

Private Pareto Optimal Exchange

SAMPATH KANNAN, University of Pennsylvania
JAMIE MORGENSTERN, Carnegie Mellon University
RYAN ROGERS, University of Pennsylvania
AARON ROTH, University of Pennsylvania

We consider the problem of implementing an individually rational, asymptotically Pareto optimal allocation in a barter-exchange economy where agents are endowed with goods and preferences over the goods of others, but may not use money as a medium of exchange. Because one of the most important instantiations of such economies is kidney exchange – where the “input” to the problem consists of sensitive patient medical records – we ask to what extent such exchanges can be carried out while providing formal privacy guarantees to the participants. We show that individually rational allocations cannot achieve any non-trivial approximation to Pareto optimality if carried out under the constraint of differential privacy – or even the relaxation of *joint*-differential privacy, under which it is known that asymptotically optimal allocations can be computed in two sided markets [Hsu et al. STOC 2014]. We therefore consider a further relaxation that we call *marginal*-differential privacy – which promises, informally, that the privacy of every agent i is protected from every other agent $j \neq i$ so long as j does not collude or share allocation information with other agents. We show that under marginal differential privacy, it is possible to compute an individually rational and asymptotically Pareto optimal allocation in such exchange economies.

1. INTRODUCTION

Consider the following exchange problem: for each $i \in [n]$, agent i arrives at a market endowed with a good of type g_i from among a finite collection of types of goods \mathcal{G} , as well as a *preference over types of goods*, represented by a total ordering \succ_i over \mathcal{G} . In settings where money can be used as the medium of transaction (and for which people have cardinal preferences over goods), this would represent an exchange economy for which we could compute market clearing prices, resulting in a Pareto optimal allocation. However, in some settings – most notably markets for kidney exchanges¹ [Roth et al. 2005] – the use of money is not permitted, and agents are limited to participating in an exchange – essentially a permutation of the goods amongst the agents, with no additional payments. In such settings, we may still require that:

- (1) The exchange be *individually rational* – i.e. every agent weakly prefers the good that they receive to the good that they were endowed with, and
- (2) the exchange be (approximately) *Pareto optimal* – that there should not be any other permutation of the goods that some agent (or, in the approximate case, many agents) strictly prefers, unless there exists another agent who strictly prefers the original allocation.

Because one of the key applications of efficient barter-exchange involves computation on extremely sensitive medical data, this paper investigates to what extent it can be accomplished while guaranteeing a formal notion of *privacy* to the agents participating in the market, without compromising the above two desiderata of individual rationality and (approximate) Pareto optimality. We wish to give algorithms that protect the privacy of agents’ initial endowment as well as their preferences over goods from the other agents in the market, using the tools of differential privacy.

In this respect, our paper continues a recent line of work exploring the power of algorithms satisfying various privacy definitions (relaxations of *differential privacy*) in

¹Kidney exchange forms one of the most notable “barter” markets, but it is not the only example. A number of startups such as “TradeYa” and “BarterQuest” act as market makers for barter exchange of consumer goods.

different kinds of exchange problems. To understand this question, note that the *input* to an algorithm clearing an exchange economy is partitioned amongst the n agents (each agent reports her initial endowment g_i and her preference ordering \succ_i), as is the output (the mechanism reports to each agent the type of good she will receive, g'_i .) This allows us to parameterize privacy guarantees in terms of adversaries who can see differing sets of outputs of the mechanism. The standard notion of *differential privacy*, requires that we guarantee privacy even against an adversary who can see all n outputs of the mechanism – i.e. the type of good that is allocated to each of the n agents. It is intuitively clear that nothing non-trivial can be done while guaranteeing differential privacy in allocation problems – “good” allocations must give individuals what they want, and this is exactly what we must keep private. An adversary who is able to see the allocation given to agent i by any mechanism which guarantees individual rationality would immediately learn the relative ranking of the good that agent i was allocated compared to the good that he was endowed with. However, this does not rule out the possibility of protecting the privacy of agent i against an adversary who can only see the allocation of *some* agents– notably, not the allocation given to agent i .

1.1. Our Results

The question of privately computing allocations where agents’ preferences are sensitive data was first studied by Hsu et al. [2014a], who showed that in two sided allocation problems (with distinguished *buyers* and *sellers*) with monetary transfers, no non-trivial allocation can be computed under the constraint of differential privacy, when we must protect the privacy of the buyers. However, Hsu et al. [2014a] showed that near-optimal results can be achieved under *joint differential privacy*, which informally requires that for every agent i simultaneously, the joint distribution on allocations given to agents $j \neq i$ be differentially private in the data of agent i . This corresponds to privacy against an adversary who can see the allocation of *all other agents* $j \neq i$, but who cannot observe agent i ’s own allocation when trying to violate i ’s privacy.

The allocation problem we study in this paper is distinct from the two sided problem studied in [Hsu et al. 2014a] in that there are no distinguished buyers and sellers – in our barter exchange problem, every agent both provides a good and receives one, and so we must protect the privacy of every agent in the market. We insist on algorithms that always guarantee *individually rational* allocations, and ask how well they can approximate *Pareto optimality*. (Informally, an allocation π is α -approximately Pareto optimal if for every other allocation π' that is strictly preferred by an α -fraction of agents, there must be some other agent who strictly prefers π to π' .) We start by showing that even under the relaxed notion of *joint* differential privacy, no individually rational mechanism can achieve any nontrivial approximation to Pareto optimality (and, since joint differential privacy is a relaxation of differential privacy, neither can any differentially private mechanism).

THEOREM. (Informal) *No ϵ -jointly differentially private algorithm for the exchange problem that guarantees individually rational allocations can guarantee with high probability that the resulting allocation will be α -approximately Pareto optimal for:*

$$\alpha \leq 1 - \frac{e^\epsilon}{e^\epsilon + 1}$$

Given this impossibility result, we consider a further relaxation of differential privacy, which we call *marginal differential privacy*. In contrast to joint differential privacy, marginal differential privacy requires that, simultaneously for every pair of agents $i \neq j$, the *marginal* distribution on agent j ’s allocation be differentially private in agent i ’s data. This corresponds to privacy from an adversary who has the ability

only to look at a single other agent’s allocation (equivalently – privacy from the other agents, assuming they do not collude). Our main result is a marginally-differentially private algorithm that simultaneously guarantees individually rational and approximately Pareto optimal allocations, showing a separation between marginal and joint differential privacy for the exchange problem:

THEOREM. *(Informal) There exists an ϵ -marginally differentially private algorithm that solves the exchange problem with n agents and $k = |\mathcal{G}|$ types of goods by producing an allocation which is individually rational and, with high probability, α -approximately Pareto optimal for*

$$\alpha = O\left(\frac{\text{poly}(k)}{\epsilon n}\right)$$

Note that the approximation to Pareto optimality depends polynomially on the number of *types* of goods in the market, but decreases linearly in the number of participants in the market n . Hence, fixing k , and letting the market size n grow, this mechanism is asymptotically Pareto optimal.

It is natural to ask whether this bound can be improved so that the dependence on k is only $\alpha = O\left(\frac{\text{polylog}(k)}{\epsilon n}\right)$, which is the dependence on the number of distinct types of goods achieved in the approximation to optimality in [Hsu et al. 2014a] (again, under *joint* differential privacy, for a two-sided market). We show that this is not the case.

THEOREM. *(Informal) For every ϵ -marginally differentially private algorithm that on every instance of the exchange problem with n agents and $k = |\mathcal{G}|$ types of goods, produces an individually rational allocation that with high probability is α -approximately Pareto optimal, we have:*

$$\alpha = \Omega\left(\frac{k}{n} \left(1 - \frac{e^\epsilon}{e^\epsilon + 1}\right)\right)$$

We also consider exchange markets in which every agent brings exactly one good to the market, but there are also a small number of *extra* goods available to distribute that are not brought to the market by any of the agents. In the kidney exchange setting, these can represent altruistic donors or non-living donors who provide kidneys for transplantation, but do not need to receive a kidney in return. The existence of extra goods intuitively makes the problem easier, and we show that this is indeed the case. In this setting we use techniques from [Hsu et al. 2014a,b] to get an individually rational, asymptotically exactly Pareto optimal allocation subject to *joint differential privacy* (i.e. circumventing our impossibility result above), under the condition that the number of extra copies of each type of good is a (slowly) growing function of the number of agents.

1.2. Related Work

Differential privacy, introduced by Dwork et al. [2006] has become a standard “privacy solution concept” over the last decade, and has spawned a vast literature too large to summarize. We here discuss only the most closely related work.

Although the majority of the differential privacy literature has considered numeric valued and continuous optimization problems, a small early line of work including Nissim et al. [2007] and Gupta et al. [2010] study combinatorial optimization problems. The dearth of work in this area in large part stems from the fact that many optimization problems cannot be nontrivially solved under the constraint of differential privacy, which requires that the entire output be insensitive to any input. This problem was first observed by Gupta et al. [2010] in the context of set cover and vertex cover, who

also noted that if the notion of a solution is slightly relaxed to include private “instructions” which can be given to the agents, allowing them (together with their own private data) to reconstruct a solution, then more is possible. Similar ideas are also present in McSherry and Mironov [2009], in the context of recommendation systems.

Joint differential privacy, which can be viewed as a generalization of the “instructions” based solution of Gupta et al. [2010], was formalized by Kearns et al. [2014], who showed that, although correlated equilibria in large games could not be computed to any nontrivial accuracy under differential privacy, they can be computed quite accurately under joint differential privacy. A similar result was shown by Rogers and Roth [2014] for *Nash* equilibria in congestion games. Hsu et al. [2014a] subsequently studied a two-sided allocation problem, in which buyers with private valuation functions over bundles of goods must be allocated goods to maximize social welfare (note that here buyers are allocated goods, but do not provide them, unlike the problem we study in this work). They also show that, although the allocation problem cannot be solved to nontrivial accuracy under differential privacy, it can be solved accurately (when buyers preferences satisfy the *gross substitutes* condition) under joint differential privacy.

In the present paper, we continue the study of private allocation problems, and consider the exchange problem in which n agents both supply and receive the goods to be allocated. This problem was first studied by Shapley and Scarf [1974], who also proposed the Top Trading Cycles algorithm for solving it (attributing this algorithm to David Gale). We show that this problem is strictly harder from a privacy perspective than the two sided allocation problem: it cannot be solved non-trivially even under joint differential privacy, but can be solved under a weaker notion which we introduce, that of marginal differential privacy. Our solution involves a privacy-preserving modification of the top trading cycles algorithm. To the best of our knowledge, we are the first to give *marginal* differential privacy a name and to demonstrate a separation from joint differential privacy. The solution concept has, however, appeared in other works – for example in He and Mu [2014], in the context of privacy preserving and incentive compatible recommendation systems.

2. MODEL

We study the setting of trading within an exchange market where there are k types of goods. We denote this set of types of goods as \mathcal{G} . The set of agents will be denoted as N where $|N| = n$. Each $i \in N$ has one copy of some type of good $g_i \in \mathcal{G}$, and some strict linear preference \succ_i over all good types in \mathcal{G} . Let $N_j = \{i : g_i = j\}$ and $n_j = |N_j|$ denote the set and number of agents, respectively, who bring good $j \in \mathcal{G}$ to the market. Since each agent brings exactly one good to the market, we have $\sum_{j \in \mathcal{G}} n_j = n$. An instance of an exchange market is given as $\mathbf{x} = (x_i)_{i \in N}$ where each $x_i = (g_i, \succ_i)$. Our goal will be to find beneficial trades amongst the agents in the market.

