

Private Matchings and Allocations

Justin Hsu* Zhiyi Huang† Aaron Roth‡
Tim Roughgarden§ Zhiwei Steven Wu¶

November 13, 2013

Abstract

We consider a private variant of the classical *allocation problem*: given m goods and n agents with individual, private valuation functions over bundles of goods, how can we partition the goods amongst the agents to maximize social welfare? Specifically, the valuation functions are sensitive information which the agents wish to keep private from arbitrary coalitions of other agents. An important special case is when each agent desires at most one good, and specifies her (private) value for each good: in this case, the problem is exactly the maximum-weight matching problem in a bipartite graph.

Private matching and allocation problems have not been considered in the differential privacy literature, and for good reason: they are plainly impossible to solve under the standard notion of differential privacy. Informally, the allocation must match agents to preferred goods in order to maximize social welfare, but this preference is exactly what agents wish to keep private! Therefore, we consider the problem under the recently introduced constraint of *joint differential privacy*: roughly, for any agent i , no coalition of agents excluding i should be able to learn about the valuation function of agent i . We first show that if there are a small number of identical copies of each good, then it is possible to efficiently and accurately solve the maximum weight matching problem while guaranteeing joint differential privacy. We then extend our techniques to the more general allocation problem, when bidder valuations satisfy the *gross substitutes* condition. Finally, we prove lower bounds demonstrating that the problem cannot be privately solved to non-trivial accuracy without requiring multiple copies of each type of good.

*Department of Computer and Information Science, University of Pennsylvania. Supported in part by NSF Grant CNS-1065060. Email: justhsu@cis.upenn.edu.

†Department of Computer Science, Stanford University. Email: hziyi@stanford.edu. This work is done in part while the author was a student at the University of Pennsylvania.

‡Department of Computer and Information Science, University of Pennsylvania. Supported in part by an NSF CAREER award, NSF Grants CCF-1101389 and CNS-1065060, and a Google Focused Research Award. Email: aaroth@cis.upenn.edu.

§Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. This research was supported in part by NSF Awards CCF-1016885 and CCF-1215965 and an ONR PECASE Award. Email: tim@cs.stanford.edu. This work done in part while visiting University of Pennsylvania.

¶Department of Computer and Information Science, University of Pennsylvania. Supported in part by NSF Grant CCF-1101389. Email: wuzhiwei@cis.upenn.edu

1 Introduction

The classic maximum-weight matching problem in bipartite graphs can be viewed as follows: there are m goods $j \in \{1, \dots, m\}$ and n buyers $i \in \{1, \dots, n\}$. Each buyer i has a value $v_{ij} \in [0, 1]$ for each good j , and the goal is to find a matching μ between goods and buyers which maximizes the social welfare: $\text{SW} = \sum_{i=1}^n v_{i, \mu(i)}$. When the goods are sensitive¹ it is natural to ask for a matching that hides the reported values of each of the players.

It is not hard to see that this is impossible under the standard notion of differential privacy, which insists that the allocation must be insensitive to the reported valuations of each player. We formalize this in Section 5, but the intuition is simple: consider the case with two types of goods (label them “0” and “1”) with n identical copies each, and suppose that each buyer has a private preference for one of the two types: value 1 for the good that he likes, and value 0 for the other good. There is no contention since the supply of each good is larger than the total number of buyers, so any allocation achieving social welfare $\text{OPT} - \alpha n$ can be used to reconstruct a $(1 - \alpha)$ fraction of the preferences. This is impossible for non-trivial values of α under differential privacy.

In light of this observation, is there any hope for privately solving max-weight matching problems? In this paper, we show that the answer is *yes*: it is possible to solve matching problems (and more general allocation problems) to high accuracy assuming a small number of identical copies of each good, while still satisfying an extremely strong variant of differential privacy. We observe that the matching problem has the following two features:

1. Both the input and solution are naturally partitioned amongst the same n people: in our case, each buyer i receives the item $\mu(i)$ he is matched to in the solution.
2. The problem is not solvable privately because the item given to a buyer must reflect his private data, but this need not (necessarily) be the case for items given to other buyers.

By utilizing these two features, we show that the matching problem can be accurately solved under the recently introduced notion of *joint differential privacy* [Kearns et al., 2014]. Informally speaking, this requires that for every buyer i , the joint distribution on items $\mu(j)$ for $j \neq i$ must be differentially private in the reported valuation of buyer i . As a consequence, buyer i ’s privacy is protected even if *all* other buyers collude against him, potentially sharing the identities of the items they receive. As long as buyer i does not reveal his own item, his privacy is protected.

We then show that our techniques generalize well beyond the max-matching problem, to the more general *allocation* problem—in this setting, each buyer i has a valuation function defined over subsets of goods $v_i : 2^{[m]} \rightarrow [0, 1]$ from some class of valuations \mathcal{C} , and the goal is to find a partition of the goods S_1, \dots, S_n maximizing social welfare. (Note that the max-weight matching problem is the special case when agents are *unit demand*, i.e., only want bundles of size 1). We generalize our algorithm to solve the allocation problem when bidders’ valuations satisfy the *gross substitutes* condition. This is an economically meaningful class of valuation functions that form a strict subclass of submodular functions, and (as we will explain) are the most general class of valuation functions for which our techniques could possibly apply.

¹ For instance, the goods might be related to the treatment of disease, or might be indicative of a particular business strategy, or might be embarrassing in nature.

1.1 Our Techniques and Results

Our approach makes a novel connection between *market clearing prices* and differential privacy. Prices have long been considered as a low information way to coordinate markets; conceptually, our paper formalizes this intuition in the context of differentially private allocation. Our algorithm is a differentially private implementation of m simultaneous ascending price auctions, one for each type of good. Following the classic analysis of Kelso and Crawford [1982], the prices in these auctions converge to *Walrasian equilibrium prices*: prices under which each buyer is simultaneously able to buy his most preferred bundle of goods. We show that although the allocation itself cannot be computed under standard differential privacy, the Walrasian equilibrium prices can be, and that the computation of these prices can be used to coordinate a high welfare allocation while satisfying joint differential privacy.

The classical ascending price auction works as follows: each good begins with a price of 0, and each agent is initially unmatched to any good. Unmatched agents i take turns bidding on the good j^* that maximizes their utility at the current prices: i.e., $j^* \in \arg \max(v_{ij} - p_j)$. When a bidder bids on a good j , he becomes the new high bidder and the price of j^* is incremented. Bidders are considered to be tentatively matched to a good while they are the high bidder. The auction continues until there are no unmatched bidders who would prefer to be matched to any of the goods at the current prices. The algorithm necessarily converges because each bid increases the sum of the prices of the goods, and prices are bounded by some finite value.² Moreover, by construction, every bidder ends up matched to their most preferred good given the prices. Finally, by the so-called “First Welfare Theorem”, any matching that corresponds to these (Walrasian equilibrium) prices is necessarily social-welfare maximizing. We emphasize that it is this final implication that is the key: “prices” play no role in our problem description, nor do we ever actually charge “prices” to the agents—the prices are purely a convenient device.

We give an approximate, private version of this algorithm based on several observations. First, in order to implement this algorithm, it is sufficient to maintain the sequence of prices of the goods privately: given a record of the price trajectory, each agent can figure out for himself what good he is matched to. Second, in order to privately maintain the prices, it suffices to maintain a private count of the number of bids each good has received over the course of the auction. Finally (and most subtly), it turns out that it is possible to halt the algorithm early without significantly harming the quality of the final matching. This guarantees that no bidder ever makes more than a small number (independent of both n and m) of total bids, which allows us to bound the sensitivity of the bid-counters. Together, these observations allow us to implement the auction privately using the work of Dwork et al. [2010a] and Chan et al. [2011], which give counters that are private under continual observation.³ The result is an algorithm that converges to a matching together with prices that form an approximate Walrasian equilibrium. We complete our analysis by proving an approximate version of the first welfare theorem, which shows that the matching must have high weight.

The algorithm of Kelso and Crawford [1982] extends to the general allocation problem when players have gross substitute preferences, and our private algorithm does as well. We note that

² Bidders do not bid on goods for which they have negative utility; in our case, and $v_{ij} \in [0, 1]$

³ These papers show that a single sensitivity-1 counter can be implemented privately. We need a stronger version of this theorem in our work: a collection of counters operating on adaptively chosen streams, whose sensitivity can be bounded jointly, should be private with a privacy parameter that depends only on the joint-sensitivity of the streams and is independent of the number of counters. This stronger statement also holds, as we show in our privacy analysis.

this class of preferences is the natural limit of our approach, which makes crucial use of equilibrium prices as a coordinating device: in general, when agents have valuations over bundles of goods that do not satisfy the gross substitutes condition, Walrasian equilibrium prices may not exist. Privately solving allocation problems in such settings would therefore require completely different techniques; whether or not such problems are solvable is an intriguing open question.

Finally, we give lower bounds showing that our results are qualitatively tight: not only is the problem impossible to solve under the standard constraint of differential privacy (rather than joint differential privacy), to get any non-trivial solution, it is necessary to assume that there are multiple copies of each type of good. Our lower bounds are all fundamentally reductions to database reconstruction attacks. Our lower bound for joint-differentially private algorithms may be of general interest, as we believe it forms a good template for other lower bounds for joint differential privacy.

Here we state an informal version of our main result in the special case of max-matchings, which we prove in Section 3. Our more general theorem for allocation problems with gross substitutes preferences, which we prove in Section 4, generalizes this result. Here, privacy is protected with respect to a single agent i changing his valuations v_{ij} for possibly *all* goods j .

Theorem (Informal). *There is a computationally efficient ε -joint differentially private algorithm which computes a matching of weight $\text{OPT} - \alpha n$ in settings in which there are n agents and k types of goods, with s copies of each good when:*

$$s \geq O\left(\frac{1}{\alpha^3 \varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}\right)\right).$$

In certain settings, the welfare guarantee can be improved to $(1 - \alpha) \text{OPT}$.

We complement this result with several lower bounds in Section 5. We show that no algorithm can solve the private max-matchings problem to non-trivial accuracy under the standard constraint of differential privacy. We also show that even under the constraint of joint differential privacy, it is necessary to assume that there are multiple copies of each item:

Theorem (Informal). *No joint differentially private algorithm can compute matchings of weight greater than $\text{OPT} - \alpha n$ on instances in which there are n agents and s copies of each good, when*

$$s \leq O\left(\frac{1}{\sqrt{\alpha}}\right).$$

In particular, no algorithm can compute matchings of weight $\text{OPT} - o(n)$ on instances for which the supply $s = O(1)$. In addition, we show that when goods have supply only $s = O(1)$, it is not even possible to compute the equilibrium prices privately.

1.2 Related Work

Differential privacy, first defined by [Dwork et al. \[2006\]](#), has become a standard “privacy solution concept” in the theoretical computer science literature. There is far too much work to survey comprehensively; we mention only the most relevant work. (For a textbook introduction, see [Dwork and Roth \[2013\]](#).)

The privacy of our algorithms relies on work by [Dwork et al. \[2010a\]](#) and [Chan et al. \[2011\]](#), who show how to release a running count of a stream of bits under *continual observation*—i.e.,

report the online count at every timestep as the stream is revealed, provide high accuracy at every point in time, while guaranteeing that the entire transcript is differentially private.

Beginning with [Dinur and Nissim \[2003\]](#), much work in differential privacy has focused on answering numeric valued queries on a private dataset (e.g., [Dwork et al. \[2006\]](#), [Blum et al. \[2013\]](#), [Hardt and Rothblum \[2010\]](#), among many others). In contrast, work on private combinatorial optimization problems has been sporadic (despite several papers, e.g., [Nissim et al. \[2007\]](#), [Gupta et al. \[2010\]](#)). Part of the reason is that many combinatorial optimization problems are impossible to solve under differential privacy (including the allocation problems we consider in this paper). To sidestep this problem, we employ the recently introduced solution concept of *joint differential privacy*. First formalized by [Kearns et al. \[2014\]](#), similar ideas are present in the vertex and set-cover algorithms of [Gupta et al. \[2010\]](#), and in the analyst private data analysis algorithms of [Dwork et al. \[2012\]](#), [Hsu et al. \[2013\]](#).

The utility of our algorithm relies on analysis due to [Kelso and Crawford \[1982\]](#), who study the problem of matching *firms* to *workers* when the firms have preferences that satisfy the *gross substitutes* condition. They give an algorithm based on simulating simultaneous ascending auctions that converge to *Walrasian equilibrium prices*, together with a corresponding matching. In this respect, our approach is complete: [Gul and Stacchetti \[1999\]](#) show that gross substitutes preferences are precisely the set of preferences for which Walrasian equilibrium prices are guaranteed to exist.

While our approximate equilibrium achieves good approximation to the optimal welfare at the expense of certain incentive properties, our work is closely related to recent work on privately computing various kinds of equilibrium in games (e.g., correlated equilibrium [[Kearns et al., 2014](#)], Nash equilibrium [[Rogers and Roth, 2013](#)], and minmax equilibrium [[Hsu et al., 2013](#)]). These works belong to a growing literature studying the interface of game theory and differential privacy; for a recent survey, see [Pai and Roth \[2013\]](#).

