Privacy-Preserving Public Information for Sequential Games

Avrim Blum*, Jamie Morgenstern†, and Ankit Sharma
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA
avrim@cs.cmu.edu,
jamiemmt@cs.cmu.edu,
ankits@cs.cmu.edu

Adam Smith ‡
Computer Science and Engineering Department
Pennsylvania State University
State College, PA
asmith@cse.psu.edu

ABSTRACT

In settings with incomplete information, players can find it difficult to coordinate to find states with good social welfare. For instance, one of the main reasons behind the recent financial crisis was found to be the lack of market transparency, which made it difficult for financial firms to accurately measure the risks and returns of their investments. Although regulators may have access to firms' investment decisions, directly reporting all firms' actions raises confidentiality concerns for both individuals and institutions. The natural question, therefore, is whether it is possible for the regulatory agencies to publish some information that, on one hand, helps the financial firms understand the risks of their investments better, and, at the same time, preserves the privacy of their investment decisions. More generally, when can the publication of privacy-preserving information about the state of the game improve overall outcomes such as social welfare?

In this paper, we explore this question in a sequential resource-sharing game where the value gained by a player on choosing a resource depends on the number of other players who have chosen that resource in the past. Without any knowledge of the actions of the past players, the social welfare attained in this game can be arbitrarily bad. We show, however, that it is possible for the players to achieve good social welfare with the help of *privacy-preserving, publicly-announced information*. We model the behavior of players in this imperfect information setting in two ways – greedy and undominated strategic behaviours, and we prove guarantees about the social welfare that certain kinds of privacy-preserving information can help attain. To achieve the social welfare guarantees, we design a counter with improved privacy guarantees under continual observation. In addition to the resource-sharing game, we study the main question for other games including sequential versions of the cut, machine-scheduling and cost-sharing games, and games where the value attained by a player on a particular action is not only a function of the actions of the past players but also of the actions of the future players.

Categories and Subject Descriptors

F.m [**Theory of Computation**]: Miscellaneous

Keywords

Privacy; Game Theory

1. INTRODUCTION

Multi-agent settings that are non-transparent (where players cannot see the current state of the system) have the potential to lead to disastrous outcomes. For example, in examining causes of the recent financial crisis and subsequent recession, the Financial Crisis Inquiry Commission [4, p. 352] concluded that "The OTC derivatives market's lack of transparency and of effective price discovery exacerbated the collateral disputes of AIG and Goldman Sachs and similar disputes between other derivatives counterparties." Even though regulators have access to detailed confidential information about financial institutions and (indirectly) individuals, current statistics and indices are based only on public data, since disclosures based on confidential information are restricted. However, forecasts based on confidential data can be much more accurate[1], prompting regulators to

*Blum, Morgenstern, and Sharma were partially supported by NSF grants CCF-1116892 and CCF-1101215.

†Morgenstern was partially supported by an NSF GRFP award and a Simons Award for Graduate Students in Theoretical Computer Science.

‡Smith was funded by NSF awards #0747294 and #0941553. Some of this work was done while on sabbatical at Boston University's Hariri Center for Computation, and at Harvard University's Center for Research on Computation & Society, supported by a Simons Investigator grant to Salil Vadhan.

ITCS'15, January 11–13, 2015, Rehovot, Israel.

ACM 978-1-4503-3333-7/15/01 ...$15.00.
http://dx.doi.org/10.1145/2688073.2688100.

[1]For example, Oet et al. [12] compared an index based on both public and confidential data with an analogous index based only on publicly available data. The former index would have been a significantly more accurate predictor of financial stress during the recent financial crisis (see Oet et al. [11, Figure 4]). See Flood et al. [5] for further discussion.

ask whether aggregate statistics can be economically useful while also providing rigorous privacy guarantees [5].

In this work, we show that such *privacy-preserving public information*, in an interesting class of sequential decision-making games, can achieve (nearly) the best of both worlds. In particular, the goal is to produce information about actions taken by previous agents that can be posted publicly, preserves all agents' (differential) privacy, and can significantly improve worst-case social-welfare. While our models do not directly speak to the highly complex issues involved in real-world financial decision-making, they do indicate that in settings involving contention for resources and first-mover advantages, privacy-preserving public information can be a significant help in improving social welfare. In the following sections, we describe the game setting and the information model.

1.1 Game Model

Consider a setting in which there are m resources and n players. The players arrive online, in an *adversarial* order, one at a time[2]. Each player i has some set A_i of resources she is interested in and that is known only to herself. An action a_i of player i is of the form $(a_{i,1}, \ldots, a_{i,m})$, where $a_{i,r} \geq 0$ represents the amount that player i invests in resource r, and moreover, $\sum_{j\in[m]} a_{i,j} = 1$. For simplicity, we assume that all $a_{i,r}$ are in $\{0, 1\}$ i.e, the unit-demand setting (we study the continuous version where $a_{i,r}$'s can be fractional, but still sum to 1, in the full version of this paper). Furthermore, we do not make the assumption that players have knowledge of their position in the sequence, that is, a player need not know how many players have acted before her.

Each resource r has some non-increasing function $V_r : \mathbb{Z}^+ \to \mathbb{R}^+$ indicating the *value, or utility, of this resource to the kth player who chooses it.* Therefore, the utility of player i is $u_i(a_i, a_{1,\ldots,i-1}) = \sum_r a_{i,r} V_r(x_{i,r})$, where $x_{i,r} = \sum_{j=1}^{i-1} a_{j,r}$ for each r. In this **resource sharing** setting, the utility for a player of choosing a certain resource is a function of the resource and (importantly) the number of players who have invested in the resource before her (and not after her)[3].

Illustrative Example.

