Privacy and Truthful Equilibrium Selection for Aggregative Games

Rachel Cummings[1], Michael Kearns[2], Aaron Roth[2], and Zhiwei Steven Wu[2]

[1] California Institute of Technology
Computing and Mathematical Sciences
rachelc@caltech.edu
[2] University of Pennsylvania
Computer and Information Science
{mkearns,aaroth,wuzhiwei}@cis.upenn.edu

Abstract. We study a very general class of games — multi-dimensional aggregative games — which in particular generalize both anonymous games and weighted congestion games. For any such game that is also *large*, we solve the equilibrium selection problem in a strong sense. In particular, we give an efficient *weak mediator*: a mechanism which has only the power to listen to reported types and provide non-binding suggested actions, such that (a) it is an asymptotic Nash equilibrium for every player to truthfully report their type to the mediator, and then follow its suggested action; and (b) that when players do so, they end up coordinating on a particular asymptotic pure strategy Nash equilibrium of the induced complete information game. In fact, truthful reporting is an *ex-post* Nash equilibrium of the mediated game, so our solution applies even in settings of incomplete information, and even when player types are arbitrary or worst-case (i.e. not drawn from a common prior). We achieve this by giving an efficient differentially private algorithm for computing a Nash equilibrium in such games. The rates of convergence to equilibrium in all of our results are inverse polynomial in the number of players n. We also apply our main results to a multi-dimensional market game.

Our results can be viewed as giving, for a rich class of games, a more robust version of the Revelation Principle, in that we work with weaker informational assumptions (no common prior), yet provide a stronger solution concept (ex-post Nash versus Bayes Nash equilibrium). In comparison to previous work, our main conceptual contribution is showing that weak mediators are a game theoretic object that exist in a wide variety of games – previously, they were only known to exist in traffic routing games. We also give the first weak mediator that can implement an equilibrium optimizing a linear objective function, rather than implementing a possibly worst-case Nash equilibrium.

Keywords: differential privacy, equilibrium computation, mechanism design

The full version of this extended abstract can be found on arXiv [9].

1 Introduction

Games with a large number of players are almost always played, but only sometimes modeled, in a setting of *incomplete information*. Consider, for example, the problem of selecting stocks for a 401k portfolio among the companies listed in the S&P500. Because stock prices are the result of the aggregate decisions of millions of investors, this is a large multi-player strategic interaction, but it is so decentralized that it is implausible to analyze it in a complete information setting (in which every player knows the types or utilities of all of his opponents), or even in a Bayesian setting (in which every agent shares common knowledge of a prior distribution from which player types are drawn). How players will behave in such interactions is unclear; even under settings of complete information, there remains the potential problem of coordinating or selecting a particular equilibrium among many.

One solution to this problem, recently proposed by Kearns et al. [20] and Rogers and Roth [28], is to modify the game by introducing a *weak mediator*, which essentially only has the power to listen and to give advice. Players can ignore the mediator, and play in the original game as they otherwise would have. Alternately, they can use the mediator, in which case they can report their type to it (although they have the freedom to lie). The mediator provides them with a suggested action that they can play in the original game, but they have the freedom to disregard the suggestion, or to use it in some strategic way (not necessarily following it). The goal is to design a mediator such that *good behavior* – that is, deciding to use the mediator, truthfully reporting one's type, and then faithfully following the suggested action – forms an ex-post Nash equilibrium in the mediated game, and that the resulting play forms a Nash equilibrium of the original *complete information* game, induced by the actual (but unknown) player types. A way to approximately achieve this goal – which was shown in Kearns et al. [20], Rogers and Roth [28] – is to design a mediator which computes a Nash equilibrium of the game defined by the reported player types under a stability constraint known as *differential privacy* [12]. Prior to our work, this was only known to be possible in the special case of large, unweighted congestion games.