2 Preliminaries

2.1 The Allocation Problem

We consider allocation problems defined by a set of goods G , and a set of n agents $[n]$. Each agent $i \in [n]$ has a *valuation function* $v_i : 2^G \rightarrow [0, 1]$ mapping bundles of goods to values. A *feasible allocation* is a collection of sets $S_1, \dots, S_n \subseteq G$ such that $S_i \cap S_j = \emptyset$ for each $i \neq j$: i.e., an assignment of goods to agents such that no good is allocated to more than one agent. The *social welfare* of an allocation S_1, \dots, S_n is defined to be $\sum_{i=1}^n v_i(S_i)$, the sum of the agent's valuations for the allocation; we are interested in finding allocations which maximize this quantity. Given an instance of an allocation problem, we write $\text{OPT} = \max_{S_1, \dots, S_n} \sum_{i=1}^n v_i(S_i)$ to denote the social welfare of the optimal feasible allocation.

A particularly simple valuation function is a *unit demand valuation*, where bidders demand at most one item. Such valuation functions take the form $v_i(S) = \max_{j \in S} v_i(\{j\})$, and can be specified by numbers $v_{i,j} = v_i(\{j\}) \in [0, 1]$, which represent the value that bidder i places on good j . When bidders have unit demand valuations, the allocation problem corresponds to computing a maximum weight matching in a bipartite graph.

Our results will hold for the more general class of *gross substitute valuations*, which include unit demand valuations as a special case. Informally, gross substitute valuations must have the following property: any set of goods S' that are in a most-demanded bundle at some set of prices

p remain in a most-demanded bundle if the prices of *other* goods are raised, keeping the prices of goods in S' fixed. Gross substitute valuations are an economically meaningful class of valuation functions, and are a strict subclass of submodular functions. They are also precisely the valuation functions with Walrasian equilibria in markets with indivisible goods [Gul and Stacchetti, 1999].

Before giving the formal definition, we first introduce some notation. Given a vector of prices $\{p_g\}_{g \in G}$, the (quasi-linear) *utility* that player i has for a bundle of goods S_i is defined to be $u_i(S_i, p) = v_i(S_i) - \sum_{j \in S_i} p_j$.⁴ Given a vector of prices p , for each agent i , we can define his set of *most demanded bundles*: $\omega(p) = \arg \max_{S \subseteq G} u_i(S, p)$. Given two price vectors p, p' , we write $p \preceq p'$ if $p_g \leq p'_g$ for all g .

Definition 1. A valuation function $v_i : 2^G \rightarrow [0, 1]$ satisfies the gross substitutes condition if for every pair of price vectors $p \preceq p'$, and for every set of goods $S \in \omega(p)$, if $S' \subseteq S$ satisfies $p'_g = p_g$ for every $g \in S'$, then there is a set $S^* \in \omega(p')$ with $S' \subseteq S^*$.

Finally, we consider markets with multiple copies of each type of good. Two goods $g_1, g_2 \in G$ are *identical* if for every bidder i and for every bundle $S \subseteq G$, $v_i(S \cup \{g_1\}) = v_i(S \cup \{g_2\})$: i.e., the two goods are indistinguishable according to every valuation function. Formally, we say that a set of goods G consists of k *types* of goods with s *supply* if there are k representative goods $g_1, \dots, g_k \in G$ such that every good $g' \in G$ is identical to one of g_1, \dots, g_k , and for each representative good g_i , there are s goods identical to g_i in G . For simplicity of presentation we assume throughout the paper that the supply of each good is identical, but this is not necessary. All of our results continue to hold when the supply s denotes the *minimum* supply of any type of good.

2.2 Differential Privacy Preliminaries

Although it is impossible to solve the allocation problem under standard differential privacy (see Section 5), standard differential privacy plays an essential role in our analysis.

Suppose agents have valuation functions v_i from a class of functions C . A database $D \in C^n$ is a vector of valuation functions, one for each of the n bidders. Two databases D, D' are *i -neighbors* if they differ in only their i 'th index: that is, if $D_j = D'_j$ for all $j \neq i$. If two databases D, D' are i -neighbors for some i , we say that they are *neighboring databases*. We will be interested in randomized algorithms that take a database as input, and output an element from some range \mathcal{R} . Our final mechanisms will output sets of n bundles (so $\mathcal{R} = (2^G)^n$), but intermediate components of our algorithms will have different ranges.

Definition 2 (Dwork et al. [2006]). An algorithm $\mathcal{M} : C^n \rightarrow \mathcal{R}$ is (ϵ, δ) -differentially private if for every pair of neighboring databases $D, D' \in C^n$ and for every set of subset of outputs $S \subseteq \mathcal{R}$,

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

If $\delta = 0$, we say that \mathcal{M} is ϵ -differentially private.

When the range of a mechanism is also a vector with n components (e.g., $\mathcal{R} = (2^G)^n$), we can define *joint differential privacy*: this requires that simultaneously for all i , the *joint* distribution on outputs given to players $j \neq i$ is differentially private in the input of agent i . Given a vector

⁴ This is a natural definition of utility if agents must pay for the bundles they buy at the given prices. In this paper we are concerned with the purely algorithmic allocation problem, so our algorithm will not actually charge prices. However, prices will be a convenient abstraction throughout our work.

$x = (x_1, \dots, x_n) \in (2^G)^n$, we write $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n) \in (2^G)^{n-1}$ to denote the vector of length $n - 1$ which contains all coordinates of x except the i 'th coordinate.

Definition 3 (Kearns et al. [2014]). *An algorithm $\mathcal{M} : C^n \rightarrow (2^G)^n$ is (ε, δ) -joint differentially private if for every i , for every pair of i -neighbors $D, D' \in C^n$, and for every subset of outputs $S \subseteq (2^G)^{n-1}$,*

$$\Pr[\mathcal{M}(D)_{-i} \in S] \leq \exp(\varepsilon) \Pr[\mathcal{M}(D')_{-i} \in S] + \delta.$$

If $\delta = 0$, we say that \mathcal{M} is ε -joint differentially private.

Note that this is still an extremely strong definition that protects i from arbitrary coalitions of adversaries—it weakens the constraint of differential privacy only in that the output given specifically to agent i is allowed to be sensitive in the input of agent i .

2.3 Differentially Private Counters

The central tool in our matching algorithm is the private streaming counter proposed by Chan et al. [2011], Dwork et al. [2010a]. Given a bit stream $\sigma = (\sigma_1, \dots, \sigma_T) \in \{0, 1\}^T$, a streaming counter $\mathcal{M}(\sigma)$ releases an approximation to $c_\sigma(t) = \sum_{i=1}^t \sigma_i$ at every time step t .

Definition 4. *A streaming counter \mathcal{M} is (α, β) -useful if with probability at least $1 - \beta$, for each time $t \in [T]$,*

$$|\mathcal{M}(\sigma)(t) - c_\sigma(t)| \leq \alpha.$$

For the rest of this paper, let **Counter** (ε, T) denote the Binary Mechanism of Chan et al. [2011], instantiated with parameters ε and T . The mechanism produces a monotonically increasing count, and satisfies the following accuracy guarantee. (Further details may be found in Appendix A.)

Theorem 1 (Chan et al. [2011]). *For $\beta > 0$, **Counter** (ε, T) is ε -differentially private with respect to a single bit change in the stream, and (α, β) -useful for*

$$\alpha = \frac{2\sqrt{2}}{\varepsilon} \ln\left(\frac{2}{\beta}\right) \left(\sqrt{\log(T)}\right)^5.$$

3 A Private Algorithm for Maximum-Weight Matching

In this section, we study the special case of unit demand valuations. Though our later algorithm for gross substitutes valuations generalizes this case, we first present our algorithm in this simpler setting to highlight the key features of our approach.

Consider a matching market with n bidders and k different types of goods, where each good has supply s and bidder i has valuation $v_{ij} \in [0, 1]$ for good j . Some agents may not end up being matched to a good: to simplify notation, we will say that unmatched agents are matched to \perp , a special dummy good.

To reach a maximum weight matching, we first formulate an intermediate goal: we want to privately compute prices $p \in [0, 1]^k$ and an allocation of the goods $\mu : [n] \rightarrow [k] \cup \{\perp\}$ such that *most* bidders are matched with their *approximately* favorite goods *given the prices* and each over-demanded good almost clears, where a good is *overdemanded* if its price is strictly positive.⁵ We will show that if this intermediate goal is met, then in fact we have computed an approximately maximum weight matching.

⁵ This will be our notion of approximate Walrasian equilibrium.

Definition 5. A price vector $p \in [0, 1]^k$ and an assignment $\mu: [n] \rightarrow [k] \cup \{\perp\}$ of bidders to goods is an (α, β, ρ) -approximate matching equilibrium if

1. All but a ρ fraction of bidders i are matched to an α -approximate favorite good: i.e. for at least $(1 - \rho)n$ bidders i , $v_{i\mu(i)} - p_{\mu(i)} \geq v_{ij} - p_j - \alpha$ for every good j ; we call these bidders satisfied, and
2. the number of bidders assigned to any type of good does not exceed its supply, and
3. each overdemanded good clears except for at most β supply; recall that a good j is said to be overdemanded if $p_j > 0$.

3.1 Overview of the Algorithm

Algorithm 1 (**PrivateMatching**) is a variant of a *deferred acceptance* algorithm first proposed and analyzed by Kelso and Crawford [1982], which runs k simultaneous ascending price auctions: one for each type of good. At any given moment, each type of good has a *proposal price* p_j . In rounds, unsatisfied bidders take turns bidding on the good that maximizes their utility at the current prices: that is, the good j that maximizes $v_{ij} - p_j$. (This is the **Propose** function.)

The s most recent bidders for a type of good are tentatively matched to that type of good (these are the current *high bidders*). A bidder tentatively matched to a good with supply s becomes unmatched to that good (he has been *outbid*) once the good he is matched to receives s subsequent bids. Bidders keep track of which good they are matched to (in the variable μ), if any, and can determine whether they are currently matched or unmatched by looking at a count of the number of bids received by the last good they bid on. Finally, every s bids on a good increases its price by a fixed increment α .

To implement this algorithm privately, we count the number of bids each good has received using private counters. Unsatisfied bidders can infer the prices of all goods based on the number of bids each has received, and from this information, they determine which good to bid on (their favorite good at the given prices). Their bid is recorded by sending a “1” to the appropriate counter. (This is the **Bid** function.) Matched bidders remember the reading of the bid counter on the good they are matched to at the time that they last bid (in the variable d_i); when the counter ticks s bids past this initial count, the bidder concludes that he has been outbid, and becomes unmatched.

Since the private counters are noisy, the number of bidders matched to each good may deviate from s . To maintain feasibility, the auction is run with some supply m withheld: i.e., it is run as if the supply of each good were $s - m$, rather than s . The reserved supply m is used to satisfy the demand of all bidders who believe themselves to be matched to each type of good, which with high probability does not exceed the actual supply s .

Our algorithm stops as soon as we have a round in which ρn bidders or fewer place bids. We show that this early stopping condition does not significantly harm the welfare guarantee of the matching, while it substantially reduces the sensitivity of the counters: no bidder ever bids more than $O(1/(\alpha\rho))$ times in total. Crucially, this is independent of both the number of types of goods k , and the number of bidders n . This greatly improves the accuracy of the prices, which depend on the bid counts: the degree to which we have to perturb the bid counts to protect privacy is proportional to the sensitivity of the counters.

To privately implement this stopping condition, we maintain a separate counter (`counter(0)`) which counts the number of unsatisfied bidders throughout the run of the algorithm. At the end of

each proposal round, bidders who are unsatisfied will send “1” to this counter, and bidders who are matched will send “0”. If this counter increases by less than roughly ρn in any round, we terminate the algorithm and output the final matching. (This is the **CountUnsatisfied** function.)

Algorithm 1 PrivateMatching($\alpha, \rho, \varepsilon$)

Input: Bidders’ valuations on the goods ($\{v_{1j}\}_{j=1}^m, \dots, \{v_{nj}\}_{j=1}^m$)

Initialize:

$$T = \frac{8}{\alpha\rho}, \quad \varepsilon' = \frac{\varepsilon}{2T}, \quad E = \frac{2\sqrt{2}}{\varepsilon'}(\log nT)^{5/2} \log\left(\frac{4k}{\gamma}\right), \quad m = 2E + 1$$

$$\text{counter}(j) = \mathbf{Counter}(\varepsilon', nT), \quad p_j = c_j = 0 \quad \text{for every good } j,$$

$$\mu(i) = \emptyset, \quad d_i = 0, \quad \text{for every bidder } i, \quad \text{counter}(0) = \mathbf{Counter}(\varepsilon', nT)$$

Call **Propose** T times; **Output:** Prices p and allocation μ .