For each resource, suppose $V_r(k) = V_r(0)/k$, where $V_r(0)$ is the initial value of resource r. The value of each resource r drops rapidly as a function of the number of players who have chosen it so far. If each player i has *perfect information* about the investment choices made by the players before her, the optimal action for player i is to greedily select the action in A_i of highest utility based on the number of players who have selected each resource so far. As shown in Section 3, the resulting social welfare of this behavior is within a factor of 4 of the optimal. In the case where each player has *no information* about other players' behaviors, some particularly disastrous sequences of actions might reasonably occur, leading to very low social welfare. For example, if each player i has access to a public resource r where $V_r(0) = 1$ and a private resource r_i where $V_{r_i}(0) = 1 - \epsilon$, each might reasonably choose greedily according to $V_\cdot(0)$, selecting the resource of highest initial value (in this case, r). This would give social welfare of $\ln(n)$, whereas the optimal assignment would give $n(1 - \epsilon)$. Without information about the game state, therefore, the players may achieve only a $O\left(\frac{\ln(n)}{n}\right)$ fraction of the possible welfare.

[2]For ease of exposition, we rename players such that player i is the ith to arrive.

[3]In Section 5, we consider a generalization where the utility to a player of investing in a particular resource is a function of the total number of players who have chosen that resource, including those who have invested after her.

1.2 Information Model

In resource sharing games, players' decisions about their actions will be best when they know how many players have chosen each resource when they arrive. The mechanisms we consider, therefore, will publicly announce some estimate of these counts. We consider the trade-off between the privacy lost by publishing these estimates and the accuracy of the counters in terms of social welfare. We consider three categories of counters for publicly posting the estimate of resource usage: perfect, private and empty counters.

Perfect Counters: At all points, the counters display the exact usage of each resource.

Privacy-preserving public counters: At all points, the counters display an approximate usage of the resources while maintaining privacy for each player. We define the privacy guarantee in Section 2.

Empty Counters: At all points, every counter displays the value 0.

1.3 Players' Behavior

Each player is a utility-maximizing agent and will choose the resource that, given their beliefs about actions taken by previous players and the publicly displayed counters, gives them maximum value. We analyze the game play under two classes of strategies – greedy and undominated strategies.

1. **Greedy strategy:** Under the greedy strategy, a player has no outside belief about the actions of previous players and chooses the resource that maximizes her utility given the *currently displayed (or announced) values of the counters.* Greedy is a natural choice of strategy to consider since it is the utility-maximizing strategy when the usage counts posted are *perfect.*

2. **Undominated Strategy(UD):** Under undominated strategies, we allow players to have any beliefs about the actions of the previous players that are consistent with the displayed value of the counters[4], and they are allowed to play *any* undominated strategy a_i under this belief. A strategy a_i is *undominated* under a belief, if no other a_i' get a strictly higher utility. [5]

[4]As will become clear in Section 2, we work with privacy-preserving public counters that display values that can be off from the true usage only in a *bounded* range. Hence with these counters, a player's belief is consistent as long as the belief implies the usage of the resource to be a number that is within the bounded range of the displayed value. Moreover, with empty counters, any belief about the actions of previous players is a consistent belief.

[5]For each counter mechanism we consider, there exists at least one undominated strategy. For example, with perfect counters, the only consistent belief is that the true value is equal to the displayed value and here the greedy strategy is always undominated; moreover, if the counter mechanism has a nonzero probability of outputting the true value, then again the greedy strategy is undominated under the belief that the displayed value is the true value; if the counter mechanism can display values that are arbitrarily off from the true value, then for equal initial values *every* strategy is undominated.

We analyze the social welfare $SW(a) = \sum_i u_i(a)$ generated by an announcement mechanism $\mathcal{M}$ for a set of strategies D and compare it to the optimal social welfare OPT. For a game setting g, constituted of a collection of players $[n]$ and their allowable actions A_i (as defined in Section 1.1), $OPT(g)$ is defined as the optimal social welfare that can be achieved by any allocation of resources to the players, where the space of feasible allocations is determined by the setting g. In the unit-demand setting, $OPT(g)$ is the maximum weight matching in the bipartite graph $G = (U \cup V, E)$ where U is the set of the n players, V has n vertices for each resource r, one of value $V_r(k)$ for each $k \in [n]$, and there is an edge between player i and all vertices corresponding to resource r if and only if $r \in A_i$ (Note that the weights are on the vertices in V). The object of our study is $CR_D(g, \mathcal{M})$, the *worst case competitive ratio* of the optimal social welfare to the welfare achieved under strategy D and counter mechanism $\mathcal{M}$. As mentioned earlier, D will either be the greedy (GREEDY) or the undominated (UNDOM) strategy, and $\mathcal{M}$ will be either the perfect ($\mathcal{M}_{Full}$), the privacy-preserving or the empty ($\mathcal{M}_\emptyset$) counter. When $\mathcal{M}$ uses internal random coins, our results will either be worst-case over all possible throws of the random coins, or will indicate the probability with which the social welfare guarantee holds.

1.4 Statement of Main Results

For sequential resource-sharing games, we prove that for all nonincreasing value curves, the greedy strategy following privacy-preserving counters has a competitive ratio *polylogarithmic* in the number of players (Theorem 5). This should be contrasted with the competitive ratio of 4 achieved by greedy w.r.t. perfect counters (Theorem 1) and the nearly-linear (in the number of players) competitive ratio of greedy with empty counters (as shown in the illustrative example in Section 1.1). For the case of undominated strategies, when the marginal values of resources drop slowly, (for example, at a polynomial rate, $V_r(k) = V_r(0)/k^p$ for constant $p > 0$), we bound the competitive ratio (w.r.t. privacy-preserving counters) (Theorem 7). With empty counters, the competitive ratio for undominated strategies is unbounded (Theorem 2) for arbitrary curves and is at least quadratic (in the number of players) if the value curve drops slowly (Theorem 3). We note here that for many of our positive results for privacy preserving counters state the competitive ratio in terms of parameters of the counter vector α and β (as detailed in Section 2) and for a particular implementation of the counter vectors, the values of α and β are mentioned in Section 4.

The key privacy tool we use is the differentially private counter under continual observation [3], which we use to publish estimates of the usage of each resource. We improve upon the existing error guarantees of differentially private counters and design a new differentially private counter in Section 4. The new counter provides a tighter additive guarantee at the price of introducing a constant multiplicative error.

