell\}$. For each layer $L_\ell, 1 \leq \ell \leq d_{\max}$, LAY assigns a probability, denoted by $\Pr^{\theta'}_{L_\ell}$. We write it as $\Pr_{L_\ell}$ when there is no ambiguity. Given an infinite decreasing sequence $\gamma = (\gamma_1, \gamma_2, \ldots)$, where $\sum \gamma_i = 1$, we define the probability for layer L_ℓ as

$$\Pr_{L_\ell} = \gamma_\ell \tag{6}$$

Intuitively, as the layer becomes deeper, the assigned probability for the layer decreases.

(b) Assignment of winning probability to the buyers on a layer. On the ℓth layer, LAY assigns buyer i with θ'_i on layer $d_i = \ell$ with probability

$$\Pr_i(\theta'_i) = \Pr_{L_\ell} \times \frac{\mathrm{Exp}(i)}{\mathrm{Exp}(L_\ell)} \tag{7}$$

Once the probability distribution over all possible outcomes is determined, LAY computes the payment and randomly selects a winner w, following Step (3) of REC.

The complete process of layered DPDM is shown in Alg. 3. Example 5.1 provides a running example of Step (2).

Algorithm 3 Layered DPDM LAY

Input: Reported global profile θ', privacy parameter ϵ and score function σ
Output: Allocation result $\pi(\theta')$ and payment result $p(\theta')$

```
Initialise π(θ') = 0, p(θ') = 0
Construct a profile digraph G_θ' = (V_θ', E_θ')
Construct a critical diffusion tree T_θ'
for 1 ≤ ℓ ≤ d_max do
    Calculate the probability of layer ℓ by Equation (6)
    for i ∈ L_ℓ do
        Calculate winning probability Pr_i by Eqn. (7)
    end for
end for
Randomly select a buyer w with the distribution
Set π_w = 1 and p_w by Equation (4)
```

Example 5.1. *Apply* LAY *paired with* $\sigma(\theta, o_i) = v'_i$ *and sequence* $\gamma = \left\{\frac{1}{2^{\kappa+1}}\right\}_{\kappa\in\mathbb{N}}$ *to the scenario in Figure 1. Then in this graph, three layers,* $L_1 = \{a, b, c\}, L_2 = \{d, e, f\}$, $L_3 = \{g\}$ *correspond to probabilities* $\frac{1}{2}, \frac{1}{4}, \frac{1}{8}$, *resp. In* L_1, *buyer* a *wins with probability* $\exp(10\epsilon)/(2(\exp(10\epsilon) + \exp(8\epsilon) + \exp(14\epsilon)))$. *Similarly, we get the probabilities for* b *and* c. *Then in* L_2, d *wins with probability* $\exp(9\epsilon)/(4(\exp(9\epsilon) + \exp(12\epsilon) + \exp(15\epsilon)))$. *The probabilities for* e, f *can be obtained in a similar way. Lastly, in* L_3, *buyer* g *wins with probability* $\frac{1}{8}$.

Next we show that layered DPDM LAY is IC, IR and DP. Lemmas 5.2 and 5.3 show LAY incentivises buyers to reveal their true profiles and thus addresses valuation asymmetry and neighbourhood asymmetry, while Lemma 5.4 shows that LAY addresses the valuation privacy issue.

Lemma 5.2. *Layered DPDM* LAY *is individually rational in terms of both valuations and neighbours.*

The proof of Lemma 5.2 follows the same reasoning as Lemma 4.3. See details in **Appendix C**.

Lemma 5.3. *Layered DPDM* LAY *is incentive compatible in terms of both valuations and neighbours.*

Proof. The IC property in terms of valuations is proved as in Lemma 4.5. We need to show is $\Pr_i$ is non-decreasing in valuations v'_i. By Eqn. (7), $\Pr_i((v'_i, r'_i))$ is proportional to $\sigma(\theta, o_i)$, which is non-decreasing in v'_i.