Propose:

for all bidders i **do**

if $\mu(i) = \emptyset$ **then**

 Let $\mu(i) := \operatorname{argmax}_j v_{ij} - p_j$.

if $v_{i\mu(i)} - p_{\mu(i)} \leq 0$ **then**

 Let $\mu(i) := \perp$ and **Bid**(0).

else Save $d_i := c_{\mu(i)}$ and **Bid**($\mathbf{e}_{\mu(i)}$).

else **Bid**(0)

CountUnsatisfied

Bid: On input bid vector \mathbf{b}

for all goods j **do**

 Feed \mathbf{b}_j to $\text{counter}(j)$.

 Update count $c_j := \text{counter}(j)$.

if $c_j \geq (p_j/\alpha + 1)(s - m)$ **then**

 Update $p_j := p_j + \alpha$.

CountUnsatisfied:

for all bidders i **do**

if $\mu(i) \neq \perp$ and $c_{\mu(i)} - d_i \geq s - m$ **then**

 Feed 1 to $\text{counter}(0)$; Let $\mu(i) := \emptyset$

else Feed 0 to $\text{counter}(0)$.

if $\text{counter}(0)$ increases by less than $\rho n - 2E$ **then**

 Halt; For each i with $\mu(i) = \emptyset$, let $\mu(i) = \perp$

3.2 Privacy Analysis

In this section, we show that the allocation output by our algorithm satisfies joint differential privacy with respect to a single bidder changing *all* of his valuations. We first show a basic but useful lemma: to show joint differential privacy, it is sufficient to show that the output sent to each agent i is an arbitrary function only of some global signal that is computed under the standard constraint of differential privacy, together with agent i ’s private data. We call this the *billboard model*: some message is viewable by all agents, as if placed on a public billboard, and this message is differentially private. If every agent can compute their own output given only what is on the billboard and their own private data, then the algorithm is joint differentially private. In our case, the price history over the course of the auction is the differentially private message posted on the billboard. From this information and their personal private valuation, each bidder can compute their personal allocation.

Lemma 1 (Billboard Lemma). *Suppose $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ is (ε, δ) -differentially private. Consider any set of functions $f_i : \mathcal{D}_i \times \mathcal{R} \rightarrow \mathcal{R}'$, where \mathcal{D}_i is the portion of the database containing i 's data. The composition $\{f_i(\Pi_i D, \mathcal{M}(D))\}$ is (ε, δ) -joint differentially private, where Π_i is the projection to i 's data.*

Proof. We need to show that for any agent i , the view of the other agents is (ε, δ) -differentially private when i 's private data is changed. Suppose databases D, D' are i -neighbors, so $\Pi_j D = \Pi_j D'$ for $j \neq i$. Let \mathcal{R}_{-i} be a set of views of the bidders besides i . Let $\mathcal{R}^* = \{r \in \mathcal{R} \mid \{f_j(\Pi_j D, r)\}_{-i} \in \mathcal{R}_{-i}\}$. Then, we wish to show

$$\begin{aligned} \Pr[\{f_j(\Pi_j D, \mathcal{M}(D))\}_{-i} \in \mathcal{R}_{-i}] &\leq e^\varepsilon \Pr[\{f_j(\Pi_j D', \mathcal{M}(D'))\}_{-i} \in \mathcal{R}_{-i}] + \delta \\ &= e^\varepsilon \Pr[\{f_j(\Pi_j D, \mathcal{M}(D'))\}_{-i} \in \mathcal{R}_{-i}] + \delta \\ \Pr[\mathcal{M}(D) \in \mathcal{R}^*] &\leq e^\varepsilon \Pr[\mathcal{M}(D') \in \mathcal{R}^*] + \delta, \end{aligned}$$

but this is true since \mathcal{M} is (ε, δ) -differentially private. \square

Theorem 2. *The sequence of prices and counts of unsatisfied bidders released by **PrivateMatching** $(\alpha, \rho, \varepsilon)$ satisfies ε -differential privacy.*

Proof Sketch. We give a rough intuition here, and defer the full proof to Appendix A. Note that the prices can be computed from the noisy counts, so it suffices to show that these counts are private. Since no bidder bids more than $T \approx 1/(\alpha\rho)$ times in total, the *total* sensitivity of the k price streams to a single bidder's valuations is only $O(1/(\alpha\rho))$ (independent of k) even though a single bidder could in principle bid $\Omega(1/\alpha)$ times on each of the k streams. Hence the analysis of these k simultaneously running counters is akin to the analysis of answering *histogram queries*—multiple queries whose joint sensitivity is substantially smaller than the sum of their individual sensitivities.

By setting the counter for each good with privacy parameter $\varepsilon' = \varepsilon/2T$, the prices should be $\varepsilon/2$ differentially private. By the same reasoning, setting the unsatisfied bidders counter with privacy parameter $\varepsilon' = \varepsilon/2T$ also makes the unsatisfied bidders count $\varepsilon/2$ private. Thus, these outputs together satisfy ε -differential privacy.

While this intuition is roughly correct, there are some technical details. Namely, [Chan et al. \[2011\]](#) show privacy for a single counter with sensitivity 1 on a non-adaptively chosen stream. Since intermediate outputs (i.e., prices) from our counters will affect the future streams (i.e., future bids) for other counters, this is not sufficient. In fact, it is possible to prove privacy for multiple counters running on adaptively chosen streams, where the privacy parameter depends only on the joint sensitivity of the streams, and not on the number of streams. We show this using largely routine arguments; details can be found in Appendix A. \square

Theorem 3. **PrivateMatching** $(\alpha, \rho, \varepsilon)$ *satisfies ε -joint differential privacy.*

Proof. Note that given the sequence of prices, counts of unsatisfied bidders, and the private valuation of any bidder i , the final allocation to that bidder can be computed by simulating the sequence of bids that bidder i would make: these are determined by the price at rounds at which bidder i is slotted to bid, and by whether the halting condition has been met. Bidder i 's final allocation is simply the final item that he bids on. The prices and halting condition are computed as a deterministic function of the noisy counts, which are ε -differentially private by Theorem 2. So, Lemma 1 shows that **PrivateMatching** is ε -joint differentially private. \square

3.3 Utility Analysis

In this section, we compare the weight of the matching produced by **PrivateMatching** with OPT. As an intermediate step, we first show that the resulting matching *paired with the prices* output by the algorithm forms an approximate matching equilibrium. We next show that any matching that can be paired with prices to form an approximate matching equilibrium must be an approximately max-weight matching.

The so-called “first welfare theorem” from general equilibrium theory guarantees that an exact (i.e., a $(0, 0, 0)$ -) matching equilibrium gives an exact maximum weight matching. Compared to this ideal, **PrivateMatching** loses welfare in three ways. First, a ρ fraction of bidders may end up unsatisfied. Second, the matched bidders are not necessarily matched to goods that maximize their utility given the prices, but only to goods that do so approximately (up to additive α). Finally, the auction sets aside part of the supply to handle over-allocation errors from the noisy counters, which may not end up being sold (say, if the counters are accurate or actually under-allocate). That is, we compute an equilibrium of a market with reduced supply, so our welfare guarantee requires that the supply s be significantly larger than the necessary reserved supply m .

The key performance metric is *how much* supply is needed to achieve a given welfare approximation in the final matching. (We show in Section 5 that such a tradeoff is unavoidable—no joint differentially private algorithm can achieve welfare $\text{OPT} - o(n)$ with supply $s = O(1)$.) So, let us consider which ranges of s are most interesting. On the one hand, we show that the problem is impossible if $s = 1$. On the other hand, the problem is trivial if $s \geq n$: every agent can be simultaneously matched to her favorite good with no coordination, which is trivially both optimal and private. Our algorithm will achieve positive results when $s \geq \text{polylog}(n)$.

Theorem 4. *Let $\alpha > 0$, and μ be the matching computed by **PrivateMatching** $(\alpha/3, \alpha/3, \varepsilon)$. Let OPT denote the weight of the optimal (max weight) matching. Then, if the supply satisfies*

$$s \geq \frac{16E' + 4}{\alpha} = O\left(\frac{1}{\alpha^3\varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\gamma}\right)\right),$$

and $n > s$, the matching μ has social welfare at least

$$\sum_{i=1}^n v_{i,\mu(i)} \geq \text{OPT} - \alpha n,$$

with probability $\geq 1 - \gamma$, where

$$E' = \frac{288\sqrt{2}}{\alpha^2\varepsilon} \left(\log\left(\frac{72n}{\alpha^2}\right)\right)^{5/2} \log\left(\frac{4k}{\gamma}\right).$$

Remark 1. *We note that our approximation guarantee here is additive. In Section 4, we show that if we are in the unweighted case where $v_{ij} \in \{0, 1\}$, the above guarantee can be made multiplicative. That is, we can find a matching μ such that $\sum_{i=1}^n v_{i,\mu(i)} \geq (1 - \alpha)\text{OPT}$, which is unusual in the context of differential privacy. Also, the second assumption $n > s$ is minimal, as the problem is trivially solvable for $s \geq n$.*

We first show a few lemmas.

Lemma 2. We call a bidder who wants to continue bidding unsatisfied; otherwise bidder i is satisfied. At termination of **PrivateMatching** $(\alpha, \rho, \varepsilon)$, all satisfied bidders i are matched to a good $\mu(i)$ such that:

$$v_{i, \mu(i)} - p_{\mu(i)} \geq \max_j (v_{i,j} - p_j) - \alpha$$

Proof. Fix any satisfied bidder i matched to $j^* = \mu(i)$. At the time that bidder i last bid on j^* , by construction, $v_{i,j^*} - p_{j^*} \geq \max_j (v_{i,j} - p_j)$. Since i remained matched to j^* , its price could only have increased by at most α , and the prices of other goods $j \neq j^*$ could only have increased. Hence, at completion of the algorithm,

$$v_{i, \mu(i)} - p_{\mu(i)} \geq \max_j (v_{i,j} - p_j) - \alpha$$

for all matched bidders i . □

Lemma 3. Assume all counters have error at most E throughout the run of **PrivateMatching** $(\alpha, \rho, \varepsilon)$. Then the number of bidders assigned to any good is at most s , and each overdemanded good clears except for at most β supply, where

$$\beta = 4E + 1 = O\left(\frac{1}{\alpha\rho\varepsilon} \cdot \text{polylog}\left(\frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}, k, n\right)\right).$$

Proof. Since the counter for each underdemanded good never exceeds $s - m$, we know that each underdemanded good is matched to no more than $s - m + E < s$ bidders.

Consider any counter c for an overdemanded good. Let t be a time step in counter c such that

$$c(nT) - c(t+1) \leq s - m < c(nT) - c(t).$$

Note that the bidders who bid after time t are the only bidders matched to this good at time nT . Let σ be the true bid stream for this good, so the total number of bidders allocated to this good at time nT is

$$\begin{aligned} c_\sigma(nT) - c_\sigma(t) &\leq c_\sigma(nT) - c_\sigma(t+1) + 1 \\ &\leq (c(nT) + E) - (c(t+1) - E) + 1 \\ &\leq s - m + 2E + 1 = s. \end{aligned}$$

Similarly, we can lower bound the number of bidders allocated to this good:

$$\begin{aligned} c_\sigma(nT) - c_\sigma(t) &= (c_\sigma(nT) - c(nT)) + (c(nT) - c(t)) + (c(t) - c_\sigma(t)) \\ &> s - m - 2E > s - 4E - 1. \end{aligned}$$

Therefore, every overdemanded good clears except for at most $\beta = 4E + 1$ supply, which gives the dependence

$$\beta = \frac{16\sqrt{2}}{\alpha\rho\varepsilon} \left(\log\left(\frac{6n}{\alpha\rho}\right)\right)^{5/2} \log\left(\frac{4k}{\gamma}\right) + 1 = O\left(\frac{1}{\alpha\rho\varepsilon} \cdot \text{polylog}\left(\frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}, k, n\right)\right).$$

□

Lemma 4. *Assume all counters have error at most E throughout the run of **PrivateMatching** $(\alpha, \rho, \varepsilon)$. Then at termination, all but a ρ fraction of bidders are satisfied, so long as $s \geq 8E+1$ and $n \geq 8E/\rho$.*

Proof. First, we claim that the total number of bids made over the course of the algorithm is bounded by $3n/\alpha$. We account separately for the under-demanded goods (those with price 0 at the end of the auction) and the over-demanded goods (those with positive price). For the under-demanded goods: since their prices remain 0 throughout the algorithm, their corresponding noisy counters never exceeded $(s - m)$. Since no bidder is ever unmatched after having been matched to an underdemanded good, the set of underdemanded goods can receive at most one bid from each agent, and hence together the under-demanded goods can receive at most n bids.

Next, we account for the over-demanded goods. Note the bidders matched to these goods are precisely the bidders who bid within $s - m$ ticks of the final counter reading. Since the counter has error bounded by E at each time step, this means at least $s - m - 2E$ bidders end up matched to each overdemanded good. Since no agent can be matched to more than one good there can be at most

$$\frac{n}{s - m - 2E}$$

over-demanded goods in total.