In Section 5, we consider other classes of games – specifically, we analyze Unrelated Machine Scheduling, Cut, and Cost Sharing games. The work of Leme et al. [10] showed these games have improved *sequential* price of anarchy over the *simultaneous* price of anarchy. For these games, we ask the question: if players do not have perfect information to make decisions, but instead have only noisy approximations (due to privacy considerations), does sequentiality still improve the quality of play? We prove that the answer is affirmative in many cases.

1.5 Related Work

A great deal of work has been done at the intersection of mechanism design and privacy; Pai and Roth [13] have an extensive survey. Our work is similar to much of the previous work in that it considers maintaining differential privacy to be a constraint. The focus of our work however is on *how useful information can be provided to players in games of imperfect information* to help achieve a good social objective while respecting the privacy constraint of the players. The work of Kearns et al. [9] is close in spirit to ours. Kearns et al. [9] consider games where players have incomplete information about other players' types and behaviors. They construct a privacy-preserving mechanism which collects information from players, computes an approximate correlated equilibria, and then advises players to play according to this equilibrium. The mechanism is approximately incentive compatible for the players to participate in the mechanism and to follow its suggestions. Several later papers [14, 7] privately compute approximate equillibria in different settings. Our main privacy primitive is the differentially private counters under continual observation [3, 2], also used in much of the related work on private equilibrium computation.

Our investigation of cut games, unrelated machine scheduling, and cost-sharing (Section 5) is inspired by work of Leme et al. [10]. Their work focuses on the improvement in social welfare of equilibria in the *sequential* versus the *simultaneous* versions of certain games. We ask a related question: when we consider sequential versions of games, and only *private, approximate information about the state of play* (as opposed to perfect) is given to players, how much worse can social welfare be?

As mentioned in Section 1.3, one class of player behavior for which we analyze the games is *greedy*. Our analysis of greedy behavior is in part inspired by the work of Balcan et al. [1], who study best response dynamics with respect to noisy cost functions for potential games. An important distinction between their setting and ours is that the noisy estimates we consider are *estimates of state, not value*, and may for natural value curves be quite far from correct in terms of the *values* of the actions.

2. PRIVACY-PRESERVING PUBLIC COUNTERS

We design announcement mechanisms $\mathcal{M}_i$ which give approximate information about actions made by the previous players to player i. Let Δ_m denote the action space for each player (the m-dimensional simplex $\Delta_m = \{a \in [0,1]^m \mid \|a\|_1 \leq 1\}$). Mechanism $\mathcal{M}_i : (\Delta_m)^{i-1} \times R \to \Delta_m$ depends upon the actions taken before i (specifically, the usage of each resource by each player), and on internal random coins R. When player i arrives, $m_i(a_1, \ldots, a_{i-1}) \sim \mathcal{M}_i(a_1, \ldots, a_{i-1})$ is publicly announced. Player i plays according to some strategy $d_i : \Delta_m \to A_i$, that is $a_i = d_i(m_1, \ldots, m_i(a_1, \ldots, a_{i-1}))$, a random variable which is a function of this announcement. When it is clear from context, we denote $m_i(a_1, \ldots, a_{i-1})$ by m_i. Formally, the counters used in this paper satisfy the following notion of privacy.

DEFINITION 1. *An announcement mechanism $\mathcal{M}$ is (ϵ, δ)-differentially private under adaptive*[6] *continual observation in the strategies of players if, for each d, for each player i, each pair of strategies d_i, d_i', and every $S \subseteq (\Delta_m)^n$:*

$$\mathbb{P}[(m_1, \ldots, m_n) \in S] \leq e^{\epsilon}\mathbb{P}[(m_1, \ldots, m_i, m_{i+1}' \ldots, m_n') \in S] + \delta,$$

where $a_j = d_j(m_1, \ldots, m_j)$, $a_i' = d_i'(m_1, \ldots, m_i)$, $m_j \sim \mathcal{M}_j(a_1, \ldots, a_{j-1})$, $m_j' \sim \mathcal{M}_j(a_1, \ldots, a_{i-1}, a_i', a_{i+1}', \ldots, a_{j-1}')$ and for all $j > i$, $a_j' = d_j(m_1, \ldots, m_{i-1}, m_i, m_{i+1}', \ldots, m_j')$

This definition requires that two worlds which differ in a single player changing her strategy from d_i to d_i' have statistically close joint distributions over all players' announcements (and thus their joint distributions over actions). Note that the distribution of $j > i$'s announcement can change slightly, causing j's distribution over actions to change slightly, necessitating the cascaded m_j', a_j' for $j > i$ in our definition. The mechanisms we use maintain approximate use counters for each resource. The values of the counters are *publicly announced* throughout the game play. We now define the notion of accuracy used to describe these counters.

DEFINITION 2 ((α, β, γ)-ACCURATE COUNTER VECTOR). *A set of counters $y_{i,r}$ is defined to be (α, β, γ)-accurate if with probability at least $1 - \gamma$, at all points of time, the displayed value of every counter $y_{i,r}$ lies in the range $[\frac{x_{i,r}}{\alpha} - \beta, \alpha x_{i,r} + \beta]$ where $x_{i,r}$ is the* true *count for resource i, and is monotonically increasing in the true count.*

We refer to a set of $(\alpha, \beta, 0)$-accurate counters as (α, β)-counters for brevity. It is possible to achieve $\gamma = 0$ (which is necessary for undominated strategies, which assumes the multiplicative and additive bounds on y are worst-case), taking an appropriate loss in the privacy guarantees for the counter (Proposition 1). Counters satisfying Definitions 1 and 2 with $\alpha = 1$ and $\beta = O(\log^2 n)$ were given in Dwork et al. [3], Chan et al. [2]; we give a different implementation in Section 4 which gives a tighter bound on $\alpha\beta$ by taking α to be a small constant larger than 1. Furthermore, the counters in Section 4 are *monotonic* (i.e., the displayed values can only increase as the game proceeds) and we use monotonicity of the counters in some of our results.