Definition 2.1 (Allocation). An *allocation* is a mapping $\pi : N \rightarrow \mathcal{G}$ where we have $|\{i \in N : \pi(i) = j\}| = n_j$ for each $j \in \mathcal{G}$.

We say an algorithm is *individually rational* if each agent is allocated a type of good she weakly prefers to her initial endowment. Formally:

Definition 2.2 (IR). An allocation π is Individually Rational (IR) if $\pi(i) \succeq_i g_i \forall i \in N$.

IR alone does not ensure that high-quality solutions are found: in particular, one IR allocation rule is to make no trades, i.e. $\pi(i) = g_i$ for all $i \in [n]$. It does, however, ensure that no agent is worse off for having participated in the mechanism. Pareto optimality (PO), on the other hand, gives a way to describe the inherent quality of an allocation. An allocation is PO if it cannot be changed to improve some agent i 's utility

without harming the utility of some other agent j . Under privacy constraints, it will be impossible to obtain exact PO, and so we instead ask for an approximate version.

Definition 2.3 (α -PO). An allocation π is α -approximately Pareto optimal (or just α -PO) if for any other allocation π' , if there exists a set $S \subset N$ with $|S| > \alpha n$ such that $\pi'(i) \succ_i \pi(i), \forall i \in S$, then there must be some $j \in N \setminus S$ such that $\pi(j) \succ_j \pi'(j)$. In other words, an allocation is α -PO if strictly improving the allocation for more than an α -fraction of agents necessarily requires strictly harming the allocation of at least 1 agent.

We say that an algorithm guarantees α -PO if, on every exchange problem instance, it outputs an α -PO allocation. If $\alpha = \alpha(n)$ is a function of the number of agents n , we say that an algorithm is asymptotically Pareto optimal if it guarantees $\alpha(n)$ -PO, and $\alpha(n) = o(1)$.

We wish to compute such allocations while guaranteeing a formal notion of privacy to each of the participating agents. The notions of privacy we consider will all be relaxations of *differential privacy*, which has become a standard privacy “solution concept”. We borrow standard notation from game theory: given a vector $\mathbf{x} \in \mathcal{X}^n$, we write $\mathbf{x}_{-i} \in \mathcal{X}^{n-1}$ to denote the vector \mathbf{x} with the i th coordinate removed, and given $x'_i \in \mathcal{X}$ we write $(\mathbf{x}_{-i}, x'_i) \in \mathcal{X}^n$ to denote the vector \mathbf{x} with its i th coordinate replaced by x'_i .

Definition 2.4 (Differential Privacy, [Dwork et al. 2006]). A mechanism $M : \mathcal{X}^n \rightarrow R$ satisfies (ϵ, δ) -differential privacy if for every $i \in [n]$, for any two types $x_i, x'_i \in \mathcal{X}$, any tuple of types $\mathbf{x}_{-i} \in \mathcal{X}^{n-1}$, and any $B \subseteq R$, we have

$$\mathbb{P}(M(\mathbf{x}_{-i}, x_i) \in B) \leq e^\epsilon \mathbb{P}(M(\mathbf{x}_{-i}, x'_i) \in B) + \delta$$

Here R denotes an arbitrary range of the mechanism. The definition of differential privacy assumes that the *input* to the mechanism \mathbf{x} is explicitly partitioned amongst n agents. In the problem we consider, the sensitive data is the exchange market \mathbf{x} , so we simultaneously want to preserve the privacy of each agent, in the form of his initial endowment as well as his preferences over other goods. Note that the input \mathbf{x} is partitioned and the range of the mechanism is also naturally partitioned between n agents (the output of the mechanism can be viewed as n messages, one to each agent i , telling them the type of good they are receiving, $\pi(i)$). In such cases, we can consider relaxations of differential privacy informally parameterized by the maximum size of collusion that we are concerned about. Joint differential privacy, defined by Kearns et al. [2014], asks that the mechanism simultaneously protect the privacy of every agent i from arbitrary collusions of up to $n - 1$ agents $j \neq i$ (who can share their own allocations, but cannot see the allocation of agent i):

Definition 2.5 (Joint Differential Privacy, [Kearns et al. 2014]). A mechanism $M : \mathcal{X}^n \rightarrow O^n$ satisfies (ϵ, δ) -joint differential privacy if for any agent $i \in [n]$, any two types $x_i, x'_i \in \mathcal{X}$, any tuple of types $\mathbf{x}_{-i} \in \mathcal{X}^{n-1}$, and any $B_{-i} \subseteq O^{n-1}$, we have

$$\mathbb{P}(M(\mathbf{x}_{-i}, x_i)_{-i} \in B_{-i}) \leq e^\epsilon \mathbb{P}(M(\mathbf{x}_{-i}, x'_i)_{-i} \in B_{-i}) + \delta$$

As we will show, it is not possible to find IR and asymptotically Pareto optimal allocations under joint differential privacy, and so we introduce a further relaxation which we call marginal differential privacy. Informally, marginal differential privacy requires that the mechanism simultaneously protect the privacy of every agent i from every other agent $j \neq i$, assuming that they do not collude (i.e. it requires the differential privacy condition only on the *marginal* distribution of allocations to other agents, not on the joint distribution).

Definition 2.6 (Marginal Differential Privacy). A mechanism $M : \mathcal{X}^n \rightarrow O^n$ satisfies (ϵ, δ) -marginal differential privacy if $\forall i \neq j, \forall \mathbf{x}_{-i} \in \mathcal{X}^{n-1}, \forall x_i, x'_i \in \mathcal{X}$, and $\forall B \subset O$, we have

$$\mathbb{P}(M(\mathbf{x}_{-i}, x_i)_j \in B) \leq e^\epsilon \mathbb{P}(M(\mathbf{x}_{-i}, x'_i)_j \in B) + \delta$$

3. LOWER BOUNDS

In this section, we show lower bounds on how well the exchange market problem can be solved subject to privacy constraints. We first show that under the constraint of joint-differential privacy, there does not exist any IR and asymptotically Pareto optimal mechanism. This motivates our relaxation of studying exchange problems subject to marginal differential privacy. We then show that under marginal differential privacy, any mechanism producing IR and α -PO allocations must have $\alpha = \Omega(k/n)$, where k is the number of distinct types of goods (that is, a linear dependence on k is necessary). We complement these impossibility results in Section 4, where we show that under marginal differential privacy, it is indeed possible to achieve both IR and α -PO simultaneously, if the number of good types k satisfies $k = o(n^{2/9})$.

Our impossibility result for joint differential privacy is based on a reduction to the following well-known claim, which we prove in Appendix A.

CLAIM 1. *There is no (ϵ, δ) -differentially private mechanism $M : \{0, 1\} \rightarrow \{0, 1\}$ such that $\mathbb{P}(M(b) = b) > \frac{e^\epsilon + \delta}{e^\epsilon + 1}$ for both $b = 0, 1$.*

In the setting of exchange markets, we let $\mathcal{X} = \mathcal{G} \times \mathcal{T}$ where \mathcal{G} is the set of good types and \mathcal{T} is the set of linear orderings over \mathcal{G} for a single agent. Now, we show that it will not be possible to guarantee privacy, IR, and $o(1)$ -PO (with constant ϵ).

THEOREM 3.1. *For any $\epsilon, \delta, \beta > 0$, if we have an (ϵ, δ) -joint differentially private mechanism $M_J : \mathcal{X}^n \rightarrow \mathcal{G}^n$ that guarantees an α -PO allocation with probability at least $1 - \beta$ and always gives an IR allocation then*

$$\alpha \geq 1 - \frac{e^\epsilon + \delta}{(1 - \beta)(e^\epsilon + 1)}$$

PROOF SKETCH. The formal proof is provided in Appendix A; here we sketch the basic intuition. We consider an exchange market with two types of goods, 0 and 1, and n agents endowed with each type of good. An adversary who has control over n agents with good 1 and $n - 1$ agents with good 0 can determine, by observing the trades of the agents she controls, whether the remaining agent with good 0 traded. When she does, the adversary can determine (reasoning based on the fact that the mechanism is IR) that she strictly preferred good 1 to good 0. We can then use Claim 1 to upper-bound the probability that she is allowed to trade in this event. The mechanism does not know to distinguish between the identified agents and the $2n - 1$ agents controlled by the adversary, so reasoning about the case in which this identified agent is selected uniformly at random gives the desired bound. \square

We also show in Appendix A that even under marginal differential privacy, the approximation to Pareto optimality must have at least a linear dependence on k .

THEOREM 3.2. *Any (ϵ, δ) -marginally differentially private mechanism $M_S : \mathcal{X}^n \rightarrow \mathcal{G}^n$ that guarantees an α -PO allocation with probability at least $1 - \beta$ and always satisfies IR must have:*

$$\alpha \geq \frac{k(1 - \beta)}{n} \left(1 - \frac{e^\epsilon + \delta}{e^\epsilon + 1} \right)$$

PROOF SKETCH. Consider an instance with k types of goods and we write $\mathcal{G} = [k]$. Each agent $i \in \{1, \dots, k-1\}$ is endowed with good $g_i = i$ and her most preferred good is $g_i + 1$, followed by g_i . Then we let agents k through n have good type k , all of which prefer their endowment the most except agent k who may prefer a different type of good. Any IR allocation must give agent $i \in \{1, \dots, k-1\}$ either good $g_i + 1$ or g_i and agents $\{k+1, \dots, n\}$ good type k .

Agent k , endowed with good $g_k = k$, is the subject of attack. If she prefers good g_1 , then k goods can be cleared, giving each person her favorite good (agent i getting good $g_i + 1$, for each $i \in [1, \dots, k-1]$ and k getting g_1). If not, no trades can be made. Thus, any agent $i \in [1, \dots, k-1]$ knows with certainty that k prefers g_1 if she receives good $g_i + 1$. With this construction, we again reduce to Claim 1 to lower bound the probability that the mechanism is allowed to clear the k -cycle, even when it is possible. \square

The proof of this lower bound also implies that asymptotic Pareto optimality requires there to be many copies of at least one type of good, with a proof given in Appendix A.

COROLLARY 3.3. *Any (ϵ, δ) -marginally differentially private mechanism that gives an IR allocation with certainty and an α -PO allocation with probability $1 - \beta$ requires there to be at least one type of good with c copies where*

$$c \geq \frac{(1 - \delta)(1 - \beta)}{\alpha(e^\epsilon + 1)}$$

4. PRIVATE TOP TRADING CYCLES

In this section we describe an algorithm which computes IR and α -PO exchanges, subject to marginal differential privacy. The algorithm PTTC (given formally in Algorithm 3 in Appendix B) takes an exchange market $\mathbf{x} = (g_i, \succ_i)_{i=1}^n$ as input and allocates a good to each agent. PTTC begins by considering the complete directed graph $G = (V, A)$ with each node $u \in V$ corresponding to a type of good in \mathcal{G} . For each arc $e = (u, v) \in A$ we define $P_{(u,v)}$ to be the set of people who have good u and consider v their favorite remaining good (and thus want to trade along arc $e = (u, v)$), i.e. $P_{(u,v)} = \{i \in N : g_i = u \ \& \ v \succeq_i g \ \forall g \in V\}$. We then define the arc weight w_e for arc $e \in A$ as $w_e = |P_e|$.