To show the IC for reporting neighbours, note that i's expected utility is $\mathbf{E}_{\mathrm{LAY}}[u_i(\theta')] = (v_i - p_i(\theta'))\Pr_i$ for $\theta' \in \Theta$. We plug in Eqn. (4) (7) into $u_i(\theta')$. Then we can see $\Pr_i$ is determined by d_i and d_i is determined by her ancestors. Therefore, misreporting neighbours does not affect her utility, i.e., $\mathbf{E}_{\mathrm{LAY}}[u_i(((v_i, r'_i), \theta_{-i}))] = \mathbf{E}_{\mathrm{LAY}}[u_i(((v_i, r_i), \theta'_{-i}))]$. □

Lemma 5.4. *Given a reported global profile* θ', *layered DPDM* LAY *is* $\epsilon\Delta\sigma$*-differential private, where* ϵ *is the privacy parameter of* LAY.

Lem. 5.4 is proved by showing in Eqn. (7), the change on a single buyer's valuation is bounded by $\epsilon\Delta\sigma$. Due to space limit, the proof of Lem. 5.4 is deferred to **App. D**. The next thm. then easily follows from Lem. 5.2, 5.3 and 5.4.

Theorem 5.5. *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. In this scenario, 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 than all DPDMs and thus is used as the benchmark.

Theorem 5.6. *Given a global profile* θ, *layered DPDM* LAY *has* $\mathbf{E}_{\mathrm{LAY}}[sw_{\mathrm{LAY}}(\theta)] \geq \gamma_{d_{\max}} \mathbf{E}_{\mathrm{EMD}}[sw_{\mathrm{EMD}}(\theta)]$.

See the proof in **App. E**. Following is an easy corollary.

Corollary 5.7. *For* $\gamma = (\frac{a-1}{a}, \frac{a-1}{a^2}, \ldots)$, *where* $a > 1$, $\mathbf{E}_{\mathrm{LAY}}[sw_{\mathrm{LAY}}(\theta)] \geq \frac{a-1}{a^{d_{\max}}} \mathbf{E}_{\mathrm{EMD}}[sw_{\mathrm{EMD}}(\theta)]$ □

6 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. For each setup, we run 5000 times and get average social welfare.

Dataset. We use three real world network datasets, including Hamsterster friendships with $1,858$ nodes and $12,534$ edges [Kunegis, 2013], Facebook with $4,039$ nodes and $88,234$ edges [Leskovec and Mcauley, 2012] and Email-Eu-core network $1,005$ nodes and $25,571$ edges [Yin et al., 2017]. For each dataset, the seller s is randomly selected.

Valuation. The network datasets contain no information about buyers' valuations. We generate random numbers as the valuations. We consider two commonly used distributions, normal distribution $v_i \sim \mu(50, 100)$ and uniform distribution $v_i \sim U[0, 100]$. We set the parameters such that the average value are same. Nevertheless, our aim is to reveal the general pattern under different distributions and these patterns are independent from these parameters.

We also consider the correlated valuations. That is, the valuation of each buyer is influenced by her neighbours. We generate such correlated valuations using DeGroot model [DeGroot, 1974], a mathematical model of social learning. The model first assigns each individual an initial valuation, which are drawn from the uniform distribution as above. Then each individual's valuation is updated by taking a weighted average of her own valuation and the valuations of her neighbours. This process of valuation updating is iterated five times, resulting in a set of correlated valuations.

Privacy parameter. To verify the performance of our mechanisms, we also vary privacy parameter $\epsilon \in \{0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$. Lem. 5.4 and 4.6 show that, under the same input ϵ, LAY and REC ensure different privacy levels. To see the performance under the same guaranteed privacy, we set the input ϵ as $\{0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$ for REC and $\{0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}d_{\max}$ for the others.

Score function. We use linear function, $\sigma(\theta, o_i) = v_i$, as the score function. The linear score function is widely used in previous DP auctions, e.g., [McSherry and Talwar, 2007, Xu et al., 2017].

Decreasing sequence. For LAY, we consider different value of $a \in \{1.25, 1.5, 2, 3\}$ in $\gamma = (\frac{a-1}{a}, \frac{a-1}{a^2}, \ldots)$, and evaluate the impact of a on expected social welfare.

Benchmark. Since there is no existing DPDM that can be applied in our problem, we design two hypothetical benchmarks. **Exponential mechanism without diffusion (EMWD)**: We apply the exponential mechanism only to the seller's neighbours. The expected social welfare of EMWD can be seen as the lower bound among all DPDMs. **Exponential mechanism diffusion (EMD)**: See the description of EMD in Section 5. We also compare with IDM [Li et al., 2017] (See **App. A**), which is not DP, to see how much social welfare is sacrificed to achieve DP.