In some settings we require counters we a more specific utility guarantee:

DEFINITION 3 ((α, β, γ)-ACCURATE UNDERESTIMATOR). *A set of counters $y_{i,r}$ is defined to be (α, β, γ)-accurate underestimator if with probability at least $1 - \gamma$, at all points of time, the displayed value of every counter $y_{i,r}$ lies in the range $[\frac{x_{i,r}}{\alpha} - \beta, x_{i,r}]$ where $x_{i,r}$ is the* true *count for resource i.*

The following observation states that a counter vector can be converted to an undercounter with small loss in accuracy.

OBSERVATION 1. *We can convert a (α, β)-counter to an $\left(\alpha^2, \frac{2\beta}{\alpha}\right)$-underestimating counter vector.*

PROOF. We can shift the counter, $\frac{1}{\alpha}x - \beta \leq y \leq \alpha x + \beta$ implies $y' = \frac{y-\beta}{\alpha} \leq x$ and $\frac{1}{\alpha^2}x - \frac{2\beta}{\alpha} \leq y'$. □

[6]Adaptivity is needed in this case because the announcements are arguments to the actions of players: when a particular action changes, this modifies the distribution over the future announcements, which in turn changes the distribution over future selected actions.

3. RESOURCE SHARING

In this section, we consider resource sharing games – the utility to a player is completely determined by the resource she chooses and the number of players who have chosen that resource before her. This section considers the case where players' actions are discrete: $a_i \in \{0, 1\}^m$ for all $i, a_i \in A_i$. We defer the analysis of the case where players' actions are continuous to the full version of this paper.

3.1 Perfect counters and empty counters

Before delving into our main results, we point out that, with perfect counters, greedy is the only undominated strategy, and the competitive ratio of greedy is a constant.

THEOREM 1. *With perfect counters, greedy behavior is dominant-strategy and all other behavior is dominated for any sequential resource-sharing game g; and $CR_{\text{GREEDY}}(\mathcal{M}_{Full}, g) \leq 4$.*

The proof of this Theorem follows from the connection between future-independent resource-sharing and online vertex-weighted matching, which we mention below.

OBSERVATION 2. *In the setting where $\|a_i\|_1 = 1$ for all $a_i \in A_i$, for all i, full-information, discrete resource-sharing reduces to online, vertex-weighted bipartite matching.*

PROOF. Construct the following bipartite graph $G = (U, V, E)$ as an instance of online vertex-weighted matching from an instance of the future-independent resource sharing game. For each resource r, create n vertices in V, one with weight $V_r(t)$ for each $t \in [n]$. As players arrive online, they will correspond to vertices in $u_i \in U$. For each $a_i \in A_i$ corresponding to a set of resources S, u_i is allowed to take any subset of V with a single copy of each $r \in S$. □

The proof of the social welfare is quite similar to the one-to-one, online vertex-weighted matching proof of [8], with the necessary extension for many-to-one matchings (losing a factor of 1/2 in the process).

PROOF OF THEOREM 1. Consider any instance of $G = (U, V, E)$, a vertex-weighted bipartite graph. Let μ be the optimal many-to-one matching, which can be applied to nodes in both U and V (where $u \in U$ has potentially multiple neighbors in V). Consider μ', the greedy many-to-one matching for a particular sequence of arrivals σ.

Consider a particular $u \in U$, and the time it arrives $\sigma(u)$ as μ' progresses. If at least 1/2 the value of $\mu(u)$ is available at that time, then $w(\mu'(u)) \geq \frac{1}{2}w(\mu(u))$ (since u can be matched to any subset of $\mu(u)$, by the downward closed assumption). If not, then $w(\mu'(\mu(u))) \geq \frac{1}{2}w(\mu(u))$ (at least half the value was taken by others). Thus, we know that, for all u,

$$w(\mu'(u)) + w(\mu'(\mu(u))) \geq \frac{1}{2}w(\mu(u))$$

summing up over all u, we get

$$\sum_u w(\mu'(u)) + w(\mu'(\mu(u))) = 2w(\mu') \geq \frac{1}{2}\sum_u w(\mu(u)) = \frac{1}{2}w(\mu)$$

Rearranging shows that $w(\mu') \geq \frac{1}{4}w(\mu)$.

Finally, the utility to a player is clearly greatest when they are greedy, so that is a dominant strategy (thus implying any non-greedy strategy is dominated). □

Recall, from our example in the introduction, that both greedy and undominated strategies can perform poorly with respect to empty counters. We defer the proof of the following results to the full version of the paper. Recall that $\mathcal{M}_\emptyset$ refers to the empty counter mechanism.

THEOREM 2. *There exist games g such that $CR_{\text{UNDOM}}(\mathcal{M}_\emptyset, g)$ cannot be bounded by any function of n.*

THEOREM 3. *There exists g such that $CR_{\text{UNDOM}}(\mathcal{M}_\emptyset, g) \geq \Omega(\frac{n^2}{\log(n)})$, when $V_r(t) = \frac{V_r(0)}{t}$.*

3.2 Privacy-Preserving Counters and Greedy Behavior

THEOREM 4. *With (α, β)-accurate underestimator counter mechanism $\mathcal{M}$, $CR_{\text{GREEDY}}(\mathcal{M}, g) = O(\alpha\beta)$ for all resource-sharing games g.*

Before we prove Theorem 4, we need a way to compare players' utilities with the utility they *think* they get from choosing resources greedily with respect to approximate counters. Let a player's *perceived value* be $V_r(y_{i,r})$ where r is the resource she chose (the value of a resource if the counter was correct, which may or may not be the *actual* value of the resource).