Let N_u be the set of people that are endowed with good u , so $N = \cup_{u \in V} N_u$. The node weights n_u for $u \in V$ are then $n_u = |N_u|$. We first give a high level presentation of our algorithm PTTC that works as follows:

- (1) We compute the arc-weights, and then add noise to them (to preserve privacy).
- (2) We “clear” cycles that have sufficient weight on each arc. Note this determines which goods will trade, but not yet *who* specifically will trade. Because of the added noise, this step satisfies differential privacy.
- (3) We actually execute the trade. This involves selecting exactly *who* will trade along the cycle selected above. We select the individuals at random from among the set that want to trade. Because we need the number of agents trading on each arc in the cycle to be identical, this step is only marginally differentially private.
- (4) We remove any goods that have few remaining copies (and allocate any remaining agent possessing that good his own good), and then return to step (1).

Now, we describe the algorithm in more formality, along with defining the necessary parameters. Note that we will use an ordering of the n agents in our algorithm, say by their index. PTTC works as follows:

- (1) (Arc Weight Noise) With failure probability $\beta > 0$ and privacy parameters $\epsilon, \delta_1, \delta_2 > 0$, we add Laplace noise Z_e with parameter $1/\epsilon'$ to each arc weight w_e where

$$\epsilon' = \frac{\epsilon \log(k^3/\beta)}{2\sqrt{8} \left(\log(k^3/\beta) \sqrt{k \log(1/\delta_1)} + k \sqrt{k \log(1/\delta_2)} \right)}. \quad (1)$$

We let E denote a high probability upper bound (see Claim 2) on all realizations of the noise Z_e , which is

$$E = \frac{\log(k^3/\beta)}{\epsilon'}. \quad (2)$$

We then set the noisy arc weights to be $\hat{w}_e = \max\{w_e + Z_e - 2E, 0\}$.

- (2) (Clear Cycle) If there exists a cycle C in the graph where all the arcs along it have weight $\lfloor \hat{w}_e \rfloor > 0$, then set $\hat{W} = \min_{e \in C} \{\lfloor \hat{w}_e \rfloor\}$ as the weight of C . Otherwise there is no cycle and we skip to step (4).
- (3) (Trade) With the cycle C of weight \hat{W} given in step (2), for each $e \in C$, we select \hat{W} agents at random from P_e to trade along e and denote this set of *satisfied* agents S_e . The way we choose the satisfied agents is via the R-SELECT(\hat{W}, P_e) procedure given in Algorithm 1 which ensures that any fixed agent at arc $e \in C$ is selected with probability \hat{W}/w_e . We then update, for every $e = (u, v) \in C$:

$$P_e \leftarrow P_e \setminus S_e \quad w_e \leftarrow w_e - \hat{W} \quad \hat{w}_e \leftarrow \hat{w}_e - \hat{W}$$

(i.e. remove the agents served and update weights) and for each $i \in S_e$, set $\pi(i) = v$ where $g_i = u$ and $e = (u, v)$ (they get their requested good). Return to step (2).

- (4) (No Cycle - Deletion) If there is no cycle with positive weight, then there must be a node v such that $\hat{n}_v = \sum_{e: e=(v,z), z \in V} \hat{w}_e < k$. We then update $V \leftarrow V \setminus \{v\}$. The arcs that were incoming $IN_v = \{e \in A : e = (u, v) \text{ for some } u \in V\}$ and outgoing from v , $OUT_v = \{e \in A : e = (v, u) \text{ for some } u \in V\}$ are also deleted: $A \leftarrow A \setminus \{OUT_v \cup IN_v\}$. We then move all agents in P_e for each $e \in IN_v$ to the arc that corresponds to trades between g_i and their next most preferred good among the remaining vertices. For every $e = (u, w) \in A$ we define these added arcs as

$$B_{(u,w)} = \{i \in P_{(u,v)} : w \succeq_i g \quad \forall g \in V\}$$

(where $v \notin V$) and update the quantities $w_e \leftarrow w_e + |B_e|$ and $P_e \leftarrow P_e \cup B_e$.

We then assign v to all agents whose endowment was v and have not been assigned a good (for all i such that $i \in P_e$ where $e \in OUT_v$, set $\pi(i) = v$). If any nodes remain, return to Step (1).

- (5) (Clean-up for IR) If we ever assign too many goods or select more agents to trade than actually want to trade, i.e. some \hat{w}_e does not satisfy (3), then undo every assignment of the goods and give everyone the good they started with.

We state the formal version of our main theorem below, which the rest of the paper is dedicated to proving.

THEOREM 4.1. *For $\epsilon, \delta_1, \delta_2, \beta > 0$, our algorithm PTTC is $(\epsilon, \delta_1 + \delta_2 + \beta)$ -marginally differentially private, IR with certainty, and with probability $1 - \beta$ outputs an allocation which is α -PO for*

$$\alpha = \mathcal{O} \left(\frac{k^3}{\epsilon n} \cdot \left(\sqrt{k \log(1/\delta_1)} \log(k^3/\beta) + k \sqrt{k \log(1/\delta_2)} \right) \right).$$

4.1. Pareto Optimality of PTTC

We now prove several lemmas about PTTC, from which Theorem 4.1 follows. We first focus on proving that PTTC is IR and asymptotically Pareto optimal.

ALGORITHM 1: Randomly Select Agents

R-SELECT (W, P);**Input:** Number to select W and ordered set of agents P .**Output:** W randomly selected agents from P Set $M \leftarrow |P|$ Relabel individuals in order of P as $1, 2, \dots, |P|$ For some large integer $Z \gg n$ \ \ For example, set Z to be the population of Earth.Select $\sigma \in [Z]$ uniformly at randomSet $S \leftarrow \{\sigma + 1, \sigma + 2, \dots, \sigma + W \bmod |P|\}$ Map back all labels in S to the original labels they had in P .**return** S .

LEMMA 4.2. *The allocation that PTTC produces is IR.*

PROOF. Note that in the event that the algorithm fails, then all agents are allocated the good they are endowed with, which is trivially IR. Now suppose that the algorithm runs to completion and at some cycle, agent i was selected to trade. In order for i to be selected, that means that his good is still available to trade. Each round, i is on the edge that points to his most preferred good that is available. With his good g_i still available, i will never be on an arc that points to a less preferred good than g_i and hence can never be selected to trade a less preferred good than g_i . \square

We next bound the number of rounds and cycles that are cleared in PTTC.

LEMMA 4.3. *For a fixed round t , there can be at most k^2 cycles selected, i.e. $\tau \leq k^2$ for PTTC. Further, the number of rounds is at most k .*

PROOF. At round t , PTTC finds the τ th cycle with noisy weight ≥ 1 . Then there is some arc e along the cycle that gets depleted, i.e. $\hat{w}_e(t, \tau) \leftarrow \hat{w}_e(t, \tau - 1) - \hat{W}(t, \tau) < 1$ and hence there will never be another cycle in the same round t that uses this arc e . Thus, each cycle that is cleared in a round t can be charged to a unique arc in the graph. There are at most k^2 arcs in any round. Hence, we can clear at most k^2 cycles before there are no arcs (and thus cycles) left at round t . Lastly, each round has some node that is depleted which causes some agents to have to select another good to trade with. Thus, after k rounds, no more nodes will exist. \square

The next claim, proved in Appendix C, allows us to assume that the error from the noise we add in PTTC is small with high probability.

CLAIM 2. *(Low Error) With probability $1 - \beta$, the noise terms of PTTC satisfy*

$$\max_{t \in [k]} \max_{e \in A} |Z_e^t| \leq E = \frac{\log(k^3/\beta)}{\epsilon'}$$

Hence, with probability $1 - \beta$,

$$E \leq w_e - \hat{w}_e \leq 3E. \tag{3}$$

We now want to bound the number of agents that may be left at a node when it has no outgoing arcs with sufficient weight.

LEMMA 4.4. *Assuming the condition in (3), if PTTC cannot find a cycle with noisy weight at least 1, then there exists a node v such that $n_v < D$ for $D = (3E + 1)k$.*

PROOF. If there is no cycle of noisy weight at least 1, then there must be some node v such that every outgoing arc has $\hat{w}_e < 1$. From our Low Error Claim, we then know

that the exact arc weight $w_e < 3E + 1$ from (3) for every outgoing arc of v . Hence we can count how many agents are on the node, $n_v = \sum_{e: e=(v,z), z \in V} w_e < k(3E + 1) = D$. \square

It remains to show that the resulting allocation is asymptotically Pareto optimal.

THEOREM 4.5. *For $\epsilon, \delta_1, \delta_2, \beta > 0$ we have that PTTC outputs an α -PO allocation for*

$$\alpha = \tilde{O}\left(\frac{k^{9/2}}{\epsilon n}\right)$$

with probability $1 - \beta$. Hence, PTTC is asymptotically Pareto optimal.

PROOF. We refer the reader to the formal description of our algorithm (Algorithm 3) in the appendix. We write good type $g = 1$ to be the first good eliminated in PTTC, good $g = 2$ the second, and so on, where the good types $g \in \mathcal{G} = [k]$. We will compare the allocation π from PTTC with any other allocation π' that Pareto dominates π , i.e. $\pi'(i) \succeq_i \pi(i)$ and $\exists j$ such that $\pi'(j) \succ \pi(j)$. We will count the number of agents who could possibly prefer π' . It suffices to bound the quantity: $\Delta = |\{i \in N : \pi'(i) \succ_i \pi(i)\}|$.

At each round, there are a small number of goods that PTTC ignores and just returns them to their owner at the end. Removing these goods from the exchange prevents feasible exchanges between other agents and the removed goods. We seek to bound the total number of agents that may get a better trade in π' but the goods they were allocated were not part of the exchange when PTTC allocated the good they got in π .

Our algorithm deletes a node only at the end of a round, when the noisy supply of at least one good falls below 1. Thus, by Claim 2 and Lemma 4.4, there are at most $D = k(3E + 1)$ copies of goods, with probability at least $1 - \beta$. We condition on this bounded error for the remainder of the proof.

Recall that $S_e(t, \tau)$ is the set of people that traded (or were satisfied) that were along arc e at round t when the τ th cycle $C(t, \tau)$ was selected in PTTC. We define the set $S(t)$ to be all the people that traded at round t : $S(t) = \cup_{\tau \in [k^2]} \cup_{e \in C(t, \tau)} S_e(t, \tau)$.

Some agents are not cleared at any round of PTTC, and these agents receive the good they were endowed with. We refer to those people that were never selected as $S(n) = N \setminus \{\cup_{t=1}^k S(t)\}$. We then have a partition of $N = \cup_{t=1}^k S(t) \cup S(n)$.