Likewise, we can account for the number of price increases per overdemanded good. Prices never rise above 1 (because any bidder would prefer to be unmatched than to be matched to a good with price larger than 1). Therefore, since prices are raised in increments of α , each overdemanded good can have its price incremented at most $1/\alpha$ times. Since there can be at most $(s - m + 2E)$ bids between each price update (again, corresponding to $s - m$ ticks of the counter), the total number of bids received by all of the over-demanded goods in total is at most

$$\frac{n}{s - m - 2E} \cdot \frac{1}{\alpha} \cdot (s - m + 2E).$$

Since each bid is either on an under or over-demanded good, we can upper bound the *total* number of bids B by

$$B \leq n + \frac{n}{\alpha} \left(\frac{s - m + 2E}{s - m - 2E} \right) = \frac{n}{\alpha} \left(\alpha + \frac{s - m + 2E}{s - m - 2E} \right).$$

We set the reserved supply to be $m = 2E + 1$ and by assumption, we have $s \geq 8E + 1$. Since we are only interested in cases where $\alpha < 1$, we conclude

$$B \leq n + \frac{n}{\alpha} \left(\frac{s - m + \alpha_2}{s - m - \alpha_2} \right) \leq \frac{3n}{\alpha}. \tag{1}$$

Now, consider the halting condition. There are two cases: either the algorithm halts early, or it does not. We claim that at termination, at most ρn bidders are unsatisfied. The algorithm halts if at any round of **CountUnsatisfied**, `counter(0)` (which counts the number of unsatisfied bidders) increases by less than $\rho n - 2E$. So if the algorithm halts, there must be at most $\rho n - 2E + 2E = \rho n$ unsatisfied bidders.

Otherwise, suppose the algorithm does not halt early. At the start of each round there must be at least $\rho n - 4E$ unsatisfied bidders. Not all of these bidders must bid during the **Propose** round since price increases while they are waiting to bid might cause them to no longer demand any item, but this only happens if bidders prefer to be unmatched at the current prices. Since prices only

increase, these bidders are satisfied for all future rounds. If the algorithm runs for R rounds and there are B true bids,

$$B \geq R(\rho n - 4E) - n.$$

Combined with our upper bound on the number of bids (Equation (1)) and our assumption $\rho n \geq 8E$, we can upper bound the number of rounds R :

$$R \leq \left(\frac{3n}{\alpha} + n\right) \cdot \left(\frac{1}{\rho n - 2E}\right) \leq \left(\frac{4n}{\alpha}\right) \left(\frac{2}{\rho n}\right) = \frac{8}{\alpha\rho} := T$$

Thus, running the algorithm for T rounds leads to all but ρn bidders satisfied. \square

Theorem 5. *With probability at least $1 - \gamma$, **PrivateMatching** $(\alpha, \rho, \varepsilon)$ computes an (α, β, ρ) -matching equilibrium, where*

$$\beta = 4E + 1 = O\left(\frac{1}{\alpha\rho\varepsilon} \cdot \text{polylog}\left(\frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}, k, n\right)\right)$$

so long as $s \geq 8E + 1$ and $n \geq 8E/\rho$.

Proof. By Theorem 1, counter(0) is $(\lambda_1, \gamma/2)$ -useful, and each of the k good counters is $(\lambda_2, \gamma/2)$ -useful, where

$$\lambda_1 = \frac{2\sqrt{2}}{\varepsilon'} (\log nT)^{5/2} \log\left(\frac{4}{\gamma}\right) \quad \text{and} \quad \lambda_2 = \frac{2\sqrt{2}}{\varepsilon'} (\log nT)^{5/2} \log\left(\frac{4k}{\gamma}\right).$$

Since we set $E = \lambda_2 > \lambda_1$, all counters are $(E, \gamma/2)$ -useful, and thus with probability at least $1 - \gamma$, all counters have error at most E . The theorem then follows by Lemmas 2 to 4. \square

Proof of Theorem 4. By Theorem 5, we know that **PrivateMatching** $(\alpha/3, \alpha/3, \varepsilon)$ calculates a matching μ that is an $(\alpha/3, \beta, \alpha/3)$ -approximate matching equilibrium with probability at least $1 - \gamma$, where $\beta = 4E' + 1$. Let p be the prices at the end of the algorithm, and S be the set of satisfied bidders. Let μ^* be the optimal matching achieving welfare $\sum_{i=1}^n v_{i, \mu^*(i)} = \text{OPT}$. We know that $|S| \geq (1 - \alpha/3)n$ and

$$\sum_{i \in S} (v_{i\mu(i)} - p_{\mu(i)}) \geq \sum_{i \in S} (v_{i\mu^*(i)} - p_{\mu^*(i)}) - \alpha|S|/3.$$

Let N_j^* and N_j be the number of goods of type j matched in matchings μ^* and μ respectively, and let G be the set of overdemanded goods at prices p .

Since each overdemanded good clears except for at most β supply, and since each of the n agents can be matched to at most 1 good, we know that $|G| \leq n/(s - \beta)$. Since the true supply in OPT is at most s , we also know $N_j^* - N_j \leq \beta$ for each overdemanded good j . Finally, by definition, underdemanded goods j have price $p_j = 0$. It follows that:

$$\begin{aligned} \sum_{i \in S} v_{i\mu^*(i)} - \sum_{i \in S} v_{i\mu(i)} &\leq \sum_{i \in S} p_{\mu^*(i)} - \sum_{i \in S} p_{\mu(i)} + \alpha|S|/3 \\ &= \sum_{j \in G} p_j (N_j^* - N_j) + \alpha|S|/3 \\ &\leq \sum_{j \in G} \beta + \alpha|S|/3 \leq \frac{n\beta}{s - \beta} + \alpha|S|/3. \end{aligned}$$

If $s \geq 4\beta/\alpha$, the first term $\beta n/(s - \beta)$ is at most $\alpha n/3$. Finally, since all but $\alpha n/3$ of the bidders are matched with goods in S , and their valuations are upper bounded by 1, we can conclude:

$$\sum_i v_{i\mu(i)} - \sum_i v_{i\mu^*(i)} \leq \alpha n/3 + \alpha|S|/3 + \alpha n/3 \leq \alpha n$$

Unpacking β from Theorem 5, we get the stated bound on supply s . □

4 Extensions

In this section, we extend our algorithm in two ways. First, we show how to compute approximately max-welfare allocations under general gross substitutes valuations. We also show how to modify and analyze the algorithm for computing max-weight matchings in the *unweighted* case when $v_{ij} \in \{0, 1\}$ to get *multiplicative* rather than additive approximation, which can be substantially better in the case when OPT is small. (More generally, the approximation depends on the minimum nonzero valuation.)

4.1 Allocations Under Gross Substitute Valuations

Let us first introduce some notation. Let $\Omega = 2^G$ denote the space of bundles (i.e., subsets of goods). Like previous sections, let k be number of types of goods, and let s be the supply of each type of good. Let d denote the *market size*—the total number of goods, including identical goods, so $d = ks$. (We remark that we assume each good has the same supply s only for convenience. In general, goods may have different supplies, if s denotes the *minimum* supply of any good. Hence, d should really be thought of as a parameter that is independent of s .) We assume each bidder has a valuation function on bundles, $v_i : \Omega \rightarrow [0, 1]$, and that this valuation satisfies the gross substitutes condition (Definition 1).

Like before, we simulate k ascending price auctions in rounds. Bidders maintain the bundle they are currently allocated to, and bid on one new good each round. For each good in a bidder's bundle, the bidder keeps track of the count of bids on that good when he was first assigned. When the current count ticks past the supply, the bidder knows that he has been outbid for that good.

The main subtlety is in how bidders decide which goods to bid on. Namely, each bidder considers goods in his bundle to be fixed in price (i.e., bidders ignore the price increment of at most α that might have occurred after winning the item). Goods outside of his bundle (even if identical to goods in his bundle) are evaluated at the true price, which are at most α higher. We call these prices the bidder's *effective* prices, so each bidder bids on an arbitrary good in his most-preferred bundle at the effective prices. The full algorithm is given in Algorithm 2.

Algorithm 2 PrivateAllocate($\alpha, \rho, \varepsilon$) (with Gross Substitute Valuations)

Input: Bidders' gross substitute valuations on the bundles $\{v_i : \Omega \rightarrow [0, 1]\}$

Initialize:

$$T = \frac{10}{\alpha\rho}, \quad \varepsilon' = \frac{\varepsilon}{2T}, \quad E = \frac{2\sqrt{2}}{\varepsilon'} (\log nT)^{5/2} \log\left(\frac{4k}{\gamma}\right) + 1, \quad m = 2E + 1,$$

$$\text{counter}(0) = \mathbf{Counter}(\varepsilon', nT),$$

$$\text{counter}(j) = \mathbf{Counter}(\varepsilon', nT), \quad p_j = c_j = 0 \quad \text{for every good type } j,$$

$$d_g = 0 \quad \text{for every good } g, \quad g(i) = \{\emptyset\} \quad \text{for every bidder } i$$

Call **Propose** T times; **Output:** Prices p and allocation g .

Propose:

for all bidders i **do**

for all goods $g \in g(i)$ **do**

if $c_{\text{type}(g)} - d_g \geq s - m$ **then**

 Remove $g(i) := g(i) \setminus g$

 Let p_0 be the original cost of $g(i)$.

 Let $\omega^* := \underset{\omega \supseteq g(i)}{\text{argmax}} v_i(\omega) - p(\omega \setminus g(i)) - p_0$.

if $v_i(\omega^*) - p(\omega \setminus g(i)) - p_0 \geq v_i(g(i)) - p_0$ **then**

 Let $j \in \omega^* \setminus g(i)$ arbitrary.

 Save $d_j := c_{\text{type}(j)}$

 Add $g(i) := g(i) \cup j$ and **Bid**(e_j)

else **Bid**(0)

CountUnsatisfied

Bid: On input bid vector \mathbf{b}

for all goods j **do**

 Feed \mathbf{b}_j to $\text{counter}(j)$.

 Update count $c_j := \text{counter}(j)$.

if c_j is a multiple of $s - m$ **then**

 Update $p_j := p_j + \alpha$.

CountUnsatisfied:

for all bidders i **do**

if i wants continue bidding **then**

 Feed 1 to $\text{counter}(0)$

else Feed 0 to $\text{counter}(0)$

 Halt if $\text{counter}(0)$ increases by less than $\rho d - 2E$

Theorem 6. **PrivateAllocate**($\alpha, \rho, \varepsilon$) satisfies ε -joint differential privacy.

Proof. Essentially the same proof as Theorem 5. □

Theorem 7. Let $0 < \alpha < n/d$, and g be the allocation computed by **PrivateAllocate**($\alpha/3, \alpha/3, \varepsilon$), and let OPT be the optimum max welfare. Then, if $d \geq n$ and

$$s \geq \frac{12E' + 3}{\alpha} = O\left(\frac{1}{\alpha^3 \varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\gamma}\right)\right),$$

the allocation g has social welfare at least

$$\sum_{i=1}^n v_i(g(i)) \geq \text{OPT} - \alpha d,$$

with probability at least $1 - \gamma$, where

$$E' = \frac{360\sqrt{2}}{\alpha^2 \varepsilon} \left(\log\left(\frac{90n}{\alpha^2}\right)\right)^{5/2} \log\left(\frac{4k}{\gamma}\right) + 1.$$

Remark 2. In comparison with Theorem 4, Theorem 7 requires a similar constraint on supply, but promises welfare only $\text{OPT} - \alpha d$ rather than $\text{OPT} - \alpha n$. Since $\text{OPT} \leq n$, this guarantee is only non-trivial for $\alpha \leq n/d$, and so the supply has a polynomial dependence on the total size of the market, d . In contrast, Theorem 4 guarantees good welfare when the supply has a logarithmic dependence on the total number of goods in the market.

However, we note that if bidders demand bundles of size at most b , then we can improve the above welfare bound to $\text{OPT} - \alpha nb$. Note that this is independent of the market size d , and strictly generalizes the matching case (where $b = 1$).

Similar to Definition 5, we define an *approximate allocation equilibrium* as an intermediate goal for showing our welfare guarantee.

Definition 6. A price vector $p \in [0, 1]^k$ and an assignment $g: [n] \rightarrow \Omega$ of bidders to goods is an (α, β, ρ) -approximate allocation equilibrium if

1. for all but ρd bidders, $v_i(g(i)) - p(g(i)) \geq \max_{\omega \in \Omega} v_i(\omega) - p(\omega) - \alpha |g(i)|$, and
2. the number of bidders assigned to any good is at most s , and
3. each overdemanded good clears except for at most β supply.

We prove the last two requirements first.