LEMMA 1. *Suppose players choose greedily according to a (α, β)-underestimator. Then, the sum of their actual values is at least a $\frac{1}{2\alpha\beta}$-fraction of the sum of their perceived values.*

PROOF. Suppose k players chose a given resource r. For ease of notation, let these be players 1 through k. We wish to bound the ratio

$$\frac{\sum_{i=1}^{k} V_r(y_{i,r})}{\sum_{c=1}^{k} V_r(c)}.$$

We start by "grouping" the counter values: it cannot take on values that are small for more than a certain number of steps. In particular, if $x_{i,r} > T\alpha\beta$, for some $T \in \mathbb{N}$,

$$y_{i,r} \geq \frac{1}{\alpha} x_{i,r} - \beta \geq \frac{T\alpha\beta}{\alpha} - \beta = (T-1)\,\beta$$

Now, we bound the ratio from above using this fact.

$$\frac{\sum_{i=1}^{k} V_r(y_{i,r})}{\sum_{c=1}^{k} V_r(c)} \leq \frac{2\alpha\beta \sum_{T=1}^{\lceil \frac{k}{\alpha\beta} \rceil} V_r((T-1)\beta)}{\sum_{c=1}^{k} V_r(c)}$$

$$\leq \frac{2\alpha\beta \sum_{T=1}^{\lceil \frac{k}{\alpha\beta} \rceil} V_r((T-1)\beta)}{\sum_{T=1}^{\lceil \frac{k}{\alpha\beta} \rceil} V_r((T-1)\beta)} \leq 2\alpha\beta$$

where the first inequality came from the fact that the value curves are non-increasing and the lower bound on the counter values from above, and the second because all terms are non-negative. □

PROOF OF THEOREM 4. The optimal value of the resource-sharing game g, denoted by $OPT(g)$, is the maximum weight matching in the bipartite graph $G = (U \cup V, E)$ where U is the set of the n players and V has n vertices for each resource r, one of value $V_r(k)$ for each $k \in [n]$. There is an edge between player i and all vertices corresponding to resource r if and only if $r \in A_i$. Note that the weights are on the vertices in V.

We now define a complete bipartite graph G' which has the same set of nodes but whose node weights differ for some nodes in G. Consider some resource r, and the collection of players who chose r in g. If there were t_k players i who chose resource r when $y_{i,r} = k$, make t_k of the nodes corresponding to r have weight $V_r(k)$. Finally, if there were F_k players who chose resource r, let the remaining $n - F_k$ nodes corresponding to r have weight $V_r(F_k + 1)$.

We first claim that the perceived utility of players choosing greedily according to the counters is identical to the weight of the greedy matching in G' (where nodes arrive in the same order). We prove, in fact, that the corresponding matching will be identical by induction. Since the counters are monotone, earlier copies of a resource appear more valuable. So, when the first player arrives in G', the most valuable node she has access to is exactly the first node corresponding to the resource she took according to the counters. Now, assume that prior to player i, all players have chosen nodes corresponding to the resource they chose according to the counters. By our induction hypothesis and monotonicity of the counters and value curves, there is a node n_i corresponding to i's selection r according to counters of weight $V_r(y_{i,r})$, and no heavier node corresponding to r. Likewise, for all other resources r', all nodes corresponding to r' have weight more than $V_{r'}(y_{i,r'})$. Thus, i will take n_i for value $V_r(y_{i,r})$. Thus, the weight of the greedy matching in G' equals the perceived utility of greedy play according to the counters.

Let $\text{GREEDY}_{\text{COUNTERS}}$ denote the set of actions players make playing greedily with respect to the counters. By Lemma 1, the social welfare of $\text{GREEDY}_{\text{COUNTERS}}$ is a $\frac{1}{\alpha\beta}$-fraction of the perceived social welfare. By our previous argument, the perceived social welfare of greedy play according to the counters is the same as the weight of the greedy matching in G'. By Theorem 1, the greedy matching in G' is a 4-approximation to the max-weight matching in G'. Finally, since the counters are underestimators, the weight of the max-weight matching in G' is at least as large as $OPT(g)$. Thus, that the social welfare of greedy play with respect to counters is a $\frac{1}{2\alpha\beta}$ fraction of the optimal welfare of g. □

THEOREM 5. *There exists (ϵ, δ)-privacy-preserving mechanism $\mathcal{M}$ such that*

$$CR_{\text{GREEDY}}(\mathcal{M}, g) \leq \min\left(O\left(\frac{\log n \log \frac{nm}{\delta}}{\epsilon}\right), O\left(\frac{m \log n \log\log \frac{1}{\delta}}{\epsilon}\right)\right)$$

for all resource-sharing games g.

PROOF. In Section 4, we prove Corollary 2 that says that we can achieve an (ϵ, δ)-differentially private counter vector achieving the better of $(1, O(\frac{(\log n)(\log(nm/\delta))}{\epsilon}))$-accuracy and $(\alpha, \tilde{O}_\alpha(\frac{m \log n \log\log(1/\delta)}{\epsilon}))$-accuracy for any constant $\alpha > 1$. This along with Theorem 4 proves the result. □

Observation 3 (whose proof can be found in the full version) states that players acting greedily according to any estimate that is *deterministically* more accurate than the values provided by the private counters also achieve similar or better social welfare guarantees. Moreover, we show that if the estimates used by the players are more accurate only in expectation, as opposed to deterministically, then we cannot make a similar claim (Observation 4).

OBSERVATION 3. *Suppose that $\mathcal{M}$ is a (α, β, γ) underestimator, giving estimates $y_{i,r}$. Furthermore, assume each player i is playing greedily with respect to a revised estimate $z_{i,r}$ such that, for each r, i, and value of $z_i r$ is always in the range $[y_{i,r}, x_{i,r}]$. Then, for g, a discrete resource-sharing game, with probability $1-\gamma$, the ratio of the optimal to the achieved social welfare is $O(\alpha\beta)$.*

PROOF. The proof follows from the proof of Theorem 4, along with the following observation. Since $z_{i,r}$'s is deterministically more accurate than the COUNTERS, we have for each i that the value gained by greedily choosing according to the estimates $z_{i,r}$ is at least as much as the value gained by greedily choosing using $y_{i,r}$. Therefore, summing over all the players, the achieved social welfare is at least as much as it would be if everyone had played greedily according to $y_{i,r}$. □