We now denote the quantity $\Delta_r = |\{i \in S(r) : \pi'(i) \succ_i \pi(i)\}|$ for $r = 1, \dots, k, n$: note that $\sum_{r \in \{1, \dots, k, n\}} \Delta_r = \Delta$. Further, we would like to refer to the number of goods of type g that were allocated to agents in $S(r)$ in π' but not in π , which we define as $\Delta_r(g)$. Note that for $g \geq r$ we know that $\Delta_r(g) = 0$ because agents cleared by PTTC receive their favorite good among the one's remaining at the round in which they are cleared, and all goods $g \geq r$ are available at round r . More formally we have:

$$\Delta_r(g) = |\{i \in S(r) : \pi'(i) = g \succ_i \pi(i)\}| \text{ for } r = 1, \dots, k, \quad g < r.$$

Thus we can write: $\Delta_r = \sum_{g=1}^{r-1} \Delta_r(g)$. Note that at the first round $r = 1$ that $\Delta_1 = 0$ because everyone that was selected gets their favorite type of good. Let us also define $\Delta_r(g, h)$ to be the number of goods of type g allocated to agents in $S(r)$ in π' who received good h in π :

$$\Delta_r(g, h) = |\{i \in S(r) : \pi'(i) = g \succ_i h = \pi(i)\}| \quad \text{for } r = 1, \dots, k, \quad h \geq r \quad g < r.$$

$$\Delta_r(g) = \sum_{h=r}^k \Delta_r(g, h) \quad g < r$$

We denote the number of initial goods of type g as $n(g)$ for $g \in [k]$: i.e. $n(g) = |i : g_i = g|$. We define $n_t(g)$ as the number of goods of type g that are not allocated to members of

$\cup_{r=1}^t S(r)$ in π' (our notation uses the round as a subscript and the good as an argument in parentheses), i.e.

$$n_t(g) = n(g) - \sum_{r=1}^t |\{i \in S(r) : \pi'(i) = g\}|.$$

We will now bound the quantities $n_t(g)$. Note that for the first round, we have

$$n_1(1) = n(1) - |\{i \in S(1) : \pi'(i) = 1\}| \leq D$$

because each person in $S(1)$ got his favorite good in PTTC, and since π' Pareto dominates π , π' must have made the same allocation as π for $S(1)$, except for the at most D agents PTTC never selected that had a good of type 1. The agents that get selected later can Pareto improve because they could have selected good type 1, but PTTC has deleted that good for future rounds. These lost copies of good 1 can potentially be used in allocation π' to improve the outcome of agents selected at future rounds by PTTC. We will account for these improvements by keeping track of the $n_r(1)$ copies of good 1 remaining at each round, where

$$n_r(1) = n_{r-1}(1) - \Delta_r(1) \quad \text{for } r = 2, \dots, k.$$

We then continue in this fashion with good type 2 in order to bound $n_2(2)$. We know that π' allocates $\Delta_2(1)$ goods of type 1 to agents in $S(2)$. Because π' Pareto dominates π , it must be that π' makes the same allocations as π among agents in $S(2)$, except for the agents that π' matches to good type 1 that got good type 2 in π (these are the only agents who are possibly not getting their favorite good among those “remaining” in π'). We then bound the number of goods of type 2 that have not been distributed to people in $S(1)$ or $S(2)$ and then keep track of the number of these goods that remain to give to the people that are selected in future rounds but get good type 2 in π' ,

$$n_2(2) \leq D + \Delta_2(1, 2)$$

$$n_r(2) = n_{r-1}(2) - \Delta_r(2) \quad r > 2.$$

We now consider $n_3(3)$, the number of goods of type 3 that π' has not allocated to members in $S(1)$, $S(2)$, and $S(3)$. This is the same as the number of goods of type 3 that π will never give to selected people (at most D) in addition to the goods of type 3 that π gave to people in $S(2)$ and $S(3)$ that π' gave a different good to, i.e. the $\Delta_2(1, 3)$ people that got good 3 at round 2 in π , but π' gave them good 1, along with the $\Delta_3(1, 3)$ and $\Delta_3(2, 3)$ people that got good 3 at round 3 in π , but π' gave them good 1 and 2 respectively. This implies both

$$n_3(3) \leq D + \Delta_2(1, 3) + \Delta_3(1, 3) + \Delta_3(2, 3)$$

$$n_r(3) = n_{r-1}(3) - \Delta_r(3) \quad r > 3.$$

We then generalize the relation for $r \geq 3$:

$$n_r(r) \leq D + \sum_{\ell=2}^r \sum_{g=1}^{\ell-1} \Delta_\ell(g, r) \tag{4}$$

$$n_t(r) = n_{t-1}(r) - \Delta_t(r) \quad t > r. \tag{5}$$

Because the number of goods remaining at each round must be nonnegative, we have

$$\Delta_t(r) \leq n_{t-1}(r). \tag{6}$$

Recall that $\Delta_1 = 0$. We also have: $\Delta_2 = \Delta_2(1) \leq n_1(1) \leq D$. For round $3 \leq t \leq k$, we use (4), (5), and (6) to get:

$$\begin{aligned}
\Delta_t &= \sum_{g=1}^{t-1} \Delta_t(g) \stackrel{(6)}{\leq} \sum_{g=1}^{t-1} n_{t-1}(g) \stackrel{(5)}{=} \sum_{r=1}^{t-1} n_r(r) - \sum_{g=1}^{t-2} \sum_{r=g+1}^{t-1} \Delta_r(g) \\
&\stackrel{(4)}{\leq} (t-1)D + \sum_{r=2}^{t-1} \sum_{\ell=2}^r \sum_{g=1}^{\ell-1} \Delta_\ell(g, r) - \sum_{g=1}^{t-2} \sum_{r=g+1}^{t-1} \Delta_r(g) \\
&= (t-1)D + \underbrace{\sum_{\ell=2}^{t-1} \sum_{g=1}^{\ell-1} \sum_{r=\ell}^{t-1} \Delta_\ell(g, r)}_{\leq \Delta_\ell(g)} - \sum_{r=2}^{t-1} \sum_{g=1}^{r-1} \Delta_r(g) \leq (t-1)D.
\end{aligned}$$

We next bound Δ_n . All agents who were never cleared might be able to get a better good in π' , which can not be more than the total number of goods that π' did not allocate to any of the selected people. Thus we have

$$\begin{aligned}
\Delta_n &\leq \sum_{g=1}^k n_k(g) = \sum_{r=1}^k n_r(r) - \sum_{g=1}^{k-1} \sum_{r=g+1}^k \Delta_r(g) \\
&\leq kD + \underbrace{\sum_{\ell=2}^k \sum_{g=1}^{\ell-1} \sum_{r=\ell}^k \Delta_\ell(g, r)}_{=\Delta_\ell(g)} - \sum_{r=2}^k \sum_{g=1}^{r-1} \Delta_r(g) = kD.
\end{aligned}$$

We then sum over Δ_t for every round t to get Δ :

$$\begin{aligned}
\Delta_1 + \Delta_2 + \Delta_3 + \cdots + \Delta_k + \Delta_n &\leq \sum_{j=1}^k jD = \mathcal{O}(k^2 D) \\
&= \mathcal{O}\left(\frac{k^3 \left(k\sqrt{k \log(1/\delta_2)} + \log(k^3/\beta)\sqrt{k \log(1/\delta_1)}\right)}{\epsilon}\right)
\end{aligned}$$

where the last equality followed from $D \leq k(3E + 1)$, from ϵ' in (1), and E in (2). \square

4.2. Privacy Analysis of PTTC

In order to prove that PTTC is marginally differentially private, we first present some known results on differential privacy and prove some lemmas that will be useful. We first state the definition of sensitivity of a function, which will help when we define the Laplace Mechanism [Dwork et al. 2006] which is a subroutine of PTTC.

Definition 4.6 (Sensitivity [Dwork et al. 2006]). A function $\phi : \mathcal{X}^n \rightarrow \mathbb{R}^m$ has *sensitivity* ν_ϕ defined as

$$\nu_\phi = \max_{i \in [n]} \max_{\mathbf{x}_{-i} \in \mathcal{X}^{n-1}} \max_{x_i \neq x'_i} \|\phi(\mathbf{x}_{-i}, x_i) - \phi(\mathbf{x}_{-i}, x'_i)\|_1.$$

The Laplace Mechanism, given in Algorithm 2, answers a numeric query $\phi : T^n \rightarrow \mathbb{R}^m$ by adding noise to each component of ϕ 's output in a way which is differentially private.

We now state the privacy guarantee for the Laplace Mechanism M_L .

ALGORITHM 2: Laplace Mechanism $M_L(\epsilon, \phi)(\mathbf{x})$

Input: Database \mathbf{x}

Output: An approximate value for ϕ

$$\hat{\phi} = \phi(\mathbf{x}) + (Z_1, \dots, Z_m) \quad Z_i \stackrel{i.i.d.}{\sim} \text{Lap}(\nu_\phi/\epsilon)$$

return $\hat{\phi}$

THEOREM 4.7 ([DWORK ET AL. 2006]). $M_L(\epsilon, \phi)$ is ϵ -differentially private for any $\phi : \mathcal{X}^n \rightarrow \mathbb{R}^n$ with bounded sensitivity.

Our algorithm PTTC uses the Laplace Mechanism to modify arc weights at each round.

One of the most useful properties given by differential privacy is its ability to compose: running a collection of private mechanisms is also private, with a loss in the privacy parameters which depends upon how the composition is done. We will need to use two composition theorems. The first shows that the privacy parameters add when we compose two differentially private mechanisms, and the second from [Dwork et al. 2010] gives a better composition guarantee even with many adaptively chosen mechanisms.

THEOREM 4.8. *If mechanism $M_1 : \mathcal{X}^n \rightarrow O$ is (ϵ_1, δ_1) -differentially private, and another mechanism $M_2 : \mathcal{X}^n \times O \rightarrow R$ is (ϵ_2, δ_2) -differentially private in its first component, then $M : \mathcal{X}^n \rightarrow R$ is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ differentially private where $M(\mathbf{x}) = M_2(\mathbf{x}, M_1(\mathbf{x}))$.*

If we were to only consider the previous composition theorem, then the composition of m mechanisms that are (ϵ, δ) -differentially private mechanisms would lead to an $(m\epsilon, m\delta)$ -differentially private mechanism. However, we can improve on the $m\epsilon$ parameter at the cost of making the additive $m\delta$ term larger. The following theorem gives this modified composition guarantee that holds even when the sequence of differentially private mechanisms is chosen adaptively by an adversary, as a function of the output of previous mechanisms.

THEOREM 4.9 (*m-FOLD ADAPTIVE COMPOSITION* [DWORK ET AL. 2010]). *Fix $\delta > 0$. The class of (ϵ', δ') differentially private mechanisms satisfies $(\epsilon, m\delta' + \delta)$ differential privacy under m -fold adaptive composition for*

$$\epsilon' = \frac{\epsilon}{\sqrt{8m \log(1/\delta)}}.$$

We next give a lemma that will be useful in proving marginal differential privacy, given an intermediate step that is differentially private.