Lemma 5. Assume all counters have error at most E throughout the run of **PrivateAllocate** $(\alpha, \rho, \varepsilon)$. Then, the number of bidders assigned to any good is at most s , and each overdemanded good clears except for at most β supply, where

$$\beta = 4E + 1 = O\left(\frac{1}{\alpha\rho\varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}\right)\right).$$

Proof. Consider any good j . If it is underdemanded, the counter corresponding to j never rise above $s - m$. Hence, by our conditioning, at most $s - m + E < s$ bidders are assigned to j . If j is overdemanded, the same reasoning as in Theorem 5 shows that the number of bidders matched to j lies in the range $[s - m - 2E, s - m + 2E + 1]$. By the choice of m , the upper bound is at most s . Likewise, at least $s - m + E = s - (4E + 1)$ bidders are assigned to j . Setting $\beta = 4E + 1$ gives the desired bound. \square

Lemma 6. We call a bidder who wants to bid more unsatisfied; otherwise, a bidder is satisfied. At termination of **PrivateAllocate** $(\alpha, \rho, \varepsilon)$, all satisfied bidders are matched to a bundle $g(i)$ that is an $\alpha \cdot |g(i)|$ -most preferred bundle.

Proof. We first claim that a bidder's bundle $g(i)$ remains a subset of their most preferred bundle at the effective prices, i.e., with prices of goods in $g(i)$ set to their price at time of assignment, and all other goods taking current prices.

The claim follows by induction on the number of timesteps (ranging from 1 to nT). The base case is clear. Now, assume the claim holds up to time t . There are three possible cases:

1. If the price of a good outside $g(i)$ is increased, $g(i)$ remains part of a most-preferred bundle by the gross substitutes condition.

2. If the price of a good in $g(i)$ is increased, some goods may be removed from the bundle leading to a new bundle $g'(i)$. The only goods that experience an effective price increase lie outside of $g'(i)$, so $g'(i)$ remains a subset of a most-preferred bundle at the effective prices.
3. If a bidder adds to their bundle, $g(i)$ is a subset of the most-preferred bundle by definition.

Hence, a bidder becomes satisfied precisely when $g(i)$ is equal to the most-preferred bundle at the effective prices. The true price is at most α more than the effective price, so the bidder must have an $\alpha|g(i)|$ -most preferred bundle at the true prices. \square

Lemma 7. *Suppose all counters have error at most E throughout the run of **PrivateAllocate** $(\alpha, \rho, \varepsilon)$. Then at termination, all but ρd bidders are satisfied, so long as*

$$n \leq d \quad \text{and} \quad d \geq \frac{8E}{\rho} = \Omega\left(\frac{1}{\alpha\rho^2\varepsilon} \cdot \text{polylog}\left(n, \frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}\right)\right).$$

Proof. Note that if the unsatisfied bidders counter increases less than $\rho d - 2E$, then at most ρd bidders are actually unsatisfied. So, it remains to handle the case where the counter increases by at least $\rho d - 2E$ bidders each round.

In this case, at least $\rho d - 4E$ bidders are unsatisfied at the beginning of the round. They may not actually bid when their turn comes, because the prices may have changed. Let the number of bids among all bidders be B , and suppose we run for R rounds. We expect at least $\rho d - 4E$ bids per round, so $R(\rho d - 4E) - B$ is a lower bound on the number of times a bidder is unsatisfied, but fails to bid.

In the matching case, if a bidder is unsatisfied at the beginning of the round but fails to bid during their turn, this must be because the prices have risen too high. Since prices are monotonic increasing, such a bidder will never be unsatisfied again.

In contrast, the gross substitutes case is slightly more complex. Bidders who are unsatisfied at the beginning of a round and don't bid on their turn may later become unsatisfied again. Clearly, this happens only when the bidder loses at least one good after they decline to bid: if they don't lose any goods, then the prices can only increase after they decline to bid. Thus, they will have no inclination to bid in the future.

There are at most n cases of the bidder dropping out entirely. Thus, the number of times bidders report wanting to reenter the bidding is at least $R(\rho d - 4E) - n - B$. Since a bidder loses at least one good each time they reenter, the number of reentries is at most the number of bids B . Hence, the number of bids in R rounds is at least

$$B \geq \frac{R(\rho d - 4E) - n}{2}. \tag{2}$$

Now, let $s' = s - m = s - (2E + 1)$ be the effective supply and consider how many bids are possible. Each of the k types of goods will accept at most $s' + 2E = s + 1$ bids at each of $1/\alpha$ price levels, so there are at most $k(s + 1)/\alpha = (d + k)/\alpha$ possible bids.

Setting the left side of Equation (2) equal to $(d + k)/\alpha$, we find

$$R \leq \frac{1}{\alpha} \left(\frac{2(d + k) + \alpha n}{\rho d - 4E} \right) := T_0,$$

so taking $T \geq T_0$ suffices to ensure that the algorithm halts with no more than ρd bidders unsatisfied. Assuming $\rho d \geq 8E$ and $d \geq n$,

$$T_0 \leq \frac{10d}{\alpha\rho d} = \frac{10}{\alpha\rho} = T.$$

The requirement on n and d is then

$$d \geq \frac{8E}{\rho} = \Omega\left(\frac{1}{\alpha\rho^2\varepsilon} \cdot \text{polylog}\left(n, \frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}\right)\right) \quad \text{and} \quad n \leq d,$$

as desired. \square

Theorem 8. *With probability at least $1 - \gamma$, $\text{PrivateAllocate}(\alpha, \rho, \varepsilon)$ computes an (α, β, ρ) -approximate allocation equilibrium, where*

$$\beta = O\left(\frac{d}{\alpha\rho d\varepsilon} \cdot \text{polylog}\left(d, \frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}\right)\right),$$

so long as

$$d \geq \frac{8E}{\rho} = \Omega\left(\frac{1}{\alpha\rho^2\varepsilon} \cdot \text{polylog}\left(n, \frac{1}{\alpha}, \frac{1}{\rho}, \frac{1}{\gamma}\right)\right), \quad \text{and} \quad n \leq d.$$

Proof. Condition on the error for each counter being at most E throughout the run of the algorithm. By Theorem 1, this holds for any single counter with probability at least $1 - \gamma/2k$. By a union bound, this holds for all counters with probability at least $1 - \gamma$. The theorem follows by Lemmas 5 to 7. \square

Proof of Theorem 7. The proof follows Theorem 4 closely. By Theorem 8, (g, p) is a $(\alpha/3, \beta, \alpha/3)$ -approximate allocation equilibrium, where $\beta = 4E' + 1$. Then all but $\alpha d/3$ bidders are satisfied and get a bundle $g(i)$ that is $\alpha|g(i)|$ optimal; let this set of bidders be B . Note that $\sum_i |g(i)| \leq d$. Let g^* be any other allocation. Then,

$$\begin{aligned} \sum_{i \in B} v_i(g(i)) - p(g(i)) &\geq \sum_{i \in B} v_i(g^*(i)) - p(g^*(i)) - \frac{\alpha}{3}|g(i)| \\ \sum_{i \in B} v_i(g^*(i)) - v_i(g(i)) &\leq \sum_{i \in B} p(g^*(i)) - p(g(i)) + \alpha d/3 = \sum_{j \in G} p_j(N_j^* - N_j) + \alpha d/3 \end{aligned}$$

where the N_j is the number of good j sold in g and N_j^* is the number of good j sold in g^* . If $p_j > 0$, we know $N_j \geq s - \beta$, hence $N_j^* - N_j \leq \beta \leq \alpha s/3$. Also $p_j \leq 1$ for each good j , we have

$$\sum_{j \in G} p_j(N_j^* - N_j) \leq \sum_j p_j(N_j^* - N_j) \leq \alpha \sum_j s = \alpha d/3.$$

Furthermore, at most $\alpha d/3$ bidders are left unsatisfied in the end; these bidders contribute at most $\alpha d/3$ welfare to the optimal matching since valuations are bounded by 1. Putting it all together,

$$\sum_i v_i(g^*(i)) - v_i(g(i)) \leq \alpha d/3 + \alpha d/3 + \alpha d/3 = \alpha d.$$

The stated supply bound s follows directly from Theorem 8. \square

4.2 Multiplicative Approximation to Welfare

In certain situations, a close variant of **PrivateMatching** (Algorithm 1) can give a multiplicative welfare guarantee. In this section, we will work with matchings, and we will assume that value of the maximum weight matching OPT is known. (It is possible to privately estimate this quantity to high accuracy.) Our algorithm is exactly the same as **PrivateMatching**, except with a different halting condition: rather than count the number of unmatched bidders each round, count the number of bids per round. Once this count drops below a certain threshold, halt the algorithm.

More precisely, we use a function **CountBids** (Algorithm 3) in place of **CountUnsatisfied** in Algorithm 1.

Algorithm 3 Modified Halting Condition **CountBids**

CountBids:
for all bidders i **do**
 if $\mu(i) \neq \perp$ and $c_{\mu(i)} - d_i \geq s - m$ **then**
 Let $\mu(i) := \emptyset$
 if i bid this round **then**
 Feed 1 to counter(0).
 else Feed 0 to counter(0).
if counter(0) increases by less than $\frac{\alpha \text{OPT}}{2\lambda} - 2E$ **then**
 Halt; For each i with $\mu(i) = \emptyset$, let $\mu(i) = \perp$

Theorem 9. *Suppose bidders have valuations $\{v_{ij}\}$ over goods such that*

$$\min_{v_{ij} > 0} v_{ij} \geq \lambda.$$

Then Algorithm 1, with

$$T = \frac{24}{\alpha^2}$$

*rounds, using stopping condition **CountBids** (Algorithm 3) in place of **CountUnsatisfied**, and stopped once the total bid counter increases by less than*

$$\frac{\alpha \text{OPT}}{2\lambda} - 2E$$

bids in a round, satisfies ε -joint differential privacy and outputs a matching that has welfare at least $O((1 - \alpha/\lambda)) \text{OPT}$, so long as

$$s = \Omega\left(\frac{1}{\alpha^3 \varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\gamma}\right)\right) \quad \text{and} \quad \text{OPT} = \Omega\left(\frac{\lambda}{\alpha^3 \varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\gamma}\right)\right).$$

Proof. Privacy follows exactly like Theorem 3. We first show that at termination, all but $\alpha \text{OPT} / \lambda$ bidders are matched to an α -approximate favorite item. The analysis is very similar to Theorem 5. Note that every matched bidder is matched to an α -approximate favorite good, since it was an exactly favorite good at the time of matching, and the price increases by at most α . Thus, it remains to bound the number of unsatisfied bidders at termination.

Condition on all counters having error bounded by E at all time steps; by Theorem 1 and a union bound over counters, this happens with probability at least $1 - \gamma$. Like above, we write

$s' = s - m$ for the effective supply of each good. Let us first consider the case where the algorithm stops early. If the total bid counter changes by less than $\frac{\alpha \text{OPT}}{2\lambda} - 2E$, the true number of bids that round is at most

$$Q = \frac{\alpha \text{OPT}}{2\lambda}.$$

We will upper bound the number of unsatisfied bidders at the end of the round. Note that the number of unsatisfied bidders at the end of the round is the number of bidders who have been rejected in the current round. Suppose there are N goods that reject bidders during this round. The total count on these goods must be at least

$$(s' - 2E) \cdot N - Q$$

at the start of the round, since each counter will increase by at most $2E$ due to error, and there were at most Q bids this round. By our conditioning, there were at least

$$(s' - 2E) \cdot N - Q - 2EN$$

bidders matched at the beginning of the round. Since bidders are only matched when their valuation is at least λ , and the optimal weight matching is OPT , at most $\frac{\text{OPT}}{\lambda}$ bidders can be matched at any time. Hence,

$$N \leq \left(\frac{\text{OPT}}{\lambda} + Q \right) \cdot \frac{1}{s' - 4E}.$$

Then, the total number of bidders rejected this round is at most $2EN + Q$. Simplifying,

$$\begin{aligned} 2EN + Q &\leq \frac{2E}{s' - 4E} \cdot \left(\frac{\text{OPT}}{\lambda} + Q \right) + Q \\ &\leq \left(\frac{6E}{s' - 4E} \right) \left(\frac{\text{OPT}}{\lambda} \right) + \frac{\alpha \text{OPT}}{2\lambda}. \end{aligned}$$

To make the first term at most $\frac{\alpha \text{OPT}}{2\lambda}$, it suffices to take

$$\begin{aligned} \frac{6E}{s' - 4E} &\leq \frac{\alpha}{2} \\ s' &\geq \frac{12E}{\alpha} + 4E \\ s &\geq \frac{12E}{\alpha} + 6E + 1, \end{aligned}$$

or $s \geq 18E/\alpha$. In this case, the algorithm terminates with at most $\frac{\alpha \text{OPT}}{\lambda}$ unsatisfied bidders, as desired.

On the other hand, suppose the algorithm does not terminate early, the bid count increasing by at least $Q - 2E$ every round. By our conditioning, this means there are at least $Q - 4E$ bids each round; let us bound the number of possible bids.

Since bidders only bid if they have valuation greater than λ for a good, and since the maximum weight matching has total valuation OPT , at most OPT/λ bidders can be matched. Goods are either underdemanded or overdemanded: they either have final price 0, or positive final price.