OBSERVATION 4. *There exists a resource-sharing game g, such that if the players play greedily according to estimates $z_{i,r}$ that are more accurate than the displayed value only in expectation – specifically for each r, i, and value of $x_{i,r}$, $\mathbb{P}[z_{i,r} < x_{i,r}] \geq 1/2$ and also $\mathbb{E}[|z_{i,r} - x_{i,r}|] = 1$, then the ratio of the optimal to the achieved social welfare can be as bad as $\Omega(\sqrt{n})$.*

PROOF. Let there be $n + \sqrt{n}$ resources, with resources $r*_{1,\ldots,\sqrt{n}}$ having $V_{r*_f}(0) = H$, $V_{r*_f}(t) = 0$ for all $t > 0$, and resource r_i such that $V_{r_i}(t) = H - \epsilon$ for all t. Player i has access to all resources $r*_f$ and r_i. Then, $OPT = H\sqrt{n} + (H-\epsilon)(n-\sqrt{n}) = Hn - (n-\sqrt{n})\epsilon$.

Consider the counter vector which is exactly correct with probability $1 - \frac{1}{\sqrt{n}}$ and undercounts by $\sqrt{n}$ with probability $\frac{1}{\sqrt{n}}$ (note that the expected error is just 1 and it undercounts with probability 1). Then, greedy behavior with respect to this counter will (in expectation) have $\sqrt{n}$ players choose $r*_f$ for each f, achieving welfare $\sqrt{n}H$. Thus, the competitive ratio is $\Omega(\sqrt{n})$ as $\epsilon \to 0$, as desired. □

3.3 Privacy-Preserving Counters and Undominated behavior

We begin with an illustration of how undominated strategies can perform poorly for arbitrary value curves, as motivation for the restricted class of value curves we consider in Theorem 7. In the case of greedy players, we were able to avoid the problem of players undervaluing resources rather easily, by forcing the counters to only underestimate $x_{i,r}$. This won't work for undominated strategies: players who know the counts are shaded downward can compensate for that fact.

THEOREM 6. *For an (ϵ, δ)–differentially private announcement mechanism $\mathcal{M}$, there exist games g for which*

$$CR_{\text{UNDOM}}(g, \mathcal{M}) = \Omega\left(\frac{1}{\delta}\right).$$

PROOF. Suppose there are two players 1 and 2, and resources r, r'. Let r have $V_r(0) = 1$, $V_r(1) = 0$, and $V_{r'}(k) = \rho$, for all $k \geq 0$. Furthermore, let player 1 have access only to resource r' but player 2 has access to both r and r'. Player 1 will choose r'. Let player 2's strategy be d_2, such that if she determines there was nonzero chance that player 1 chose r according to her signal m_2, she will choose resource r'. This is undominated: if 1 *did* choose r, r' will be more valuable for 2. Thus, if 2 sees any signal that can occur when r is chosen by 1, she will choose r'. The collection of signals 2 can see if 1 chooses r has probability 1 in total. So, because m_2 is (ϵ, δ)-differentially private in player 1's action, the set of signals *reserved* for the case when 1 chooses r' (that cannot occur when r is chosen by 1) may occur with probability at most δ (they can occur with probability 0 if 1 chose r, implying they can occur with probability at most δ when 1 chooses r'). Thus, with this probability $1-\delta$, player 2 will choose r', implying $\mathbb{E}[SW] \leq (1-\delta)2\rho + \delta(1+\rho) = \delta + (2-\delta)\rho$, which for ρ sufficiently small approaches δ, while $1+\rho$ is the optimal social welfare. □

Given the above example, we cannot hope to have a theorem as general as Theorem 4 when analyzing undominated strategies with privacy-preserving counters. Instead, we show that, for a class of well-behaved value curves, we can bound the competitive ratio of undominated strategies (Theorem 7).

Again, along the lines of the greedy case, we show that any player who chooses any undominated resource r' over resource r gets a reasonable fraction of the utility she would get from choosing r. Then, by the analysis of greedy players, we have an analogous argument implying the bound of Theorem 7.

THEOREM 7. *If each value curve V_r has the property that $\psi(\alpha, \beta)V_r(x) \geq V_r((\max\{0, \frac{x}{\alpha^2} - \frac{2\beta}{\alpha}\}))$ and also $V_r((\alpha^2 x + 2\alpha\beta)) \geq \phi(\alpha, \beta)V_r(x)$, then an action profile a of undominated strategies according to (α, β)-counter vector $\mathcal{M}$ has $CR_{\text{UNDOM}}(g, \mathcal{M}) = O(\psi(\alpha, \beta)\phi(\alpha, \beta))$.*

In particular, Theorem 7 shows that, for games where $V_r(i) = \frac{V_r(0)}{g_r(x_{i,r})}$, where g_r is a polynomial, the competitive ratio of undominated strategies degrades gracefully as a function of the maximum degree of those polynomials. A simple calculation implies the following corollary, whose proof we relegate to the full version.

COROLLARY 1. *Suppose for a resource-sharing game g, each resource r has a value curve of the form $V_r(x) = \frac{V_r(0)}{g_r(x)}$, where g_r is a monotonically increasing degree-d polynomial and $V_r(0)$ is some constant. Then, $CR_{\text{UNDOM}}(g, \mathcal{M}) \leq O(2\alpha^3\beta)^d$ with $\mathcal{M}$ providing $(\alpha, \beta)-$counters.*

4. PRIVATE COUNTERS WITH SMALLER ERROR AT SMALLER VALUES

In this section, we describe a counter for the model of differential privacy under continual observation that has improved guarantees when the value of the counter is small. Recall the basic counter problem: given a stream $\vec{a} = (a_1, a_2, ..., a_n)$ of numbers $a_i \in [0, 1]$, we wish to release at every time step t the partial sum $x_t = \sum_{i=1}^{t} a_i$. We require a generalization, where one maintains a vector of m counters. Each player's update contribution is now a vector $a_i \in \Delta_m = \{a \in [0,1]^m \mid \|a\|_1 \leq 1\}$. That is, a player can add non-negative values to all counters, but the total value of her updates is at most 1. The partial sums x_t then lie in $(\mathbb{R}^+)^m$ and have ℓ_1 norm at most t.