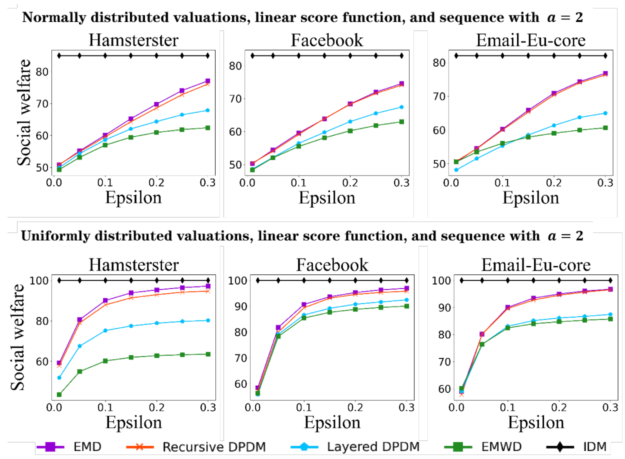

Figure 2: Average social welfare of LAY, REC, EMD, EMWD and IDM with different distributions under linear function and fixed sequence with $a = 2$. Normal distribution is shown in the first row and uniform distribution is shown in the second row.

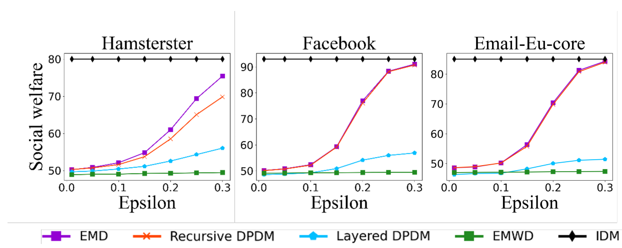

Figure 3: Average social welfare of LAY, REC, EMD, EMWD and IDM under correlated valuations, linear function and fixed sequence with $a = 2$.

Results. 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 lines of REC even coincide with those of EMD in some cases, e.g., on Facebook & Email-Eu-core in Fig. 2. The deviation of REC from EMD is at most 2.62%. REC performs better than the layered counterpart. EMWD returns the worst expected social welfare. The reason why REC has better expected social welfare than LAY is that in LAY, a probability of $1 - \sum_{\ell=1}^{d_{\max}} \gamma_\ell$ is not distributed to any buyer, which means that the seller does not sell the item and the social welfare is 0 with this probability.

Next we show the effect of different parameters. **(1) Dataset.** As shown in each column of Fig. 2, the same pattern can be found for different datasets. **(2) Privacy parameter.** The

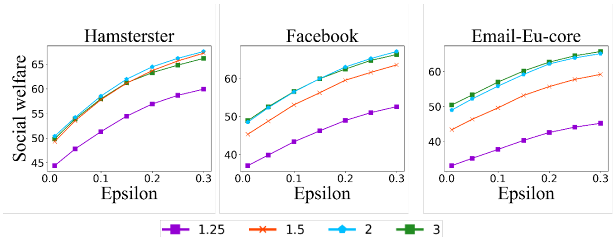

Figure 4: Average social welfare of LAY with different values of a, under normally distributed valuations and linear function.

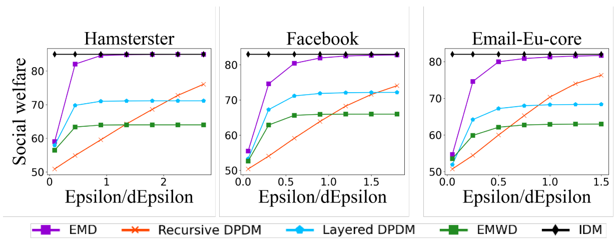

Figure 5: Average social welfare of LAY, REC, EMD, EMWD and IDM under normal distribution, linear function and sequence with $a = 2$. Horizontal axis represents the value of ϵ for EMD, EMWD & LAY, and $d_{\max}\epsilon$ for REC.

expected sw increases with ϵ. The less privacy is required, the less noise is added, and thus the higher probability of returning a result with good social welfare. **(3) Valuation.** The 1st and the 2nd row of Fig. 2 show the results with normal and uniform distributions, resp. Under both distributions, REC performs better than LAY. Under correlated valuations, as shown in Fig. 3, REC also performs better than LAY. **(4) Sequence.** Fig. 4 shows the average social welfare is best when $a = 1.5, 2$ for Hamsterster and when $a = 2, 3$ for Facebook and Email-Eu-core. When a buyer i with the highest valuation is on a deeper layer, a smaller a leads to a larger probability for the layer where i is and also a larger probability for i. The results verify this argument. In Hamsterster (Facebook, Email-Eu-core), the buyers with the highest valuation are on the 4th (3rd, 2nd) layer. **(5) same DP.** Fig. 5 shows when the realised privacy is large, the avg. social welfare of REC is greater than that of LAY, while when the realised privacy is small, LAY is better.