There are at most OPT/λ true bids on the goods of the first type; this is because bidders are never rejected from these goods. Like before, write $s' = s - m$. Each counter of a overdemanded good shows s' people matched, so at least $s' - 2E$ bidders end up matched. Thus, there are at most

$$\frac{\text{OPT}}{\lambda(s' - 2E)}$$

overdemanded goods. Each such good takes at most $s' + 2E$ bids at each of $1/\alpha$ price levels. Putting these two estimates together, the total number of bids B is upper bounded by

$$B \leq \frac{\text{OPT}}{\lambda} \cdot \left(1 + \frac{s' + 2E}{s' - 2E}\right) \leq \frac{6 \text{OPT}}{\lambda\alpha}$$

if $s' \geq 4E$, which holds since we are already assuming $s' \geq 4E + \frac{12E}{\alpha}$. Hence, we know the number of bids is at most

$$\begin{aligned} T \cdot (Q - 4E) &\leq B \leq \frac{6 \text{OPT}}{\lambda\alpha} \\ T &\leq \frac{6 \text{OPT}}{\lambda} \cdot \left(\frac{2\lambda}{\alpha \text{OPT} - 8\lambda E}\right). \end{aligned}$$

Assuming $\alpha \text{OPT} \geq 16\lambda E$, we find $T \leq 24/\alpha^2$.

With this choice of T , the supply requirement is

$$s \geq \frac{18E}{\alpha} = \Omega\left(\frac{1}{\alpha^3 \varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\gamma}\right)\right). \quad (3)$$

Likewise, the requirement on OPT is

$$\text{OPT} \geq \frac{16\lambda E}{\alpha} = \Omega\left(\frac{\lambda}{\alpha^3 \varepsilon} \cdot \text{polylog}\left(n, k, \frac{1}{\alpha}, \frac{1}{\gamma}\right)\right).$$

Now, we can follow the analysis from Theorem 4 to bound the welfare. Suppose the algorithm produces a matching μ , and consider any other matching μ^* . For each bidder who is matched to an α -approximate favorite good,

$$v_{i\mu(i)} - p_{\mu(i)} \geq v_{i\mu^*(i)} - p_{\mu^*(i)} - \alpha.$$

Each such bidder is matched to a good with value at least λ , so there are at most OPT/λ such bidders. Summing over these bidders (call them S),

$$\sum_{i \in S} v_{i\mu(i)} - p_{\mu(i)} \geq \sum_{i \in S} v_{i\mu^*(i)} - p_{\mu^*(i)} - \frac{\alpha \text{OPT}}{\lambda}.$$

Letting N_j, N_j^* be the number of goods of type j matched in μ, μ^* and rearranging,

$$\sum_{i \in S} v_{i\mu^*(i)} - v_{i\mu(i)} \leq \sum_{j \in S} p_j (N_j^* - N_j) + \frac{\alpha \text{OPT}}{\lambda}.$$

Exactly the same as Theorem 4, each overdemanded good ($p_j > 0$) clears except for at most $\beta = 4E + 1$ supply. Since at most $\frac{\text{OPT}}{\lambda}$ bidders can be matched, the number of goods with $p_j > 0$ is at most

$$\frac{\text{OPT}}{\lambda(s - \beta)}.$$

Like before, $N_j^* - N_j \leq \beta$. Since there are at most $\alpha \text{OPT} / \lambda$ bidders not in S and each has valuation in $[0, 1]$, when summing over all bidders,

$$\sum_i v_{i\mu^*(i)} - v_{i\mu(i)} \leq \frac{\text{OPT} \beta}{\lambda(s - \beta)} + \frac{\alpha \text{OPT}}{\lambda} + \frac{\alpha \text{OPT}}{\lambda}.$$

The first term is at most $\alpha \text{OPT} / \lambda$ for $s \geq \beta(1 + 1/\alpha)$, when the algorithm calculates a matching with weight $O((1 - \alpha/\lambda) \text{OPT})$. Since $\beta = 4E + 1$, this reduces to the supply constraint Equation (3). \square

Remark 3. For a comparison with Theorem 4 and **PrivateMatching**, consider the “unweighted” case where bidders have valuations in $\{0, 1\}$ (i.e., $\lambda = 1$). Note that both **PrivateMatching** and the multiplicative version require the same lower bound on supply. Ignoring log factors, **PrivateMatching** requires $n = \tilde{\Omega}(1/\alpha^3 \varepsilon)$ for an additive αn approximation, while Theorem 9 shows $\text{OPT} = \tilde{\Omega}(1/\alpha^3 \varepsilon)$ is necessary for a multiplicative α , hence additive αOPT , approximation. Hence, Theorem 9 gives a stronger guarantee if $\text{OPT} = \tilde{o}(n)$ in the unweighted case, ignoring log factors.

5 Lower Bounds

Our lower bounds all reduce to a basic database reconstruction lower bound for differential privacy.

Theorem 10. Let mechanism $\mathcal{M}: \{0, 1\}^n \rightarrow \{0, 1\}^n$ be (ε, δ) -differentially private, and suppose that for all database D , with probability at least $1 - \beta$, $\|\mathcal{M}(D) - D\|_1 \leq \alpha n$. Then,

$$\alpha \geq 1 - \frac{e^\varepsilon + \delta}{(1 + e^\varepsilon)(1 - \beta)} := c(\varepsilon, \delta, \beta).$$

In other words, no (ε, δ) -private mechanism can reconstruct more than a fixed constant fraction of its input database. For $\varepsilon, \delta, \beta$ small, $c(\varepsilon, \delta, \beta) \sim 1/2$. Informally, this theorem states that a private reconstruction mechanism can’t do much better than guessing a random database. Note that this holds even if the adversary doesn’t know which fraction was correctly reconstructed. This theorem is folklore; a proof can be found in Appendix B.

Our lower bounds will all be proved using the following pattern:

- First, we describe how to convert a database $D \in \{0, 1\}^n$ to a market, by specifying the bidders, the goods, and the valuations $v_{ij} \in [0, 1]$ on goods.
- Next, we analyze how these valuations change when a single bit in D is changed. This will control how private the matching algorithm is with respect to the original database, when applied to this market.
- Finally, we show how to output a database guess \hat{D} from the matching produced by the private matching algorithm.

This composition of three steps will be a private function from $\{0, 1\}^n \rightarrow \{0, 1\}^n$, so we can apply Theorem 10 to lower bound the error. This will in turn imply a lower bound on the error of the matching algorithm.

5.1 Lower Bounds for Standard Differential Privacy

Note that Algorithm 1 produces market clearing prices under standard differential privacy. We will first show that this is not possible if each good has unit supply. Recall that prices correspond to an (α, β, ρ) -approximate matching equilibrium if all but ρ bidders can be allocated to a good such that their utility (valuation less price) is within α of their favorite good (Definition 5). We will ignore the β parameter, which controls how many goods are left unsold.

Theorem 11. *Let n bidders have valuations $v_{ij} \in [0, 1]$ for n goods. Suppose that mechanism \mathcal{M} is (ε, δ) -differentially private, and calculates prices corresponding to an (α, β, ρ) -approximate matching equilibrium for $\alpha < 1/2$ and some β with probability $1 - \gamma$. Then,*

$$\rho \geq \frac{1}{2}c(2\varepsilon, \delta(1 + e^\varepsilon), \gamma).$$

Note that this is independent of α .

Proof. Let $D \in \{0, 1\}^{n/2}$ be a private database and construct the following market. For each bit i , we construct the following gadget, consisting of two goods $0_i, 1_i$ and two bidders, b_i, \bar{b}_i . Both bidders have valuation D_i for good 1_i , $1 - D_i$ for good 0_i , and 0 for the other goods. Evidently, there are n bidders and n goods.

Note that changing a bit i in D changes the valuation of two bidders in the market: b_i and \bar{b}_i . Therefore, mechanism \mathcal{M} is $(2\varepsilon, \delta(1 + e^\varepsilon))$ -differentially private with respect to D . Let the prices be p_{0i}, p_{1i} . To guess the database \hat{D} , we let $\hat{D}_i = 1$ if $p_{1i} > 1/2$, otherwise $\hat{D}_i = 0$.

By assumption, \mathcal{M} produces prices corresponding to an (α, β, ρ) -approximate matching equilibrium, with probability $1 - \gamma$. We do not have access to the matching, but we know the prices must correspond to *some* matching μ . Then, for all but ρn gadgets, μ matches both bidders to their α -approximate favorite good, and both goods are matched to bidders who receive α -approximate favorite goods.

Consider such a gadget i . We will show that for this gadget, exactly one of p_{0i} or p_{1i} is greater than $1/2$, and this expensive good corresponds to bit D_i . Consider one of the bidders in this gadget, and suppose he prefers good g_+ with price p_+ , while he received good g_- with price p_- . Since he receives an α -approximate favorite good,

$$(1 - p_+) - (0 - p_-) \leq \alpha, \quad \text{so} \quad p_+ - p_- \geq 1 - \alpha > 1/2.$$

So $p_+ > 1/2$ and $p_- < 1/2$. Note that good g_+ is in the gadget, while good g_- may not be. So, one of the goods in the gadget has price strictly greater than $1/2$. The other good in the gadget is an α -approximate favorite good for some bidder. All bidders have valuation 0 for the good, hence its price must be strictly less than $1/2$.

Thus, the reconstruction procedure will correctly produce bit for each such gadget, and so will miss at most ρn bits with probability at least $1 - \gamma$. The combined reconstruction algorithm is a map from $\{0, 1\}^{n/2} \rightarrow \{0, 1\}^{n/2}$, and $(2\varepsilon, \delta(1 + e^\varepsilon))$ -differentially private. By Theorem 10,

$$2\rho \geq c(2\varepsilon, \delta(1 + e^\varepsilon), \gamma).$$

□

5.2 Separation Between Standard and Joint Differential Privacy

While we can compute an approximate maximum welfare matching under joint privacy when the supply of each good is large (Theorem 5), this is not possible under standard differential privacy even with infinite supply. (In fact, it is not possible with finite supply either.)

Theorem 12. *Let n bidders have valuations $v_{ij} \in \{0, 1\}$ for 2 goods with infinite supply. Suppose that mechanism \mathcal{M} is (ε, δ) -differentially private, and computes a matching with welfare at least $\text{OPT} - \alpha n$ with probability $1 - \gamma$. Then,*

$$\alpha \geq c(\varepsilon, \delta, \gamma).$$

Proof. Let $D \in \{0, 1\}^n$. We assume two goods, $\mathbf{0}$ and $\mathbf{1}$. We have one bidder b_i for each bit $i \in [n]$, who has valuation D_i for $\mathbf{1}$, and valuation $1 - D_i$ for $\mathbf{0}$. Since changing a bit changes a single bidder's valuation, applying \mathcal{M} to this market is (ε, δ) -private with respect to D . To guess the database \widehat{D} , we let \widehat{D}_i be 0 if b_i is matched to $\mathbf{0}$, 1 if b_i is matched to $\mathbf{1}$, and arbitrary if b_i isn't matched.

Note that the maximum welfare matching assigns each b_i the good corresponding to D_i , and achieves social welfare $\text{OPT} = n$. If \mathcal{M} computes a matching with welfare $\text{OPT} - \alpha n$, it must give all but an α fraction of bidders b_i the good corresponding to D_i . So, the reconstructed database will miss at most αn bits with probability $1 - \gamma$, and by Theorem 10,

$$\alpha \geq c(\varepsilon, \delta, \gamma).$$

□

Note that this gives a separation: under joint differential privacy, Algorithm 1 can release a matching with welfare $\text{OPT} - \alpha n$ for any α , provided supply s is large enough (by Theorem 4). This is not possible under standard differential privacy, even with *infinite* supply.

5.3 Lower Bounds for Joint Differential Privacy

Finally, we show that a large supply assumption is necessary in order to compute an additive α maximum welfare matching under joint differential privacy.

Theorem 13. *Let n bidders have valuations $v_{ij} \in [0, 1]$ for k types of goods with supply s each. Suppose mechanism \mathcal{M} is (ε, δ) -joint differentially private for $\varepsilon, \delta < 0.1$, and calculates a matching with welfare at least $\text{OPT} - \alpha n$ with probability $1 - \gamma$ for $\gamma < 0.01$, and all n, k, s . Then, $s = \Omega(\sqrt{1/\alpha})$.*

Proof. Let $k = n/(s + 1)$. Given a private database $D \in \{0, 1\}^k$, construct the following market. For each bit i , we construct a gadget with two goods $\mathbf{0}_i, \mathbf{1}_i$, each with supply s . Each gadget has a distinguished bidder b_i and s identical bidders, all labeled \overline{b}_i . Let bidder b_i , who we call the *real bidder*, have valuation D_i for $\mathbf{1}_i$, and $1 - D_i$ for $\mathbf{0}_i$. Bidders \overline{b}_i , which we call the *spy bidders*, all have the same valuation: $\eta = \frac{1}{4s}$ for $\mathbf{0}_i$ or $\mathbf{1}_i$ drawn at random, and 0 for all other goods (in and out of the gadget). We say a bidder *prefers* a good if they have positive valuation for the good.