In this paper, we extend this approach to a much more general class of games known as *multi-dimensional aggregative games* (which among other things, generalize both anonymous games and weighted congestion games). In such a game, there is a vector of linear functions of players' joint actions called an *aggregator*. Each player's utility is then a possibly *non-linear* function of the aggregator vector and their own action. For example, in an investing game, the imbalance between buyers and sellers of a stock, which is a linear function of actions, may be used in the utility functions to compute prices, which are a non-linear function of the imbalances (see the full version). In an anonymous game, the aggregator function represents the number of players playing each action. In a weighted congestion game, the aggregator function represents the total weight of players on each of the facilities. Our results apply to any *large* aggregative game, meaning that any player's unilateral change in action can have at most a bounded influence on the utility of any other player, and the bound on this influence should be a diminishing function in the number of players in the game. Conceptually, our paper is the first to show that *weak mediators* are a game-theoretic object that exists in a large, general class of games: previously, although defined, weak mediators were only known to exist in traffic routing games [28].

This line of work can be viewed as giving robust versions of the Revelation Principle, which can implement Nash equilibria of the complete information game using a "direct revelation mediator," but without needing the existence of a prior type distribution. Compared to the Revelation Principle, which generically requires such a distribution and implements a Bayes Nash equilibrium, truth-telling forms an ex-post Nash equilibrium in our setting. We include a comparison to previous work in Table 1.

Finally, another important contribution of our work is that we are the first to demonstrate the existence of weak mediators (in *any* game) that have the power to optimize over an arbitrary linear function of the actions, and hence able to implement near optimal equilibria under such objective functions, rather than just implementing worst-case Nash equilibria.

Mechanism	**Class of Games**	**Common Prior?**	**Mediator Strength**	**Equilibrium Implemented**
Revelation Principle [24]	Any Finite Game	Yes	Weak	Bayes Nash
Kearns et al. [20]	Any Large Game	No	Strong	Correlated
Rogers and Roth [28]	Large Congestion Games	No	Weak	Nash
This Work	Aggregative Games	No	Weak	Nash

Table 1. Summary of truthful mechanisms for various classes of games and solution concepts. Note that a "weak" mediator does not require the ability to verify player types. A "strong" mediator does. Weak mediators are preferred.

1.1 Our Results and Techniques

Our main result is the existence of a mediator which makes truthful reporting of one's type and faithful following of the suggested action (which we call the "good behavior" strategy) an ex-post Nash equilibrium in the mediated version of any aggregative game, thus implementing a Nash equilibrium of the underlying game of complete information. Unlike the previous work in this line [20, 28], we do not have to implement an arbitrary (possibly worst-case) Nash equilibrium, but can implement a Nash equilibrium which optimizes any linear objective (in the player's actions) of our choosing. We here state our results under the assumption that any player's action has influence bounded by $(1/n)$ on other's utility, but our results hold more generally, parameterized by the "largeness" of the game.

Theorem 1 (Informal). *In a d-dimensional aggregative game of n players and m actions, there exists a mediator that makes good behavior an η-approximate ex-post Nash equilibrium, and implements a Nash equilibrium of the underlying complete information game that optimizes any linear objective function to within η, where*

$$\eta = O\left(\frac{\sqrt{d}}{n^{1/3}} \cdot \text{polylog}(n, m, d)\right).$$

It is tempting to think that the fact that players only have small influence on one another's utility function is sufficient to make any algorithm that computes a Nash equilibrium of the game a suitable weak mediator, but this is not so (see Kearns et al. [20] for an example). What we need out of a mediator is that any single agent's report should have little effect on the *algorithm* computing the Nash equilibrium, rather than on the payoffs of the other players.

The underlying tool that we use is differential privacy, which enforces the stability condition we need on the equilibrium computation algorithm. Our main technical contribution is designing a (jointly) differentially private algorithm for computing approximate Nash equilibria in aggregative games. The algorithm that we design runs in time polynomial in the number of players, but exponential in the dimension of the aggregator function. We note that since aggregative games generalize anonymous games, where the dimension of the aggregator function is the number of actions in the anonymous game, this essentially matches the best known running time for computing Nash equilibria in anonymous games, even non-privately [10]. Computing exact Nash equilibria in these games is known to be PPAD-complete [7]. Recent work of Barman and Ligett [5] showed that the equilibrium selection problem is also hard, even for more general solution concepts — it is NP-hard to compute a coarse correlated equilibrium that achieves a non-trivial approximation to the optimal welfare.