7 CONCLUSION AND FUTURE WORK

We consider designing diffusion auction mechanisms that sells a single item on social networks while preserving valuation privacy. We propose two DPDMs, recursive DPDM and layered DPDM. Also, we theoretically show their incentive and privacy properties and empirically show their good performances in social welfare. We could extend this study by considering: (1) How to design a DPDM for multi-item auctions? (2) How to design a DPDM that preserves both valuation and neighbourhood privacy? and (3) How to design a DPDM that is group IC where no group of buyers can benefit from joint misreporting?

Acknowledgements

This work is partially supported by National Natural Science Foundation of China No. 62172077.

References

Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. *Proceedings 2001 IEEE International Conference on Cluster Computing*, pages 482–491, 2001.

Jeremy Bulow and Paul Klemperer. Auctions versus negotiations. *The American Economic Review*, 86(1):180–194, 1996.

Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical association*, 69(345):118–121, 1974.

Emily Diana, Hadi Elzayn, Michael Kearns, Aaron Roth, Saeed Sharifi-Malvajerdi, and Juba Ziani. Differentially private call auctions and market impact. In *Proceedings of the 21st ACM Conference on Economics and Computation*, page 541–583, 2020.

Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.

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

Sam Fletcher and Md Zahidul Islam. Decision tree classification with differential privacy: A survey. *ACM Computing Surveys (CSUR)*, 52(4):1–33, 2019.

Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

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

Lin Jian, Yang Dejun, Li Ming, Xu Jia, and Xue Guoliang. Frameworks for privacy-preserving mobile crowdsensing incentive mechanisms. In *IEEE Trans. Mob. Comput*, page 1851–1864, 2018.

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

Jérôme Kunegis. Konect: The koblenz network collection. In *Proceedings of the 22nd International Conference on World Wide Web*, WWW '13 Companion, page 1343–1350. Association for Computing Machinery, 2013. doi: 10.1145/2487788.2488173. URL https://doi.org/10.1145/2487788.2488173.

Jure Leskovec and Julian Mcauley. Learning to discover social circles in ego networks. *Advances in neural information processing systems*, 25, 2012.

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

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

Jiamou Liu and Ziheng Wei. Network, popularity and social cohesion: a game-theoretic approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

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

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

David Xiao. Is privacy compatible with truthfulness? *IACR Cryptol. ePrint Arch.*, 2011:5, 2013.

Jinlai Xu, Balaji Palanisamy, Yuzhe Tang, and S.D. Madhu Kumar. Pads: Privacy-preserving auction design for allocating dynamically priced cloud resources. In *2017 IEEE 3rd International Conference on Collaboration and Internet Computing (CIC)*, pages 87–96, 2017. doi: 10.1109/CIC.2017.00023.

Hao Yin, Austin R. Benson, Jure Leskovec, and David F. Gleich. Local higher-order graph clustering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 555–564. Association for Computing Machinery, 2017. doi: 10.1145/3097983.3098069. URL https://doi.org/10.1145/3097983.3098069.

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

Mengxiao Zhang, Jiamou Liu, Kaiyu Feng, Fernando Beltran, and Zijian Zhang. Smartauction: A blockchain-based secure implementation of private data queries. *Future Generation Computer Systems*, 138:198–211, 2023.

Wen Zhang, Dengji Zhao, and Hanyu Chen. Redistribution mechanism on networks. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1620–1628, 2020b.

Wen Zhang, Dengji Zhao, and Yao Zhang. Incentivize diffusion with fair rewards. In *ECAI 2020*, pages 251–258. IOS Press, 2020c.

Dengji Zhao, Bin Li, Junping Xu, Dong Hao, and Nicholas R Jennings. Selling multiple items via social networks. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 68–76, 2018.

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

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