Note that changing a bit in D changes a single bidder's valuation. Also note that the spy bidders' valuations do not depend on D . Hence, by joint differential privacy of \mathcal{M} , the function that maps

the above market through \mathcal{M} to the allocation of just the spy bidders is (ε, δ) -differentially private with respect to an entry change in D .

We will describe how to guess \widehat{D} based on just the spy bidders' joint view, i.e., the goods they are assigned. This reconstruction procedure will then be (ε, δ) -differentially private, and we can apply Theorem 10 to lower bound the error of \mathcal{M} . For every bit $i \in [k]$, let \widehat{D}_i be 1 if the spy bidders in gadget i are all assigned to $\mathbf{0}_i$, 0 if the spy bidders in gadget i are all assigned to $\mathbf{1}_i$, and uniformly random otherwise.

We'll say that a gadget *agrees* if the spy bidders and real bidder prefer the same good. Gadgets that don't agree, *disagree*. Let w be the number of gadgets that agree. By construction, gadgets agree independently at random with probability $1/2$. Hence, Hoeffding's inequality gives

$$\Pr \left[\left| w - \frac{k}{2} \right| \leq \lambda k \right] \geq 1 - 2 \exp(-2\lambda^2 k)$$

for some λ to be chosen later; condition on this event. With probability at least $1 - \gamma$, mechanism \mathcal{M} computes a matching with welfare at least $\text{OPT} - \alpha n$; condition on this event as well. Note that the optimum welfare is $1 + (s - 1)\eta$ for gadgets that agree, and $1 + s\eta$ for gadgets that disagree, hence $\text{OPT} = w(1 + (s - 1)\eta) + (k - w)(1 + s\eta)$ in total.

For each gadget, there are several possible allocations. Intuitively, an assignment gives social welfare, but may also lead to a bit being reconstructed. Let $RB(\mu) = \|D - \widehat{D}\|_1$ be the error of the reconstruction when the matching is μ . We'll argue that any matching μ with nearly optimal social welfare must result in large expected reconstruction $\mathbb{E}[RB(\mu)]$. Note that

$$\mathbb{E}[RB(\mu)] = \sum_{i \in [k]} \Pr \left[D_i = \widehat{D}_i \right],$$

so we argue gadget by gadget.

First, suppose the gadget i agrees. The matching μ can give the preferred good to the bidder, the spies, or neither. If the preferred good goes to the bidder, this gives at most $1 + (s - 1)\eta$ social welfare. Not all the spies get the same good, so

$$\Pr \left[D_i = \widehat{D}_i \right] = \frac{1}{2}.$$

If the preferred good goes to the spies, then this contributes $s\eta$ to social welfare, and

$$\Pr \left[D_i = \widehat{D}_i \right] = 0.$$

Note that it doesn't matter whether the bidder is assigned in μ , since the social welfare is unchanged, and the reconstruction algorithm doesn't have access to the bidder's allocation. There are other possible allocations, but they are dominated by these two choices (they get less social welfare for higher reconstruction probability).

Now, suppose gadget i disagrees. There are several possible allocations. First, both the bidder and the spies may get their favorite good. This leads to $1 + s\eta$ welfare, and

$$\Pr \left[D_i = \widehat{D}_i \right] = 1.$$

Second, the bidder may be assigned their favorite good, and at most $s - 1$ spies may be assigned their favorite good. This leads to $1 + (s - 1)\eta$ welfare, with

$$\Pr [D_i = \widehat{D}_i] = \frac{1}{2}.$$

Again, there are other possible allocations, but they lead to less social welfare or higher reconstruction probability. We call these four allocations *optimal*.

Let a_1, a_2 be the fractions of agreeing gadgets with the two optimal agreeing allocations, and d_1, d_2 be the fractions of disagreeing gadgets with the two optimal disagreeing allocations. Let t be the fraction of agreeing pairs. The following linear program minimizes $(1/k)\mathbb{E}[RB(\mu)]$ over all matchings μ achieving an α -approximate maximum welfare matching, for supply s .

$$\begin{aligned} LP_s := \quad & \text{minimize: } \frac{1}{2}a_1 + d_1 + \frac{1}{2}d_2 \\ & \text{such that: } a_1 + a_2 \leq t \\ & d_1 + d_2 \leq 1 - t \\ & \frac{1}{2} - \lambda \leq t \leq \frac{1}{2} + \lambda \\ & (1 + (s - 1)\eta)a_1 + s\eta a_2 + (1 + s\eta)d_1 + (1 + (s - 1)\eta)d_2 \\ & \geq t(1 + (s - 1)\eta) + (1 - t)(1 + s\eta) - \frac{\alpha n}{k} \end{aligned}$$

The last constraint is the welfare requirement, the second to last constraint is from conditioning on the number of agreeing gadgets, and the objective is $(1/k)\mathbb{E}[RB(\mu)]$.

Plugging in $\eta = \frac{1}{4s}$, $\lambda = 1/128$, $\alpha = \frac{k}{16ns}$ and solving, we find

$$(a_1, a_2, d_1, d_2, t) = \left(\frac{65}{128}, 0, \frac{31}{128}, \frac{1}{4}, \frac{65}{128} \right)$$

is a feasible solution for all s , with objective $\alpha' = 159/256$. To show that this is optimal, consider the dual problem:

$$\begin{aligned} DUAL_s := \quad & \text{maximize: } -\rho_2 + \left(\frac{1}{2} - \lambda\right)\rho_3 - \left(\frac{1}{2} + \lambda\right)\rho_4 + \left(1 + s\eta - \frac{\alpha n}{k}\right)\rho_5 \\ & \text{such that: } -\rho_1 + (1 + (s - 1)\eta)\rho_5 \leq \frac{1}{2} \\ & -\rho_1 + s\eta\rho_5 \leq 0 \\ & -\rho_2 + (1 + s\eta)\rho_5 \leq 1 \\ & -\rho_2 + (1 + (s - 1)\eta)\rho_5 \leq \frac{1}{2} \\ & \rho_1 - \rho_2 + \rho_3 - \rho_4 + \eta\rho_5 \leq 0 \end{aligned}$$

We can directly verify that

$$(\rho_1, \rho_2, \rho_3, \rho_4, \rho_5) = \left(\frac{5}{2}s - 1, \frac{5}{2}s - 1, 0, \frac{1}{2}, 2s \right)$$

is a dual feasible solution with objective $\alpha' = 159/256$.

We know that \mathcal{M} calculates an additive α -approximate maximum welfare matching. While the allocations to each gadget may not be an optimal allocation, suboptimal allocations all have less social welfare and larger RB . So, we know the objective of LP_m is a lower bound for $RB(\mathcal{M})$.

Thus, $\mathbb{E}[RB(\mathcal{M})] \geq k\alpha'$ for any supply s . Since RB is the sum of k independent, 0/1 random variables, another Hoeffding bound yields

$$\Pr [RB(\mathcal{M})/k \geq \alpha' - \lambda'] \geq 1 - 2 \exp(-2\lambda'^2 k).$$

Set $\lambda' = 1/256$, and condition on this event. Taking everything together, any matching mechanism \mathcal{M} which finds a matching with weight at least $\text{OPT} - \alpha n$ failing with at most γ probability gives an (ε, δ) -private mechanism taking database D to \widehat{D} , such that

$$\frac{1}{k} \cdot \|D - \widehat{D}\|_1 \geq \alpha' - \lambda' = 79/128.$$

with probability at least $1 - \gamma - 2 \exp(-2\lambda^2 k) - 2 \exp(-2\lambda'^2 k)$.

For $\varepsilon, \delta < 0.1$ and $\gamma < 0.01$, this contradicts Theorem 10 for large k . Note that the failure probability and accuracy do not depend directly on s , since $\lambda, \lambda', \alpha'$ are constants. Hence,

$$\alpha \gg \frac{k}{16ns} = \frac{1}{16s(s+1)}$$

uniformly for all s , and $s = \Omega(\sqrt{1/\alpha})$ as desired. \square

6 Conclusion and Open Problems

In this paper we have given algorithms to accurately solve the private allocation problem when bidders have gross substitute valuations. Our results are qualitatively tight: it is not possible to strengthen our solution concept to standard differential privacy (from joint differential privacy), nor is it possible to solve even max-matching problems to non-trivial accuracy under joint differential privacy with constant supply. Moreover, our approach cannot be pushed any further: our algorithm fundamentally relies on computing Walrasian equilibrium prices for the underlying market, and such prices are not guaranteed to exist for valuation functions beyond the gross substitutes class. This does not mean that the allocation problem cannot be solved for more general valuation functions, only that fundamentally new ideas would be needed; we leave this extremely interesting problem open.

Taken together with [Kearns et al. \[2014\]](#), our work provides compelling evidence that substantially more is possible under the relaxation of *joint* differential privacy, as compared to the standard notion of differential privacy. For both the allocation problem studied here, and the equilibrium computation problem studied in [Kearns et al. \[2014\]](#), non-trivial results are impossible under differential privacy, while strong results can be derived under joint-differential privacy. Characterizing the power of joint differential privacy, as compared to the standard differential privacy, continues to be a fascinating direction for future work.

More specifically, in this paper we achieve joint differential privacy via the *billboard lemma*: we show that the allocation given to each player can be derived as a deterministic function only of 1) a differentially private message revealed to all players, and 2) their own private data. However, this isn't necessarily the only way to achieve joint differential privacy. How much further does the power of joint differential privacy extend beyond the billboard model?

Acknowledgments

The authors would like to thank Cynthia Dwork, Sudipto Guha, Moritz Hardt, Sanjeev Khanna, Scott Kominers, Mallesh Pai, David Parkes, Adam Smith, and Kunal Talwar for helpful discussions. In particular, we would like to thank Scott Kominers for suggesting the connection to Kelso and Crawford at an early stage of this work, and Adam Smith for discussions on the “billboard model” of privacy.

References

- Avrim Blum, Katrina Ligett, and Aaron Roth. [A learning theory approach to noninteractive database privacy](#). *Journal of the ACM*, 60(2):12, 2013.
- T.-H. Hubert Chan, Elaine Shi, and Dawn Song. [Private and continual release of statistics](#). *ACM Transactions on Information and System Security*, 14(3):26, 2011.
- Irit Dinur and Kobbi Nissim. [Revealing information while preserving privacy](#). In *ACM SIGACT–SIGMOD–SIGART Symposium on Principles of Database Systems (PODS), San Diego, California*, pages 202–210, 2003.
- Cynthia Dwork and Aaron Roth. [The algorithmic foundations of differential privacy](#). 2013.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. [Calibrating noise to sensitivity in private data analysis](#). In *IACR Theory of Cryptography Conference (TCC), New York, New York*, 2006.
- Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N. Rothblum. [Differential privacy under continual observation](#). In *ACM SIGACT Symposium on Theory of Computing (STOC), Cambridge, Massachusetts*, pages 715–724, 2010a.
- Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. [Boosting and differential privacy](#). In *IEEE Symposium on Foundations of Computer Science (FOCS), Las Vegas, Nevada*, pages 51–60, 2010b.
- Cynthia Dwork, Moni Naor, and Salil Vadhan. [The privacy of the analyst and the power of the state](#). In *IEEE Symposium on Foundations of Computer Science (FOCS), New Brunswick, New Jersey*, pages 400–409, 2012.
- Faruk Gul and Ennio Stacchetti. [Walrasian equilibrium with gross substitutes](#). *Journal of Economic Theory*, 87(1):95–124, 1999.
- Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. [Differentially private combinatorial optimization](#). In *ACM–SIAM Symposium on Discrete Algorithms (SODA), Austin, Texas*, pages 1106–1125, 2010.
- Moritz Hardt and Guy N. Rothblum. [A multiplicative weights mechanism for privacy-preserving data analysis](#). In *IEEE Symposium on Foundations of Computer Science (FOCS), Las Vegas, Nevada*, pages 61–70, 2010.

Justin Hsu, Aaron Roth, and Jonathan Ullman. [Differential privacy for the analyst via private equilibrium computation](#). In *ACM SIGACT Symposium on Theory of Computing (STOC)*, Palo Alto, California, pages 341–350, 2013.

Michael Kearns, Malleesh Pai, Aaron Roth, and Jonathan Ullman. [Private equilibrium release, large games, and no-regret learning](#). In *ACM SIGACT Innovations in Theoretical Computer Science (ITCS)*, Princeton, New Jersey, 2014.

Alexander Kelso and Vincent Crawford. [Job matching, coalition formation, and gross substitutes](#). *Econometrica*, 50(6):22, 1982.

Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. [Smooth sensitivity and sampling in private data analysis](#). In *ACM SIGACT Symposium on Theory of Computing (STOC)*, San Diego, Illinois, pages 75–84, 2007.

Malleesh Pai and Aaron Roth. [Privacy and mechanism design](#). *ACM SIGecom Exchanges*, 2013.