In the process of proving this result, we develop several techniques which may be of independent interest. First, we give the first algorithm for computing equilibria of multi-dimensional aggregative games (efficient for constant dimensional games) even in the absence of privacy constraints — past work in this area has focused on the single dimensional case [4, 19]. Second, in order to implement this algorithm privately, we develop the first technique for solving a certain class of linear programs under the constraint of joint differential privacy.

We also give similar results for a class of one-dimensional aggregative games that permit a more general aggregation function and rely on different techniques, and we show how our main result can be applied to equilibrium selection in a multi-commodity market. The details are deferred to the full version [9].

1.2 Related Work

Conceptually, our work is related to the classic Revelation Principle of Myerson [24], in that we seek to implement equilibrium behavior in a game via a "mediated" direct revelation mechanism. Our work is part of a line, starting with Kearns et al. [20] and continuing with Rogers and Roth [28], that attempts to give a more robust reduction, without the need to assume a prior on types. Kearns et al. [20] showed how to privately compute correlated equilibria (and hence implement this agenda) in arbitrary large games. The private computation of correlated equilibrium turns out to give the desired reduction to a direct revelation mechanism only when the mediator has the power to verify types. Rogers and Roth [28] rectified this deficiency by privately computing Nash equilibria, but their result is limited to large unweighted congestion games. In this paper, we substantially generalize the class of games in which we can privately compute Nash equilibria (and hence solve the equilibrium selection problem with a direct-revelation mediator).

This line of work is also related to "strategyproofness in the large," introduced by Azevedo and Budish [3], which has similar goals. In comparison to this work, we do not require that player types be drawn from a distribution over the type-space, do not require any smoothness condition on the set of equilibria of the game, are algorithmically constructive, and do not require our game to be nearly as large. Generally, their results require the number of agents n to be larger than the size of the action set and the size of the type set. In contrast, we only require n to be as large as the *logarithm* of the number of actions, and require no bound at all on the size of the type space (which can even be infinite).

Our work is also related to the literature on mediators in games [22, 23]. In contrast to our main goal (which is to implement solution concepts of the complete information game in settings of incomplete information), this line of work aims to modify the equilibrium structure of the complete information game. It does so by introducing a mediator, which can coordinate agent actions if they choose to opt in using the mediator. Mediators can be used to convert Nash equilibria into dominant strategy equilibria [22], or implement equilibrium that are robust to collusion [23]. Ashlagi et al. [2] considers mediators in games of incomplete information, in which agents can misrepresent their type to the mediators. Our notion of a mediator is related, but our mediators require substantially less power than the ones from this literature. For example, our mechanisms do not need the power to make payments [22], or the power to enforce suggested actions [23]. Like the mediators of Ashlagi et al. [2], ours are designed to work in settings of incomplete information and so do not need the power to verify agent types — but our mediators are weaker, in that they can only make suggestions (i.e. players do not need to cede control to our weak mediators).

The computation of equilibria in aggregative games (also known as summarization games) was studied in Kearns and Mansour [19], which gave efficient algorithms and learning dynamics converging to equilibria in the 1-dimensional case. Babichenko [4] also studies learning dynamics in this class of games and shows that in the 1-dimensional setting, sequential best response dynamics converge quickly to equilibrium. Our paper is the first to give algorithms for equilibrium computation in the multi-dimensional setting, which generalizes many well studied classes of games, including anonymous games. The running time of our algorithm is polynomial in the number of players n and exponential in the dimension of the aggregation function d, which essentially matches the best known running time for equilibrium computation in anonymous games [10].

We use a number of tools from differential privacy [12], as well as develop some new ones. In particular, we use the advanced composition theorem of Dwork et al. [14], the exponential mechanism from McSherry and Talwar [21], and the sparse vector technique introduced by Dwork et al. [13] (refined in Hardt and Rothblum [16] and abstracted into its current form in Dwork and Roth [11]). We introduce a new technique for solving linear programs under joint differential privacy, which extends a line of work (solving linear programs under differential privacy) initiated by Hsu et al. [17].