LEMMA 4.10. *Let $M^1 : \mathcal{X}^n \rightarrow O$ be (ϵ_1, δ_1) -differentially private and $M_j^2 : \mathcal{X}^{n-1} \times \mathcal{X} \times O \rightarrow R$ for $j = 1, \dots, n$ be (ϵ_2, δ_2) -differentially private in its first argument. Then $M : \mathcal{X}^n \rightarrow R^n$ is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -marginally differentially private where*

$$M(\mathbf{x}) = (M_j^2(\mathbf{x}_{-j}, x_j, M^1(\mathbf{x})))_{j=1}^n.$$

PROOF. To prove marginal differential privacy we need to prove that the component $M(\mathbf{x}_{-j}, x_j)_j$ is differentially private in its first argument, for every j . Fix any index i and j such that $i \neq j$. We then consider a $\mathbf{x}_{-i} \in \mathcal{X}^{n-1}$, $x'_i \neq x_i$, and $S \subset R$, then we use

composition of differentially private mechanisms to get the following:

$$\begin{aligned}
\mathbb{P}(M_j^2[\mathbf{x}_{-j}, x_j, M^1(\mathbf{x})] \in S) &= \int_O \mathbb{P}(M_j^2((\mathbf{x}_{-(i,j)}, x_i), x_j, o) \in S) \mathbb{P}(M^1(\mathbf{x}_{-i}, x_i) = o) do \\
&\leq \int_O \mathbb{P}(M_j^2((\mathbf{x}_{-(i,j)}, x_i), x_j, o) \in S) (e^{\epsilon_1} \mathbb{P}(M^1(\mathbf{x}_{-i}, x'_i) = o) + \delta_1) do \\
&\leq \int_O (e^{\epsilon_2} \mathbb{P}(M_j^2((\mathbf{x}_{-(i,j)}, x'_i), x_j, o) \in S) + \delta_2) e^{\epsilon_1} \mathbb{P}(M^1(\mathbf{x}_{-i}, x'_i) = o) do + \delta_1 \\
&\leq e^{\epsilon_1 + \epsilon_2} \mathbb{P}(M_j^2[(\mathbf{x}_{-(i,j)}, x'_i), x_j, M^1(\mathbf{x}_{-i}, x'_i)] \in S) + \delta_1 + \delta_2
\end{aligned}$$

Where in the last inequality we use the fact that if $e^{\epsilon_1} \mathbb{P}(M^1(\mathbf{x}_{-i}, x'_i) = o) > 1$ then we just replace the bound by 1. \square

We now present two lemmas specific to our setting: one dealing with the privacy of the noisy arc weights computed in PTTC and the other dealing with how agents are selected to trade. We leave both proofs to Appendix C.

LEMMA 4.11. *The mechanism $\mathcal{M} : \mathcal{X}^n \rightarrow \mathbb{R}^{k \times k^2}$ that outputs all the noisy arc weights $(\hat{w}(t, 0))_{t \in [k]}$ that is used in PTTC is (ϵ_1, δ_1) -differentially private for $\delta_1 > 0$ and*

$$\epsilon_1 = 2\epsilon' \cdot \sqrt{8k \log(1/\delta_1)}. \quad (7)$$

PROOF SKETCH. At each round of PTTC, we add Laplace noise with parameter $1/\epsilon'$ to each arc weight. Consider agent i changing her input from x_i to x'_i and fix all the other agents' inputs. We want to bound the sensitivity of the (exact) arc weights at any particular round. If we fix the randomness used in R-SELECT that determines which agents are cleared, and the previous rounds' noisy arc weights, then we can consider several cases in how agent i changing reports can affect the arc counts at the next round. We show that no matter whether agent i is selected earlier or later on a different input, she can affect at most 2 entries in the vector of arc weights at a particular round by at most 1. Hence, the sensitivity of the arc weights at a particular round is at most 2, conditioned on the past rounds' randomness. We then apply advanced composition in Theorem 4.9 over k rounds to get ϵ_1 . \square

For the next lemma, we will condition on knowing the noisy arc weights that are computed throughout PTTC and compute the difference in distributions between agent j being selected at any round when agent $i \neq j$ changes reports from x_i to x'_i .

LEMMA 4.12. *Let \hat{w} be the noisy arc weights that are computed for every round of our algorithm PTTC. We also assume that each entry of \hat{w} satisfies (3). The mechanism $\mathcal{M}_j : \mathcal{X}^{n-1} \times \mathcal{X} \times \mathbb{R}^{k \times k^2} \rightarrow \{0, \mathcal{G}\}$ is (ϵ_2, δ_2) -differentially private in its first argument given \hat{w} , where*

$$\mathcal{M}_j(\mathbf{x}_{-j}, x_j, \hat{w}) = \begin{cases} g & \text{If } \exists t \in [k], \tau \in [k^2] \text{ s.t. } j \in S_e(t, \tau) \text{ for } e = (g_j, g) \in C(t, \tau) \\ 0 & \text{else} \end{cases}.$$

with $\delta_2 > 0$ and

$$\epsilon_2 = \frac{2k \sqrt{8k \log(1/\delta_2)}}{E}. \quad (8)$$

PROOF SKETCH. We consider a round t and cycle C such that agent j is on an arc $e \in C$ and so may or may not be selected at that cycle. We then bound the ratio between the probability that agent j is selected when $i \neq j$ reports x_i to the probability of the same event when i 's input is x'_i . We are given the number of people that are being

selected (this is a function of the noisy arc weights). As we argued in the previous lemma, agent i 's influence on the weight of any arc at any round by at most 1. We do a similar analysis when the event is j not being selected. Knowing this, and the fact that the noisy arc weights satisfy (3), we get that the mechanism that determines whether j is selected or not at a given round, as part of a given cycle is $2/E$ -differentially private in reports \mathbf{x}_{-j} .

We then apply advanced composition (Theorem 4.9) over the at most k^2 different cycles that might need to be cleared on a single round, and over the k different rounds, to get the value for ϵ_2 . \square

We are now ready to prove that PTTC is marginally differentially private.

THEOREM 4.13. *For parameters $\epsilon, \delta_1, \delta_2, \beta > 0$, PTTC : $\mathcal{X}^n \rightarrow \mathcal{G}^n$ is $(\epsilon, \delta_1 + \delta_2 + \beta)$ -marginally differentially private.*

PROOF. We fix agents i and j , where $i \neq j$. Let us define M^1 as \mathcal{M} that outputs all the noisy arc weights for each round of PTTC from Lemma 4.11 and M_j^2 as \mathcal{M}_j the good that j ends up being matched with (or zero if never matched) from Lemma 4.12.

We first condition on the event that all the noisy arc weights computed by M^1 satisfy (3). We then apply Lemma 4.10 to get the composed mechanism $M : \mathcal{X}^n \rightarrow \{0, \mathcal{G}\}^n$ of our mechanisms M^1 and $(M_j^2)_{j=1}^n$ is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2)$ -marginally differentially private. Note that M and PTTC have the same distribution of outcomes where if M outputs 0 to agent j , then we know PTTC will give agent j his own good type g_j .

We have yet to consider the case when the noisy arc weights do not satisfy (3). However, this occurs with probability at most β . We conclude then that M (and thus PTTC) is $(\epsilon_1 + \epsilon_2, \delta_1 + \delta_2 + \beta)$ -marginally differentially private. We then plug in the values for ϵ_1 in (7), ϵ_2 in (8), ϵ' in (1), and E in (2) to get

$$\epsilon_1 + \epsilon_2 = \frac{2\epsilon'\sqrt{8k}}{\log(k^3/\beta)} \cdot \left(\log(k^3/\beta) \sqrt{\log(1/\delta_1)} + k \sqrt{\log(1/\delta_2)} \right) = \epsilon$$

\square

We want δ_1, δ_2 and β to be as small as possible because this causes an additive difference in the probability distributions between neighboring exchange markets. We can then set $\delta_1, \delta_2, \beta = \text{poly}(1/n)$ to still get $\alpha = \tilde{O}\left(\frac{k^{9/2}}{\epsilon n}\right)$.

5. ALLOWING A SMALL SUPPLY OF GOODS TO BE INJECTED

Without privacy constraints, rather than running the top trading cycles algorithm, we could solve the following linear program to obtain an IR and PO exchange:

$$\max_{\mathbf{z}} \quad \sum_{i \in [n]} \sum_{j \in \mathcal{G}} r_{ij} z_{ij} \quad (9)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{G}} z_{ij} = n_j \quad \forall i \in [n] \quad (10)$$

$$\forall j \in \mathcal{G} \text{ s.t. } g_i \succ_i j \quad z_{i,j} = 0 \quad \forall i \in [n] \quad (11)$$

$$\sum_{i \in [n]} z_{ij} = 1 \quad \forall j \in \mathcal{G} \quad (12)$$

$$z_{ij} \in \{0, 1\}. \quad (13)$$

where we define $r_{ij} = k - r + 1$ if $j \succeq_i g_i$ and j is the r th top choice of agent i . If $g_i \succ_i j$ then $w_{ij} = 0$. Note that since this is just a max-weight matching problem, the optimal

solution x^* will be integral, even if we relax the integer constraint. The constraint in (11) ensures that no agent is matched to a good that is preferred less than her endowed good, which enforces IR. Finally, the optimal solution to this linear program is an (exactly) Pareto optimal allocation: any other allocation that gave some agents strictly preferred goods, without decreasing the quality of goods given to other agents would have a strictly improved objective value, contradicting optimality. We thus have the following theorem:

THEOREM 5.1. *A solution to the integer program (IP) in (9) - (13) is an IR and PO allocation.*

We leverage recent work [Hsu et al. 2014a] and [Hsu et al. 2014b] on computing max-weight matchings subject to joint differential privacy. We will use the results of the latter paper to conclude that we can obtain an asymptotically Pareto optimal, IR allocation that is also jointly differentially private under a relaxation of our problem that allows us to inject a small number of extra goods into the system. We present the following theorem in the context of the IP of (9) - (13).