Given an algorithm $\mathcal{M}$, we define the output stream $(y_1, ..., y_n) = \mathcal{M}(\vec{a})$ where $y_t = \mathcal{M}(t, a_1, ..., a_{t-1})$ lies in $\mathbb{R}^m$. We seek counters that are private (Definition 1) and satisfy a mixed multiplicative and additive accuracy guarantee

(Definition 2). Proofs of all the results in this section can be found in the full verison of this paper.

The original works on differentially private counters [3, 2] concentrated on minimizing the additive error of the estimated sums, that is, they sought to minimize $\|x_t - y_t\|_\infty$. Both papers gave a binary tree-based mechanism, which we dub "TreeSum", with additive error approximately $(\log^2 n)/\epsilon$. Some of our algorithms use TreeSum, and others use a new mechanism (FTSum, described below) which gets a better additive error guarantee at the price of introducing a small multiplicative error. Formally, they prove:

LEMMA 2. *For every $m \in \mathbb{N}$ and $\gamma \in (0,1)$: Running m independent copies of TreeSum [3, 2] is $(\epsilon, 0)$-differentially private and provides an $(1, C_{tree} \cdot \frac{\log n \log \frac{nm}{\gamma}}{\epsilon}, \gamma)$-approximation to partial vector sums, where $C_{tree} > 0$ is an absolute constant.*

Even for $m = 1, \alpha = 1$, this bound is slightly tighter than those in Chan et al. [2] and Dwork et al. [3]; however, it follows directly from the tail bound in Chan et al. [2].

Our new algorithm, FTSum (for Flag/Tree Sum), is described in Algorithm 1. For small m ($m = o(log(n))$), it provides lower additive error at the expense of introducing an arbitrarily small constant multiplicative error.

LEMMA 3. *For every $m \in \mathbb{N}$, $\alpha > 1$ and $\gamma \in (0,1)$, FTSum (Algorithm 1) is $(\epsilon, 0)$-differentially private and $(\alpha, \tilde{O}_\alpha(\frac{m \log \frac{n}{\gamma}}{\epsilon}), \gamma)$-approximates partial sums (where $\tilde{O}_\alpha(\cdot)$ hides polylogarithmic factors in its argument, and treats α as constant).*

FTSum proceeds in two phases. In the first phase, it increments the reported output value only when the underlying counter value has increased significantly. Specifically, the mechanism outputs a public signal, which we will call a "flag", roughly when the true counter achieves the values $\log n$, $\alpha \log n$, $\alpha^2 \log n$ and so on, where α is the desired *multiplicative* approximation. The reported estimate is updated each time a flag is raised (it starts at 0, and then increases to $\log n$, $\alpha \log n$, etc). The privacy analysis for this phase is based on the "sparse vector" technique of Hardt and Rothblum [6], which shows that the cost to privacy is proportional to the number of times a flag is raised (but not the number of time steps between flags).

When the value of the counter becomes large (about $\frac{\alpha \log^2 n}{(\alpha-1)\epsilon}$), the algorithm switches to the second phase and simply uses the TreeSum protocol, whose additive error (about $\frac{\log^2 n}{\epsilon}$) is low enough to provide an α multiplicative guarantee (without need for the extra space given by the additive approximation).

If the mechanism were to raise a flag *exactly* when the true counter achieved the values $\log n$, $\alpha \log n$, $\alpha^2 \log n$, etc, then the mechanism would provide a $(\alpha, \log n, 0)$ approximation during the first phase, and a $(\alpha, 0, 0)$ approximation thereafter. The rigorous analysis is more complicated, since flags are raised only near those thresholds.

Enforcing Additional Guarantees.

Finally, we note that it is possible to enforce to additional useful properties of the counter. First, we may insist that the accuracy guarantees be satisfied with probability 1 (that is, set $\gamma = 0$), at the price of increasing the additive term δ in the privacy guarantee:

Algorithm 1: FTSum — A Private Counter with Low Multiplicative Error

Input: Stream $\vec{a} = (a_1, ..., a_n) \in ([0,1]^m)^n$, parameters $m, n \in \mathbb{N}$, $\alpha > 1$ and $\gamma > 0$
Output: Noisy partial sums $y_1, ..., y_n \in \mathbb{R}^m$

$k \leftarrow \lceil \log_\alpha(\frac{\alpha}{\alpha-1} \cdot C_{tree} \cdot \frac{\log(nm/\gamma)}{\epsilon}) \rceil$;
/* C_{tree} is the constant from Lemma 2 */
$\epsilon' \leftarrow \frac{\epsilon}{2m(k+1)}$;
for $r = 1$ **to** m **do**
  $\text{flag}_r \leftarrow 0$;
  $x_{0,r} \leftarrow 0$;
  $\tau_r \leftarrow (\log n) + \mathsf{Lap}(2/\epsilon')$;
for $i = 1$ **to** n **do**
  for $r = 1$ **to** m **do**
    if $flag_r \leq k$ **then** (First phase still in progress for counter r)
      $x_{i,r} \leftarrow x_{i-1,r} + a_{i,r}$;
      $\tilde{x_{i,r}} \leftarrow x_{i,r} + \mathsf{Lap}(\frac{2}{\epsilon'})$;
      if $\tilde{x_{i,r}} > \tau_r$ **then** (Raise a new flag for counter r)
        $\text{flag}_r \leftarrow \text{flag}_r + 1$;
        $\tau_r \leftarrow (\log n) \cdot \alpha^{\text{flag}_r} + \mathsf{Lap}(2/\epsilon')$;
      Release $y_{i,r} = (\log n) \cdot \alpha^{\text{flag}_r - 1}$;
    else (Second phase has been reached for counter r)
      Release $y_{i,r} = r$-th counter output from TreeSum$(\vec{a}, \epsilon/2)$;

PROPOSITION 1. *If $\mathcal{M}$ is (ϵ, δ)-private and (α, β, γ)-accurate, then one can modify $\mathcal{M}$ to obtain an algorithm $\mathcal{M}'$ with the same efficiency that is $(\epsilon, \delta + \gamma)$-private and $(\alpha, \beta, 0)$-accurate.*

Second, as in [3], we may enforce the requirement that the reported values be monotone, integral values that increase at each time step by at most 1. The idea is to simply report the nearest integral, monotone sequence to the noisy values (starting at 0 and incrementing the reported counter only when the noisy value exceeds the current counter).