Finally, our work relates to a long line of work initiated by McSherry and Talwar [21] using differential privacy as a tool and desideratum in mechanism design. In addition to works already cited, this includes Blum et al. [6], Chen et al. [8], Ghosh and Ligett [15], Kannan et al. [18], Nissim et al. [25, 26], Xiao [29] among others. For a survey of this area see Pai and Roth [27].

2 Model and Preliminaries

2.1 Aggregative Games

Consider an n-player game with action set $\mathcal{A}$ consisting of m actions and a (possibly infinite) type space $\mathcal{T}$ indexing utility functions. Let $\boldsymbol{x} = (x_i, \boldsymbol{x}_{-i})$ denote a strategy profile in which player i plays action x_i and the remaining players play strategy profile $\boldsymbol{x}_{-i}$. Each player i has a utility function, $u\colon \mathcal{T} \times \mathcal{A}^n \to [-1, 1]$, where a player with type t_i experiences utility $u(t_i, \boldsymbol{x})$ when players play according to $\boldsymbol{x}$. When it is clear from context, we will use shorthand and write $u_i(\boldsymbol{x})$ to denote $u(t_i, \boldsymbol{x})$, the utility of player i at strategy profile $\boldsymbol{x}$.

The utility functions in *aggregative games* can be defined in terms of a multi-dimensional *aggregator* function $S\colon \mathcal{A}^n \to [-W, W]^d$, which represents a compact "sufficient statistic" to compute player utilities. In particular, each player's utility function can be represented as a function only of her own action x_i and the aggregator of the strategy profile $\boldsymbol{x}$: $u_i(\boldsymbol{x}) = u_i(x_i, S(\boldsymbol{x}))$. We also assume W to be polynomially bounded by n and m. In aggregative games, the function S_k for each coordinate $k \in [d]$, is an additively separable function: $S_k(\boldsymbol{x}) = \sum_{i=1}^n f_i^k(x_i)$.[3]

Similar to the setting of Kearns and Mansour [19] and Babichenko [4], we focus on γ-*aggregative games*, in which each player has a *bounded influence* on

[3] In the economics literature, aggregative games have more restricted aggregator function: $S_k(\boldsymbol{x}) = \sum_{i=1}^n x_i$. The games we study are more general, and sometimes referred to as *generalized aggregative games.*

the aggregator:

$$\max_i \max_{x_i, x_i' \in \mathcal{A}} \|S(x_i, \boldsymbol{x}_{-i}) - S(x_i', \boldsymbol{x}_{-i})\|_\infty \leq \gamma, \text{ for all } \boldsymbol{x}_{-i} \in \mathcal{A}^{n-1}.$$

That is, the greatest change a player can unilaterally cause to the aggregator is bounded by γ. With our motivation to study large games, we assume γ diminishes with the population size n. We also assume that all utility functions are 1-Lipschitz with respect to the aggregator: for all $x_i \in \mathcal{A}$, $|u_i(x_i, s) - u_i(x_i, s')| \leq \|s - s'\|_\infty$.[4]

For γ-aggregative games, we can express the aggregator more explicitly as

$$S_k(\boldsymbol{x}) = \gamma \sum_{i=1}^{n} f_i^k(x_i),$$

where $f_i^k(x_i)$ is the influence of player i's action x_i on the k-th aggregator function, and also $|f_i^k(x_i)| \leq 1$ for all actions $i \in [n]$ and $x_i \in \mathcal{A}$. Let $f_{ij}^k = f_i^k(a_j)$, where a_j denotes the j-th action in $\mathcal{A}$.

We say that player i is playing an *η-best response* to $\boldsymbol{x}$ if $u_i(\boldsymbol{x}) \geq u_i(x_i', \boldsymbol{x}_{-i}) - \eta$, for all $x_i' \in \mathcal{A}$. A strategy profile $\boldsymbol{x}$ is an *η-pure strategy Nash equilibrium* if all players are playing an η-best response in $\boldsymbol{x}$. We also consider *mixed strategies*, which are defined by probability distributions over the action set. For any profile of mixed strategies, given by a product distribution $\boldsymbol{p}$, we can define expected utility $u_i(\boldsymbol{p}) = \mathbb{E}_{\boldsymbol{x} \sim \boldsymbol{p}} u_i(\boldsymbol{x})$ and the expected aggregator

$$S_k(\boldsymbol{p}) = \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \boldsymbol{p}} S_k(\boldsymbol{x}) = \gamma \sum_{i=1}^{n} \sum_{j=1}^{m} f_{ij}^k \, p_{ij} = \gamma \, \langle f^k, \boldsymbol{p} \rangle. \tag{1}$$