THEOREM 5.2 (HSU ET AL. 2014B). *Let OPT be the optimal objective value of (9). For $\epsilon, \delta, \beta > 0$ there exists an (ϵ, δ) -jointly differentially private algorithm that produces fractional solution \bar{z} such that $\sum_{j \in \mathcal{G}} \bar{z}_{ij} = 1$ and with probability $1 - \beta$*

- We get a solution close to OPT , i.e. $\sum_{i \in [n]} \sum_{j \in \mathcal{G}} r_{ij} \bar{z}_{ij} \geq OPT - \eta$ where

$$\eta = \mathcal{O} \left(\frac{k^2 \log(nk/\beta) \log^{1/2}(n/\delta)}{\epsilon} \right)$$

- The total amount all the constraints in (12) are violated by \bar{z} is small, i.e. we have for each good type $j \in \mathcal{G}$, $n_j - \lambda_j \leq \sum_{i \in [n]} \bar{z}_{ij} \leq n_j + \lambda_j$ where

$$\sum_{j \in \mathcal{G}} \lambda_j = \mathcal{O} \left(\frac{k \log(nk/\beta) \log^{1/2}(n/\delta)}{\epsilon} \right)$$

Note that the solution \bar{z} we obtain is fractional. We then use a randomized rounding technique from Raghavan and Thompson [1987] that has each agent i choose good j with probability \bar{z}_{ij} . By applying a Chernoff bound, the resulting solution after applying randomized rounding \hat{z}_{ij} gives with probability at least $1 - \beta$

- It is the case that $\sum_{i \in [n]} \sum_{j \in \mathcal{G}} r_{ij} \hat{z}_{ij} \geq OPT - \hat{\eta}$ where

$$\hat{\eta} = \mathcal{O} \left(\frac{k^2 \sqrt{n} \log^{3/2}(nk/\beta) \log^{1/2}(n/\delta)}{\epsilon} \right)$$

- The total amount all the constraints in (12) are violated by \hat{z} is small, i.e. $n_j - \hat{\lambda}_j \leq \sum_{i \in [n]} \hat{z}_{ij} \leq n_j + \hat{\lambda}_j$ where

$$\sum_{j \in \mathcal{G}} \hat{\lambda}_j = \mathcal{O} \left(\frac{k \sqrt{n} \log^{3/2}(nk/\beta) \log^{1/2}(n/\delta)}{\epsilon} \right)$$

Since the supply constraints are violated in the above solution, it is infeasible, and cannot be implemented if the market is closed. Moreover, this is inherent – everything we have done here is subject to joint differential privacy, for which we have proven a lower bound. However, if we have a supply of *extra* goods of each type (e.g. non-living

kidney donors), then we can use these extra goods to restore feasibility (Note that it is important that the “extra” goods are not attached to agents who have IR constraints). This allows us to circumvent our lower bound, and leads to the following theorem:

THEOREM 5.3. *There exists an (ϵ, δ) joint differentially private algorithm that allocates goods in an exchange market that is α -PO with probability $1 - \beta$ and always IR which needs at most a total of Λ extra goods to ensure everyone gets a good where*

$$\alpha = \mathcal{O}\left(\frac{k^2 \log^{3/2}(nk/\beta) \log^{1/2}(n/\delta)}{\sqrt{n\epsilon}}\right) \quad \text{and} \quad \Lambda = \mathcal{O}\left(\frac{k\sqrt{n} \log^{3/2}(nk/\beta) \log^{1/2}(n/\delta)}{\epsilon}\right)$$

6. CONCLUSION/OPEN PROBLEMS

In this paper we have continued the study of the accuracy to which *allocation problems* can be solved under parameterized relaxations of *differential privacy*. Generically, these kinds of problems cannot be solved under the standard constraint of differential privacy. Unlike two sided allocation problems which can be solved under *joint-differential privacy*, we show that Pareto optimal exchanges cannot be solved even under this relaxation, but can be solved asymptotically exactly under marginal differential privacy whenever the number of types of goods $k = o(n^{2/9})$. (We note that in many applications, such as kidney exchange, k will be constant).

The two privacy solution concepts we have considered are only two extremes along a spectrum: informally, joint differential privacy protects the privacy of agent i against an adversarial collusion of possibly all of the $n - 1$ other agents in the market, acting against agent i . Similarly, marginal differential privacy protects agent i 's privacy only against a single agent $j \neq i$ in the computation, assuming she does not collude with anyone else. We propose a definition for future work which smoothly interpolates between joint and marginal differential privacy, which we call m -coalition differential privacy. The case of $m = 1$ recovers marginal differential privacy and the case of $m = n - 1$ recovers joint differential privacy: for $1 < m < n - 1$, we get a sequence of privacy definitions smoothly interpolating between the two.

Definition 6.1 (Coalition Differential Privacy). We say that a mechanism $M : \mathcal{X}^n \rightarrow A^n$ is m -coalition (ϵ, δ) -differentially private if for any set $S \subseteq [n]$ with $|S| \leq m$ and for any $\mathbf{x} = (x_1, \dots, x_n)$ and $x'_i \neq x_i$ where $i \notin S$ we have for any $B \subset O^{|S|}$

$$\mathbb{P}(M(\mathbf{x}_{-i}, x_i)_S \in B) \leq e^\epsilon \mathbb{P}(M(\mathbf{x}_{-i}, x'_i)_S \in B) + \delta$$

where $M(\mathbf{x})_S = (M(\mathbf{x})_j)_{j \in S}$

We note that it is not true in general that ϵ -marginal differential privacy implies to m -coalition $(m\epsilon)$ -differential privacy, because the marginal distributions between players may be correlated, and so this study may require new tools and techniques.

It would also be interesting to give a privacy preserving algorithm that is not only individually rational and asymptotically Pareto optimal, but makes truthful reporting a dominant strategy. One difficulty is that in our setting (in which there are multiple copies of identical goods), agents do not have strict preferences over goods, and even the top trading cycles algorithm without privacy is not incentive compatible. However, there are other algorithms such as [Saban and Sethuraman 2013; Alcalde-Unzu and Molis 2009; Jaramillo and Manjunath 2012], that are incentive compatible in exchange markets that allow indifferences, so it may be possible. (It would also be interesting to find a connection between marginal differential privacy and incentive compatibility, like the known connections with differential privacy [McSherry and Talwar 2007] and joint differential privacy [Kearns et al. 2014]).

REFERENCES

- ALCALDE-UNZU, J. AND MOLIS, E. 2009. Exchange of indivisible goods and indifference: the top trading absorbing sets mechanisms. CORE Discussion Papers 2009062, Universit catholique de Louvain, Center for Operations Research and Econometrics (CORE).
- DWORK, C., MCSHERRY, F., NISSIM, K., AND SMITH, A. 2006. Calibrating noise to sensitivity in private data analysis. In *TCC '06*. 265–284.
- DWORK, C., ROTHBLUM, G. N., AND VADHAN, S. P. 2010. Boosting and differential privacy. In *FOCS*. 51–60.
- GUPTA, A., LIGETT, K., MCSHERRY, F., ROTH, A., AND TALWAR, K. 2010. Differentially private combinatorial optimization. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 1106–1125.
- HE, K. AND MU, X. 2014. Differentially private and incentive compatible recommendation system for the adoption of network goods. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 949–966.
- HSU, J., HUANG, Z., ROTH, A., ROUGHGARDEN, T., AND WU, Z. S. 2014a. Private matchings and allocations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. STOC '14. ACM, New York, NY, USA, 21–30.
- HSU, J., HUANG, Z., ROTH, A., AND WU, Z. S. 2014b. Jointly private convex programming. *CoRR abs/1411.0998*.
- JARAMILLO, P. AND MANJUNATH, V. 2012. The difference indifference makes in strategy-proof allocation of objects. *Journal of Economic Theory* 147, 5, 1913–1946.
- KEARNS, M., PAI, M., ROTH, A., AND ULLMAN, J. 2014. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th ACM SIGact Innovations in Theoretical Computer Science (ITCS)*.
- MCSHERRY, F. AND MIRONOV, I. 2009. Differentially private recommender systems: building privacy into the net. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 627–636.
- MCSHERRY, F. AND TALWAR, K. 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.
- NISSIM, K., RASKHODNIKOVA, S., AND SMITH, A. 2007. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. ACM, 75–84.
- RAGHAVAN, P. AND THOMPSON, C. D. 1987. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7, 4, 365–374.
- ROGERS, R. M. AND ROTH, A. 2014. Asymptotically truthful equilibrium selection in large congestion games. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 771–782.
- ROTH, A. E., SÖNMEZ, T., AND UTKU ÜNVER, M. 2005. Pairwise kidney exchange. *Journal of Economic Theory* 125, 2, 151–188.
- SABAN, D. AND SETHURAMAN, J. 2013. House allocation with indifference: A generalization and a unified view. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. ACM, New York, NY, USA, 803–820.
- SHAPLEY, L. AND SCARF, H. 1974. On cores and indivisibility. *Journal of Mathematical Economics* 1, 1, 23–37.

A. FORMAL PROOFS FROM LOWER BOUNDS SECTION 3

PROOF OF CLAIM 1. Let M be (ϵ, δ) -differentially private such that $\mathbb{P}(M(0) = 0) > \frac{e^\epsilon + \delta}{e^\epsilon + 1}$ and $\mathbb{P}(M(1) = 1) > \frac{e^\epsilon + \delta}{e^\epsilon + 1}$. By the definition of ϵ -differential privacy, we have

$$\begin{aligned} \frac{e^\epsilon + \delta}{e^\epsilon + 1} &< \mathbb{P}(M(0) = 0) \leq e^\epsilon \mathbb{P}(M(1) = 0) + \delta \\ &= e^\epsilon(1 - \mathbb{P}(M(1) = 1)) + \delta < e^\epsilon \left(1 - \frac{e^\epsilon + \delta}{e^\epsilon + 1}\right) + \delta, \end{aligned}$$

giving a contradiction. \square

PROOF OF THEOREM 3.1. The proof proceeds by a reduction to Claim 1. Suppose we had such an (ϵ, δ) -joint differentially private mechanism M_J , with $\alpha < 1 - \frac{e^\epsilon + \delta}{(1-\beta)(e^\epsilon + 1)}$. We show that we could use it to construct an ϵ -differentially private mechanism $M : \{0, 1\} \rightarrow \{0, 1\}$ that contradicts Claim 1. We design an exchange market parameterized by the input bit b received by mechanism M (see Figure 1). The market has two types of goods, g_0 and g_1 and $2n$ agents partitioned into two sets, N_0 and N_1 of size n each. The agents $j \in N_1$ are endowed with good g_1 and strictly prefer good g_0 (i.e. all such agents j have preference $g_0 \succ_j g_1$). The agents in N_0 are endowed with good g_0 . We assume $n - 1$ of them strictly prefer good g_1 (i.e. all such agents j have preference $g_1 \succ_j g_0$). A distinguished agent, $i \in N_0$, selected among the n agents in N_0 uniformly at random, has preference determined by bit b : $g_b \succ_i g_{1-b}$. (i.e. the i 'th agent wishes to trade if $b = 1$, but prefers keeping her own good if $b = 0$.) We denote the vector of linear preferences that depends on i 's bit b as $\succ(b)$. We will refer to this exchange market as $\mathbf{x}(b) = (\mathbf{g}, \succ(b)) \in \mathcal{X}^{2n}$ where $\mathbf{g} \in \{0, 1\}^{2n}$. We remark that when $b = 1$, the agents in N_0 are identical.




Fig. 1. Depicting the exchange market considered in the proof of Theorem 3.1.

Let $M_J : \mathcal{X}^{2n} \rightarrow \mathcal{G}^{2n}$ be the (ϵ, δ) -joint differentially private mechanism given in the statement of the theorem. Note first that by the definition of joint differential privacy, the mechanism $M' : \{0, 1\} \rightarrow \mathcal{G}^{n-1}$ defined as $M'(b) = M_J(\mathbf{x}(b))_{-i}$ (which takes as input b and outputs the allocation of all $n - 1$ agents $j \neq i$) is (ϵ, δ) -differentially private.