PROPOSITION 2 ([3]). *If $\mathcal{M}$ is (ϵ, δ)-private and (α, β, γ)-accurate, then one can modify $\mathcal{M}$ to obtain an algorithm $\mathcal{M}'$ which reports monotone, integral values that increase by 0 or 1 at each time step, with the same privacy and accuracy guarantees as $\mathcal{M}$.*

COROLLARY 2. *Algorithm 1 is an (ϵ, δ)-differentially private vector counter algorithm providing a*

1. *$(1, O(\frac{(\log n)(\log(nm/\delta))}{\epsilon}), 0)$-approximation (using modified TreeSum); or*
2. *$(\alpha, \tilde{O}_\alpha(\frac{m \log n \log \log(1/\delta)}{\epsilon}), 0)$-approximation for any constant $\alpha > 1$ (using FTSum).*

5. EXTENSIONS

In the full version of this paper, we also consider settings where players' utility when choosing a resource r depends upon the *total number* of players choosing r, not just the players who chose r before. In addition, we study several other classes of games: namely, cut games, consensus games,

and unrelated machine scheduling, and consider whether or not private synopses of the state of play is sufficient to improve social welfare over simultaneous play, as perfect synopses have been proven to be in Leme et al. [10].

6. DISCUSSION AND OPEN PROBLEMS

In this work, we considered how public dissemination of information in sequential games can guarantee a good social welfare while maintaining differential privacy of the players' strategies. We considered two 'extreme' cases – the greedy strategy and the class of all undominated strategies. While analyzing the class of undominated strategies gives guarantees that are robust, in many games that we considered, the competitive ratios were significantly worse than greedy strategies, and in some cases they were unbounded. It is interesting to note that many of the examples in this paper that show the poor performance with undominated strategies also hold when the players know their position in the sequence, an assumption we have not made for any of the positive results in this work. It is an interesting direction for future research to consider classes of strategies that more restricted than undominated strategies yet are general enough to be relevant for games where players play with imperfect information.

As mentioned in the introduction, we note here that, while players are making choices subject to approximate information, our results are not a direct extension of the line of thought that approximate information implies approximate optimization. In particular, for greedy strategies, while there may be a bound on the error of the counters, that *does not imply*, for arbitrary value curves, playing greedily according to the counters will be *approximately optimal for each individual*. In particular, consider one resource r with value H for the first 10 investors, and value 0 for the remaining investors, and a second resource r' with value $H/2$ for all investors. With (α, β, γ), as many as β players might have unbounded ratio between their value for r as r', but will pick r over r'. The analysis of greedy shows, despite this anomaly, the total social welfare is still well-approximated by this behavior.

All of our results relied on using differentially private counters for disseminating information. For the differentially-private counter, a main open question is "what is the optimal trade-off between additive and multiplicative guarantees?". Furthermore, as part of future research, one can consider other privacy techniques for announcing information that can prove useful in helping players achieve a good social welfare. And more generally, we want to understand what features of games lend themselves to be amenable to public dissemination of information that helps achieve good welfare and simultaneously preserves privacy of the players' strategies.

References

[1] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. The Price of Uncertainty. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, pages 285–294, 2009.

[2] T.-H. Hubert Chan, Elaine Shi, and Dawn Song. Private and Continual Release of Statistics. *ACM Trans. Inf. Syst. Secur.*, 14(3):26, 2011.

[3] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential Privacy under Continual Observation. In *Symposium on Theory of Computing*, STOC '10, pages 715–724. ACM, 2010.

[4] Financial Crisis Inquiry Commission. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. U.S. Government Printing Office, 2011. URL `http://fcic.law.stanford.edu/report`.

[5] Mark D. Flood, Jonathan Katz, Stephen J. Ong, and Adam Smith. Cryptography and the Economics of Supervisory Information: Balancing Transparency and Confidentiality. Working Paper #11, Office of Financial Research, US Department of Treasury, August 2013.

[6] Moritz Hardt and Guy N. Rothblum. A Multiplicative Weights Mechanism for Privacy-Preserving Data Analysis. In *FOCS '10*, 2010.

[7] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private Matchings and Allocations. *CoRR*, abs/1311.2828, 2013.

[8] Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An Optimal Algorithm for On-line Bipartite Matching. In *STOC '90*, pages 352–358, 1990.

[9] Michael Kearns, Mallesh M. Pai, Aaron Roth, and Jonathan Ullman. Mechanism Design in Large Games: Incentives and privacy. *CoRR*, abs/1207.4084, 2012.

[10] Renato Paes Leme, Vasilis Syrgkanis, and Éva Tardos. The Curse of Simultaneity. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 60–67, 2012.

[11] Mikhail V. Oet, Timothy Bianco, Dieter Gramlich, and Stephen J. Ong. Safe: An early warning system for systemic banking risk. Working Paper 11-29, Federal Reserve Bank of Cleveland, 2011. URL `http://www.clevelandfed.org/research/workpaper/2011/wp1129.pdf`.

[12] Mikhail V. Oet, Timothy Bianco, Dieter Gramlich, and Stephen J. Ong. Financial Stress Index: A Lens for Supervising the Financial System. Working Paper 12-37, Federal Reserve Bank of Cleveland, 2012. URL `http://www.clevelandfed.org/research/workpaper/2012/wp1237.pdf`.

[13] Mallesh M. Pai and Aaron Roth. Privacy and mechanism design. *CoRR*, abs/1306.2083, 2013.

[14] Ryan M Rogers and Aaron Roth. Asymptotically Truthful Equilibrium Selection in Large Congestion Games. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 771–782. ACM, 2014.