Ryan Rogers and Aaron Roth. [Asymptotically truthful equilibrium selection in large congestion games](#). 2013.

A Privacy Analysis for Counters

Chan et al. [2011] show that $\mathbf{Counter}(\varepsilon, T)$ is ε -differentially private with respect to single changes in the input stream, when the stream is generated non-adaptively. For our application, we require privacy to hold for a large number of streams whose joint-sensitivity can nevertheless be bounded, and whose entries can be chosen adaptively. To show that $\mathbf{Counter}$ is also private in this setting (when ε is set appropriately), we first introduce some differential privacy notions.

We will make use of a basic differentially private mechanism originally due to Dwork et al. [2006].

Theorem 14 (Dwork et al. [2006]). *For a function $f : \mathcal{D} \rightarrow \mathbb{R}$, let*

$$\Delta_1 = \max_{D, D' \in \mathcal{D}} \frac{|f(D) - f(D')|}{|\{i : D_i \neq D'_i\}|}$$

denote the ℓ_1 sensitivity of f . Then the Laplace Mechanism which on input D outputs $f(D) + \text{Lap}(\Delta_1/\varepsilon)$ is ε -differentially private. Here, $\text{Lap}(x)$ denotes a random variable drawn from the Laplace distribution with variance $2x^2$.

A.1 Composition

An important property of differential privacy is that it degrades gracefully when private mechanisms are composed together, even adaptively. We recall the definition of an adaptive composition experiment due to Dwork et al. [2010b].

Definition 7 (Adaptive composition experiment).

- Fix a bit $b \in \{0, 1\}$ and a class of mechanisms \mathcal{M} .

- For $t = 1 \dots T$:
 - The adversary selects two databases $D^{t,0}, D^{t,1}$ and a mechanism $\mathcal{M}_t \in \mathcal{M}$.
 - The adversary receives $y_t = \mathcal{M}_t(D^{t,b})$

The “output” of an adaptive composition experiment is the view of the adversary over the course of the experiment. The experiment is said to be ε -differentially private if

$$\max_{S \subseteq \mathcal{R}} \frac{\Pr[V^0 \in S]}{\Pr[V^1 \in S]} \leq \exp(\varepsilon),$$

where V^0 is the view of the adversary with $b = 0$, V^1 is the view of the adversary with $b = 1$, and \mathcal{R} is the range of outputs.

Any algorithm that can be described as an instance of this adaptive composition experiment (for an appropriately defined adversary) is said to be an instance of the class of mechanisms \mathcal{M} under *adaptive T -fold composition*. We now state a straightforward consequence of a composition theorem of [Dwork et al. \[2010b\]](#).

Lemma 8 ([Dwork et al. \[2010b\]](#)). *Let $\Delta_1 \geq 0$. The class of $\frac{\varepsilon}{\Delta_1}$ -private mechanisms satisfies ε -differential privacy under adaptive composition, if the adversary always selects databases satisfying*

$$\sum_{t=1}^T |D^{t,0} - D^{t,1}| \leq \Delta_1.$$

In other words, the privacy parameter of each mechanism should be calibrated for the total distance between the databases, over the whole composition (the ℓ_1 sensitivity).

A.2 Binary Mechanism

We reproduce Binary mechanism here in order to refer to its internal workings in our privacy proof.

First, it is worth explaining the intuition of the **Counter**. Given a bit stream $\sigma: [T] \rightarrow \{0, 1\}$, the algorithm releases the counts $\sum_{i=1}^t \sigma(i)$ for each t by maintaining a set of partial sums $\Sigma[i, j] := \sum_{t=i}^j \sigma(t)$. More precisely, each partial sum has the form $\Sigma[2^i + 1, 2^i + 2^{i-1}]$, corresponding to powers of 2.

In this way, we can calculate the count $\sum_{i=1}^t \sigma(i)$ by summing at most $\log t$ partial sums: let $i_1 < i_2 \dots < i_m$ be the indices of non-zero bits in the binary representation of t , so that

$$\sum_{i=1}^t \sigma(i) = \Sigma[1, 2^{i_m}] + \Sigma[2^{i_m} + 1, 2^{i_m} + 2^{i_m-1}] + \dots + \Sigma[t - 2^{i_1} + 1, t].$$

Therefore, we can view the algorithm as releasing partial sums of different ranges at each time step t and computing the counts is simply a post-processing of the partial sums. The core algorithm is presented in [Algorithm 4](#).

Algorithm 4 Counter(ε, T)

Input: A stream $\sigma \in \{0, 1\}^T$

Output: $B(t)$ as estimate for $\sum_{i=1}^t \sigma(i)$ for each time $t \in [T]$

for all $t \in [T]$ **do**

Express $t = \sum_{j=0}^{\log t} 2^j \text{Bin}_j(t)$.

Let $i \leftarrow \min_j \{\text{Bin}_j(t) \neq 0\}$

$a_i \leftarrow \sum_{j < i} a_j + \sigma(t)$, ($a_i = \sum [t - 2^i + 1, t]$)

for $0 \leq j \leq i - 1$ **do**

Let $a_j \leftarrow 0$ and $\hat{a}_j \leftarrow 0$

Let $\hat{a}_j = a_j + \text{Lap}(\log(T)/\varepsilon)$

Let $B(t) = \sum_{i: \text{Bin}_i(t) \neq 0} \hat{a}_i$

A.3 Counter Privacy Under Adaptive Composition

We are now ready to provide a prove that the prices released by our mechanism satisfy ε -differential privacy.

Theorem 2. *The sequence of prices and counts of unsatisfied bidders released by **PrivateMatching**($\alpha, \rho, \varepsilon$) satisfies ε -differential privacy.*

Proof. Chan et al. [2011] show this for a single sensitivity 1 counter for a non-adaptively chosen stream. We here show the generalization to multiple counters run on adaptively chosen streams with bounded ℓ_1 sensitivity, and bound the ℓ_1 sensitivity of the set of streams produced by our algorithm. We will actually show that the sequence of noisy partial sums released by **Counter** satisfy ε -differential privacy. This is only stronger: the running counts are computed as a function of these noisy partial sums.

To do so, we first define an adversary for the adaptive composition experiment (Definition 7), and then show that the view of this adversary is precisely the sequence of noisy partial sums. The composition theorem (Lemma 8) will then show that the sequence of noisy partial sums are differentially private with respect to a change in a bidder's valuation.

Let the two runs $b = 0, 1$ correspond to any two neighboring valuations (v_i, v_{-i}) and (v'_i, v_{-i}) that differ only in bidder i 's valuation. We first analyze the view on all of the counter(j) for $j = 1, \dots, k$.

The adversary will operate in phases. There are two kinds of phases, which we label P_t and P'_t : one phase per step of the good counters, and one phase per step of the halting condition counter. Both counters run from time 1 to nT , so there are $2nT$ phases in total.

At each point in time, the adversary maintains histories $\{b_i\}, \{b'_i\}$ of all the bids prior to the current phase, and histories $\{e_i\}, \{e'_i\}$ of all prior reports to special counter counter(0), when bidder i has valuation v_i, v'_i respectively. These histories are initially empty.

Let us consider the first kind of phase. One bidder bids per step of the counter, so one bidder bids in each of these phases. Each step of the experiment the adversary will observe a partial sum. Suppose the adversary is in phase P_t . Having observed the previous partial sums, the adversary can simulate the action of the current bidder q from the histories of previous bids by first computing

the prices indicated by the previous partial sums. The adversary will compute q 's bid when the valuations are (v_i, v_{-i}) , and when the valuations are (v'_i, v_{-i}) . Call these two bids b_t, b'_t (which may be \perp if q is already matched in one or both of the histories).

Note that for bidders $q \neq i$, it is always the case that $b_t = b'_t$. This holds by induction: it is clearly true when no one has bid, and bidder q 's decision depends only on her past bids, the prices, and her valuation. Since these are all independent of bidder i 's valuation, bidder q behaves identically.

After the adversary calculates b_t, b'_t , the adversary simulates update and release of the counters. More precisely, the adversary spends phase P_t requesting a set of partial sums

$$\Sigma = \{\sigma_I^j \mid j \in [k], I \in S_t\},$$

where $S_t \subseteq [1, nT]$ is a set of intervals ending at t , corresponding to partial sums that **Counter** releases at step t .

For each $\sigma_I^j \in \Sigma$, $D^0, D^1 \in \{0, 1\}_I$ are defined by

$$D_k^0 = \begin{cases} 1 & : \text{if } b_k = j \\ 0 & : \text{otherwise} \end{cases}$$

and similarly for D^1 , with bid history $\{b'_i\}$. Informally, a database D for σ_I^j encodes whether a bidder bid on good j at every timestep in I . The adversary will pick \mathcal{M} to sum the database and add noise $Lap(1/\varepsilon_0)$, an ε_0 -differentially private operation. Once the partial sums for P_t are released, the adversary will advance to the next phase.

Now, suppose the adversary is in the second kind of phase, say P'_t . This corresponds to a step of the halting condition counter. We use exactly the same construction as above: the adversary will request the partial sums corresponding to each timestep. The adversary will simulate each bidder's action by examining the history of bids and prices. Now suppose the two runs differ in bidder i 's valuation. Following the same analysis, the reports to this halting condition counter differ only in bidder i 's reports.

With this definition, the view of the adversary on database $\{D^0\}$ and $\{D^1\}$ is precisely the noisy partial sums when the valuations are (v_i, v_{-i}) and (v'_i, v_{-i}) , respectively. So, it suffices to show that these views have almost the same probability.

We apply Lemma 8 by bounding the distance between the databases for counter(1) to counter(k). Note that the sequence of databases $\{D^0\}, \{D^1\}$ chosen correspond to streams of bids that differ only in bidder i 's bid, or streams of reports to counter(0) that differ only in bidder i 's report. This is because the bid histories $\{b_i\}, \{b'_i\}$ and report histories $\{e_i\}, \{e'_i\}$ differ only on timesteps where i acts. Thus, it suffices to focus on bidder i when bounding the distance between these databases.

Consider a single good j , and suppose c_j of i 's bids on good j differ between the histories. Each of bidder i 's bids on j show up in $\log(nT)$ databases, so

$$\sum |D_j^0 - D_j^1| \leq c_j \log nT,$$

where the sum is taken over all databases corresponding to good j . The same is true for the halting condition counter: if there are c_0 reports that differ between the histories, then

$$\sum |D_0^0 - D_0^1| \leq c_0 \log nT.$$

Since we know that a bidder can bid at most T times over T proposing rounds, and will report at most T times, we have ℓ_1 sensitivity bounded by

$$\Delta_1 \leq c_0 \log nT + \sum_j c_j \log nT \leq 2T \log nT.$$

By Lemma 8, setting

$$\varepsilon_0 = \frac{\varepsilon}{2T \log nT}$$

suffices for ε -differential privacy, and this is precisely running each **Counter** with privacy level $\varepsilon' = \varepsilon/2T$. \square

B Reconstruction Lower Bound

Here, we detail a basic lower bound about differential privacy. Intuitively, it is impossible for an adversary to recover a database better than random guessing from observing the output of a private mechanism. The theorem is folklore.

Theorem 10. *Let mechanism $\mathcal{M}: \{0, 1\}^n \rightarrow \{0, 1\}^n$ be (ε, δ) -differentially private, and suppose that for all database D , with probability at least $1 - \beta$, $\|\mathcal{M}(D) - D\|_1 \leq \alpha n$. Then,*

$$\alpha \geq 1 - \frac{e^\varepsilon + \delta}{(1 + e^\varepsilon)(1 - \beta)} := c(\varepsilon, \delta, \beta).$$

Proof. Fix a database $D \in \{0, 1\}^n$ and sample an index i uniformly at random from $[n]$. Let D' be a neighboring database of D that differs at the i -th bit. By assumption, we have that with probability at least $1 - \beta$

$$\|\mathcal{M}(D) - D\|_1 \leq \alpha n, \quad \|\mathcal{M}(D') - D'\|_1 \leq \alpha n.$$

Since i is chosen uniformly, we then have

$$\Pr[\mathcal{M}(D)_i = D_i] \geq (1 - \alpha)(1 - \beta), \quad \Pr[\mathcal{M}(D')_i = D'_i] \geq (1 - \alpha)(1 - \beta).$$

It follows that $\Pr[\mathcal{M}(D')_i = D_i] \leq 1 - (1 - \alpha)(1 - \beta)$ because $D_i \neq D'_i$. By definition of (ε, δ) -differential privacy, we get

$$(1 - \alpha)(1 - \beta) \leq \Pr[\mathcal{M}(D)_i = D_i] \leq e^\varepsilon \Pr[\mathcal{M}(D')_i = D_i] + \delta \leq e^\varepsilon (1 - (1 - \alpha)(1 - \beta)) + \delta.$$

Then we have

$$1 - \alpha \leq \frac{e^\varepsilon + \delta}{(1 + e^\varepsilon)(1 - \beta)},$$

as desired. \square