In this construction, when $b = 0$, the IR constraint requires that $M_J(\mathbf{x}(0))_i = 0$ with probability 1. Note also that from the output of $M'(b)$, we can determine whether agent i engaged in trade: if not, we must have:

$$|\{j \in N_0 \setminus i : M'(b)_j = g_1\}| = |\{j \in N_1 : M'(b)_j = g_0\}|$$

but if so, we must have:

$$|\{j \in N_0 \setminus i : M'(b)_j = g_1\}| = |j \in N_1 : M'(b)_j = g_0| - 1$$

Define $f : \mathcal{G}^{2n-1} \rightarrow \{0, 1\}$ to be the indicator function of the event $|\{j \in N_0 \setminus i : M'(b)_j = g_1\}| \neq |j \in N_1 : M'(b)_j = g_0|$. Define $M : \{0, 1\} \rightarrow \{0, 1\}$ to be $M(b) = f(M'(b))$. Note that M is (ϵ, δ) -differentially private by the post-processing guarantee of differential privacy, and is an indicator variable determining whether agent i has traded in our exchange economy.

If $M(b) = 1$, therefore, it must be that $b = 1$ by the individual rationality guarantee of M_J . (i.e. $\mathbb{P}(M(0) = 1) = 0 \implies \mathbb{P}(M(0) = 0) = 1$). Hence, by Claim 1 we must have $\mathbb{P}(M(1) = 1) \leq \frac{e^\epsilon + \delta}{e^\epsilon + 1}$. We know by hypothesis that M_J finds an α -PO allocation with probability $1 - \beta$. If $b = 1$ then every agent wishes to trade in $\mathbf{x}(b)$ and hence with probability $1 - \beta$, M_J must produce an allocation in which at least $2n(1 - \alpha)$ people trade. Since all agents in N_1 are identical, and i was selected uniformly at random, it must therefore be that agent i engages in trade with probability at least $(1 - \beta)(1 - \alpha)$. Thus, we have

$$\frac{e^\epsilon + \delta}{e^\epsilon + 1} \geq \mathbb{P}(M(1) = 1) \geq (1 - \beta)(1 - \alpha)$$

which gives us the conclusion of the theorem. \square

PROOF OF THEOREM 3.2. Suppose we have an (ϵ, δ) -marginally differentially private mechanism M_S which is IR and α -PO. We will use it to construct some (ϵ, δ) -differentially private mechanism M , which will give a lower bound on α by Claim 1.

Suppose there are k types of goods in the market, $1, \dots, k$. Let g_i represent the type of good with which agent i is initially endowed. Figure 2 shows the favorite and second favorite goods of all n agents². We will refer to this exchange market as $\mathbf{x}(b) = (\mathbf{g}, \succ(b)) \in \mathcal{X}^n$ where $\mathbf{g} \in \{1, \dots, k\}^n$. A single agent k has preferences which are determined by bit b : if $b = 0$, her favorite good is $g_k = k$ her own, and if $b = 1$, her favorite good is good 1. In the case that $b = 1$, agents $1, \dots, k$ form a cycle with their favorite preferences: agent $i = 1, \dots, k - 1$ gets good $i + 1$ and agent k gets good 1, which would give each agent her favorite good. If, on the other hand, $b = 0$, the uniquely IR trade is the $\pi(i) = g_i$, or no trade, since that instance contains no cycles, other than self loops.

Agent	Endowment	Favorite Good	Second Favorite Good
1	1	2	1
2	2	3	2
...
$k - 1$	$k - 1$	k	$k - 1$
k (if $b = 0$)	k	k	Does not matter
k (if $b = 1$)	k	1	k
$k + 1 \dots n$	k	k	Does not matter

Fig. 2. The endowments and preferences for Theorem 3.2

Consider the mechanism $M_i(b) = M_S(\mathbf{x}(b))_i$ for $i \neq k$ (which is (ϵ, δ) -differentially private, since M_S is marginally differentially private). Let $f_1 : \mathcal{G} \rightarrow \{0, 1\}$ be the indicator function for the event that 1 receives good 2; e.g. $f_1(1) = 0$ and $f_1(2) = 1$. Then,

²Note that, since we are only considering IR mechanisms, preferences need only be specified to the level where i ranks g_i , and each agent i in our example has good g_i as her first or second choice.

define $M'(b) = f_1(M_1(b))$, which is also (ϵ, δ) -differentially private, due to f being a post processing function. By individual rationality, $M'(0) = 0$ with probability 1. Thus, by Claim 1,

$$\mathbb{P}[M'(1) = 1] \leq \frac{e^\epsilon + \delta}{e^\epsilon + 1}$$

Thus,

$$\mathbb{P}[M'(1) = 0] \geq 1 - \frac{e^\epsilon + \delta}{e^\epsilon + 1} \quad (14)$$

When $M'(1) = 0$, $M_i(1) = g_i$ (each agent was allocated her initial endowment). Consider the allocation $\pi(i) = i + 1 \pmod k$ for $i \in [k]$ and $\pi(i) = g_i = k$ for $i = k + 1, \dots, n$. Agents $1, \dots, k$ prefer π to $M_S(\mathbf{x}(1))$ when $M'(1) = 0$, and all other agents are indifferent. Since M_S is α -Pareto optimal, but there exists some π which k agents prefer and no agent likes less,

$$\alpha \geq (1 - \beta) \frac{k}{n} \mathbb{P}[M'(1) = 0] \geq (1 - \beta) \frac{k}{n} \left(1 - \frac{e^\epsilon + \delta}{e^\epsilon + 1}\right)$$

by Equation 14, which completes the proof. \square

PROOF OF COROLLARY 3.3. With the same hypotheses as Theorem 3.2 we know that the number k of types of goods must satisfy

$$k \leq \frac{n\alpha(e^\epsilon + 1)}{(1 - \delta)(1 - \beta)}. \quad (15)$$

Hence, we must have some good with at least $\frac{(1 - \delta)(1 - \beta)}{\alpha(e^\epsilon + 1)}$ copies, otherwise if we let n_j be the number of goods of type $j \in [k]$ we have

$$n = \sum_{j=1}^k n_j < k \left(\frac{(1 - \delta)(1 - \beta)}{\alpha(e^\epsilon + 1)} \right) \underbrace{\leq}_{(15)} n$$

\square

B. APPENDIX - FORMAL ALGORITHM PTTC

ALGORITHM 3: Private Top Trading Cycles

PTTC (\mathbf{g}, \succ) ;

Input: Exchange Market $\mathbf{x} = (\mathbf{g}, \succ)$.

Output: An allocation to all agents

Parameters: $\beta, \delta, \epsilon > 0$

Set $\epsilon' = \frac{\epsilon \log(k^3/\beta)}{2\sqrt{8}(\log(k^3/\beta)\sqrt{k \log(1/\delta_1)} + k\sqrt{k \log(1/\delta_2)})}$ and $E = \frac{\log(k^3/\beta)}{\epsilon'}$

Initialize $t \leftarrow 1$, $\tau \leftarrow 0$, and $P_e(1, 0) \leftarrow P_e$ is an ordered set of all n agents by their index.

(1) (Arc Weight Noise) For each $e \in A$ we set

$$\hat{w}_e(t, \tau) = w_e(t, \tau) + Z_e^t - 2E \quad \text{where } Z_e^t \sim \text{Lap}(1/\epsilon').$$

(2) (Clear Cycle) **while** there is a cycle with positive weight **do** $\tau \leftarrow \tau + 1$

Denote the cycle by $C(t, \tau)$ and let $\hat{W}(t, \tau) \leftarrow \min_{e \in C(t, \tau)} \{\lfloor \hat{w}_e(t, \tau) \rfloor\}$.

(3) (Trade) **for** $e = (u, w) \in C(t, \tau)$ **do**

Set $S_e(t, \tau) \leftarrow \text{R-SELECT}(\hat{W}(t, \tau), P_e(t, \tau))$. Update:

$$P_e(t, \tau) \leftarrow P_e(t, \tau - 1) \setminus S_e(t, \tau)$$

$$w_e(t, \tau) \leftarrow w_e(t, \tau - 1) - \hat{W}(t, \tau) \quad \& \quad \hat{w}_e(t, \tau) \leftarrow \hat{w}_e(t, \tau - 1) - \hat{W}(t, \tau)$$

$$\pi(i) = v \quad \forall i \in S_e(t, \tau) \text{ where } g_i = u.$$

(4) (No Cycle - Deletion) If there is no cycle with positive rounded down noisy weight, then there exists a node v s.t. $\hat{n}_v(t, \tau) = \sum_{e: e=(v, w)} \hat{w}_e(t, \tau) < k$ (Lemma 4.4). Let

$$IN_v = \{e \in A : e = (u, v) \text{ some } u \in V^t\} \quad \& \quad OUT_v = \{e \in A : e = (v, u) \text{ some } u \in V^t\}$$

We then update:

$$V^{t+1} \leftarrow V^t \setminus \{v\} \quad \& \quad A \leftarrow A \setminus \{OUT_v \cup IN_v\}.$$

for all $e = (u, w) \in A$ **do**

Define $B_e(t) = \{i \in P_{(u, v)}(t, \tau) : w \succeq_i g \quad \forall g \in V^{t+1}\}$ and update:

$$w_e(t+1, 0) \leftarrow w_e(t, \tau) + |B_e(t)| \quad \& \quad P_e(t+1, 0) \leftarrow P_e(t, \tau) \cup B_e(t)$$

and assign goods

$$\pi(i) = g_i \quad \forall i \in P_e(t, \tau) \text{ where } e \in OUT_v.$$

Set $t \leftarrow t + 1$ and $\tau \leftarrow 0$. Return to Step 1.

(5) (Clean-up for IR)

if $(w_e(t, \tau) < 0)$ **then** set $\pi(i) = g_i \forall i$ and HALT.

else Return to Step 1.

return π

C. FORMAL PROOFS FROM PRIVATE TOP TRADING CYCLES SECTION 4

PROOF OF CLAIM 2. Recall that for a random variable $Z \sim \text{Lap}(b)$ we have

$$\mathbb{P}(|Z| \geq \mu \cdot b) = e^{-\mu}$$

We then have, for $\mu = \log(k^3/\beta)$ and $b = \frac{1}{\epsilon'}$

$$\mathbb{P}\left(|Z| \geq \frac{\log(k^3/\beta)}{\epsilon'}\right) = \frac{\beta}{k^3}.$$

Now our algorithm in a fixed round t will sample a new Laplace random variable Z_e^t at most k^2 times - for each of the $O(k^2)$ arcs, one Laplace random variable is sampled

and there are as many as k rounds according to Lemma 4.3. Hence, we can obtain the following bound,

$$\mathbb{P}\left(|Z_e^t| \leq \frac{\log(k^3/\beta)}{\epsilon'} \quad \forall e \in A, \forall t \in [k]\right) \geq 1 - \beta.$$

Hence we can then lower bound and upper bound the difference between the error in the noisy arc weights and the actual arc weights to get the relation in (3):

$$E \leq w_e - \hat{w}_e = 2E - Z_e \leq 3E.$$

□

PROOF OF LEMMA 4.11. We first fix any agent i and agent types $\mathbf{x} = (\mathbf{x}_{-i}, x_i)$ and $\mathbf{x}' = (\mathbf{x}_{-i}, x'_i)$. Let us first consider the noisy arc weights $\hat{w}(1, 0)$ for round $t = 1$. We define $M^1 : \mathcal{X}^n \rightarrow \mathbb{R}^{k^2}$ as

$$M^1(\mathbf{x}) = \hat{w}(1, 0).$$

We know that this computation just uses the Laplace Mechanism from Algorithm 2 where we add Laplace noise with parameter $1/\epsilon'$ to each component of the exact edge weights $w(1, 0)$. When agent i changes her data from x_i to x'_i , she can change at most 2 entries in $w(1, 0)$ by at most 1. Thus, we have shown that M^1 is $2\epsilon'$ -differentially private.

Throughout round 1, agents will trade goods and be removed from consideration once they have a good. To figure out how much the arc weights will be impacted in future rounds, we use the algorithm R-SELECT that selects a given number of agents to trade (that is determined by $\hat{w}(1, 0)$) from those that want to trade in a random way. We denote here the random vector σ_1^τ where each of its entries correspond to the random value selected internally in R-SELECT. This vector of values σ_1^τ determines who trades from each arc in cycle $C(1, \tau)$. We write the randomness from round one as $r_1 = (\hat{w}(1, 0), \sigma_1)$ where $\sigma_1 = (\sigma_1^\tau)_{\tau \in [k^2]}$.

We will now assume that we have the randomness from all previous rounds, i.e. $\vec{r}_t = (r_1, \dots, r_t)$, and consider round $t + 1$. Our algorithm PTTC can compute the exact arc weights at round $t + 1$ as a function of these past random vectors and the data of the agents. We now ask, what is the sensitivity of the exact arc weights at round $t + 1$ when agent i changes reports and we condition on all the prior randomization \vec{r}_t ? We consider the following cases when agent i reports x_i and compare to when i reports x'_i instead:

- Agent i has yet to be selected after reporting x_i at the start of round $t + 1$.
 - Agent i has yet to be selected after reporting x'_i at the start of round $t + 1$. In this case when we fix all the prior information, then the sensitivity in the actual arc count vector $w(t + 1, 0)$ is at most 2 because agent i may move from one arc to another.
 - Agent i was selected at an earlier round $t' \leq t$ with report x'_i . Fixing all the prior information, when agent i was selected on cycle $C(t', \tau)$ for some τ when she reported x'_i , she must have caused some agent $i_{t', \tau}$ to not get chosen that was chosen when i reported x_i . This can further cascade to when agent $i_{t', \tau}$ gets chosen at a later cycle; someone that was selected when i reported x_i will not be selected when i reports x'_i instead. However, once agent i was selected when reporting x'_i , she has no further impact on the exact arc weights other than the person she displaced from her trade. Further, when the people that were displaced by i get matched, they are eliminated from further consideration. Note that we are assuming that i is still contributing weight to an arc when i reported x_i , because

i has yet to be chosen when reporting x_i . Thus in this case, conditioning on all the prior information, the sensitivity of the exact arc weight vector at round $t + 1$ is at most 2.

- Agent i was selected at an earlier round, say $t' \leq t$ with report x_i .
 - Agent i was selected at an earlier round with report x'_i . This follows a similar analysis as above when agent i reports x'_i and gets selected at round $t' \leq t$ when compared with i getting chosen after round $t + 1$ with report x_i .
 - Agent i has yet to be selected at round $t + 1$ with report x'_i . We have already considered this case above, with the roles of x_i and x'_i switched.

Thus, we have shown that conditioning on all the prior randomization, the exact arc counts $w(t + 1, 0)$ are 2-sensitive.

We now want to compute the noisy arc weights $\hat{w}(t + 1, 0)$ for round $t + 1$ given \vec{r}_t . We will denote \mathcal{R}^* to be the space where there can be an arbitrary number of random vectors, like \vec{r}_t . We define $M^{t+1} : \mathcal{X}^n \times \mathcal{R}^* \rightarrow \mathbb{R}^{k^2}$ to be

$$M^{t+1}(\mathbf{x}, \vec{r}_t) = \hat{w}(t + 1, 0).$$

Hence for any given randomization vectors \vec{r}_t the mechanism $M^{t+1}(\mathbf{x}, \vec{r}_t)$ is $2\epsilon'$ -differentially private in the database \mathbf{x} .

The full mechanism $\mathcal{M} : \mathcal{X}^n \rightarrow \mathbb{R}^{k \times k^2}$ that computes all the noisy arc weights in every round of PTTC can be defined as

$$\mathcal{M}(\mathbf{x}) = (\hat{w}(t, 0))_{t \in [k]}.$$

We then apply advanced composition from Theorem 4.9 and use the fact that the random integers σ_t computed in R-SELECT are chosen independent of the data to get that \mathcal{M} is (ϵ_1, δ_1) -differentially private where $\delta_1 > 0$ and

$$\epsilon_1 = 2\epsilon' \cdot \sqrt{8k \log(1/\delta_1)}$$

□

PROOF OF LEMMA 4.12. We assume that we have the arc weights \hat{w} for every round. Consider round t and cycle τ that has player $j \in P_e(t, \tau)$ on an arc $e \in C(t, \tau)$. The probability that agent j is selected depends on the exact number of people on edge e , that is $w_e(t, \tau)$ and the number of people $\hat{W}(t, \tau)$ that are being cleared along that cycle $C(t, \tau)$. Recall in the proof of Lemma 4.11 we showed that the weights $w_e(t, 0)$ at the beginning of round t were a function of the reported types \mathbf{x} and the prior randomization, that we called \vec{r}_{t-1} . The randomization terms include the noisy arc weights for rounds before t , which are part of \hat{w} , but it also includes the random integers σ calculated in R-SELECT which is computed independent of the data.

In order to find the exact number of people at an arc along cycle $C(t, \tau)$, we need to know the prior randomization terms and the random values $\sigma_t^1, \dots, \sigma_t^{\tau-1}$. Further, we showed in the previous lemma that each entry $w_e(t, \tau)$ for $e \in C(t, \tau)$ may change by at most 1 when agent i changes her reported type. We let $w'_e(t, \tau) \in \{w_e(t, \tau) - 1, w_e(t, \tau), w_e(t, \tau) + 1\}$ denote the number of people on arc $e \in C(t, \tau)$ when agent i changes her type to x'_i , but the types of the other agents remains the same. We will write the randomization terms as $\vec{r}_{t, \tau} = (\vec{r}_{t-1}, \hat{w}(t, 0), \sigma_t^1, \dots, \sigma_t^{\tau-1})$. We define the mechanism $M_j^{t, \tau} : \mathcal{X}^{n-1} \times \mathcal{X} \times \mathcal{R}^* \rightarrow \{0 \cup \mathcal{G}\}$ as

$$M_j^{t, \tau}(\mathbf{x}_{-j}, x_j, \vec{r}_{t, \tau}) = \begin{cases} g & \text{If } j \in S_e(t, \tau) \text{ for } e = (g_j, g) \in C(t, \tau) \\ 0 & \text{otherwise} \end{cases}$$

We will now make use of our assumption that the difference between the noisy arc weights and the actual arc weights satisfies (3). This gives us

$$\begin{aligned} \frac{\mathbb{P}(M_j^{t,\tau}(\mathbf{x}_{-j}, x_j, \vec{r}_{t,\tau}) = g)}{\mathbb{P}(M_j^{t,\tau}((\mathbf{x}_{-i,j}, x'_i), x_j, \vec{r}_{t,\tau}) = g)} &= \frac{\frac{\hat{W}(t,\tau)}{w_e(t,\tau)}}{\frac{\hat{W}(t,\tau)}{w'_e(t,\tau)}} \leq \frac{\frac{\hat{W}(t,\tau)}{w_e(t,\tau)}}{\frac{\hat{W}(t,\tau)}{w_e(t,\tau)+1}} \\ &= \frac{w_e(t,\tau)+1}{w_e(t,\tau)} \leq 1 + \frac{1}{E} \leq e^{1/E} \end{aligned} \quad (16)$$

The first inequality follows from $w'_e(t,\tau) \leq w_e(t,\tau) + 1$. The second to last inequality comes from the fact that the cycle always has an arc with $\hat{w}_e = w_e + Z_e - 2E \geq 1 \implies w_e \geq E + 1$.

We now consider the case when j is not selected at round t , cycle τ when his arc is on the cycle.

$$\begin{aligned} \frac{\mathbb{P}(M_j^{t,\tau}(\mathbf{x}_{-j}, x_j, \vec{r}_{t,\tau}) = 0)}{\mathbb{P}(M_j^{t,\tau}((\mathbf{x}_{-i,j}, x'_i), x_j, \vec{r}_{t,\tau}) = 0)} &= \frac{1 - \frac{\hat{W}(t,\tau)}{w_e(t,\tau)}}{1 - \frac{\hat{W}(t,\tau)}{w'_e(t,\tau)}} \leq \frac{1 - \frac{\hat{W}(t,\tau)}{w_e(t,\tau)}}{1 - \frac{\hat{W}(t,\tau)}{w_e(t,\tau)-1}} \\ &= \frac{(w_e(t,\tau) - \hat{W}(t,\tau))(w_e(t,\tau) - 1)}{w_e(t,\tau) \cdot (w_e(t,\tau) - \hat{W}(t,\tau) - 1)} \leq \frac{w_e(t,\tau) - \hat{W}(t,\tau)}{w_e(t,\tau) - \hat{W}(t,\tau) - 1}. \end{aligned} \quad (17)$$

The first inequality holds because $w'_e(t,\tau) \geq w_e(t,\tau) - 1$. Recall that we have $w_e(t,\tau) - \hat{w}_e(t,\tau) \geq E$ from (3) $\implies w_e(t,\tau) - \hat{W}(t,\tau) \geq E$. Hence, we can further bound our ratio in (17) by

$$\frac{\mathbb{P}(M_j^{t,\tau}(\mathbf{x}_{-j}, x_j, \vec{r}_{t,\tau}) = 0)}{\mathbb{P}(M_j^{t,\tau}((\mathbf{x}_{-i,j}, x'_i), x_j, \vec{r}_{t,\tau}) = 0)} \leq \frac{E}{E-1} \leq 1 + \frac{2}{E} \leq e^{2/E}. \quad (18)$$

Hence, fixing all randomness prior to selecting people at cycle $C(t,\tau)$, each $M_j^{t,\tau}$ is $2/E$ -differentially private with respect to the database \mathbf{x}_{-j} for all $\tau \in [k^2]$ and $t \in [k]$. The mechanism \mathcal{M}_j in the statement of the lemma is then just a composition of the mechanisms $M_j^{t,\tau}$ over all rounds $t \in [k]$ and cycles $\tau \in [k^2]$ ³. We have shown that $M_j^{t,\tau}$ is $2/E$ -differentially private in data \mathbf{x}_{-j} . Thus, applying advanced composition in Theorem 4.9, we get that \mathcal{M}_j is (ϵ_2, δ_2) differentially private for $\delta_2 > 0$ and

$$\epsilon_2 = 2k\sqrt{8k\log(1/\delta_2)}/E$$

□

³In the theorem statement we are conditioning on the weights \hat{w} instead of all prior randomness \vec{r} . The only additional terms in \vec{r} that are not \hat{w} are the values σ that determine who trades at each cycle in all prior rounds. However, these values σ are chosen independently of our data and so do not change our privacy analysis if we condition on \vec{r} or \hat{w} .