The *support* of a mixed strategy p, denoted $\mathrm{Supp}(\boldsymbol{p}_i)$, is the set of actions that are played with non-zero probabilities. A mixed strategy profile $\boldsymbol{p}$ is a *mixed strategy Nash equilibrium* if $u_i(\boldsymbol{p}) \geq \mathbb{E}_{\boldsymbol{x}_{-i} \sim \boldsymbol{p}_{-i}} u_i(x_i', \boldsymbol{x}_{-i})$ for all $i \in [n]$ and $x_i' \in \mathcal{A}$.

For each aggregator s, we define the *aggregative best response*[5] for player i to s as $\mathrm{BA}_i(s) = \arg\max_{x_i \in A}\{u_i(x_i, s)\}$, breaking ties arbitrarily. We define the *η-aggregative best response set* for player i to s as

$$\eta\text{-}\mathrm{BA}_i(s) = \{x_i \in \mathcal{A} | u_i(x_i, s) \geq \max_{x_i'} u_i(x_i', s) - \eta\}$$

to be the set of all actions that are at most η worse than player i's exact aggregative best response.

[4] Note that the influence that any single player's action has on the utility of others is also bounded by γ. If $\gamma = o(1/n)$, then any player's utility is essentially independent of other players' actions. Therefore, we further assume that $\gamma = \Omega(1/n)$ for the problem to be interesting. This will also simplify some statements.

[5] Sometimes called *best react* [4], and *apparent best response* [19].

Remark 1. Note that best response is played against the other players' actions x_{-i}, but aggregative best response is played against the aggregator value s. Aggregative best response ignores the effect of the player's action on the aggregator, which is bounded by γ; the player reasons about the utility of playing different actions as if the aggregator value were *promised* to be s. Nevertheless, aggregative best response and best response can translate to each other with only an additive loss of γ in the approximation factor. Furthermore, aggregative best responses to different aggregators can translate to each other as long as the corresponding aggregators are close. If $\|s - s'\|_\infty \leq \alpha$, then the actions in η-BA(s) are also in $(\eta + 2\alpha)$-BA(s').

2.2 Mediated Games

We now define games modified by the introduction of a mediator. A mediator is an algorithm $M : (\mathcal{T} \cup \{\perp\})^n \to \mathcal{A}^n$ which takes as input reported types (or $\perp$ for any player who declines to use the mediator), and outputs a suggested action to each player. Given an aggregative game G, we construct a new game G_M induced by the mediator M. Informally, in G_M, players have several options: they can *opt-out* of the mediator (i.e. report $\perp$) and select an action independently of it. Alternately they can *opt-in* and report to it some type (not necessarily their true type), and receive a suggested action r_i. They are free to follow this suggestion or use it in some other way: they play an action $f_i(r_i)$ for some arbitrary function $f_i : \mathcal{A} \to \mathcal{A}$. Formally, the game G_M has an action set $\mathcal{A}_i$ for each player i defined as $\mathcal{A}_i = \mathcal{A}'_i \cup \mathcal{A}''_i$, where

$$\mathcal{A}'_i = \{(t_i, f_i) : t_i \in \mathcal{T}, f_i : \mathcal{A} \to \mathcal{A}\} \quad \text{ and } \mathcal{A}''_i = \{(\perp, f_i) : f_i \text{ is constant}\}.$$

Players' utilities in the mediated game are simply their expected utilities induced by the actions they play in the original game. Formally, they have utility functions u'_i: $u'_i(t, f) = \mathbb{E}_{\boldsymbol{x} \sim M(t)}[u_i(f(\boldsymbol{x}))]$. We are interested in finding mediators such that *good behavior* is an ex-post Nash equilibrium in the mediated game. We first define an ex-post Nash equilibrium.

Definition 1 (Ex-Post Nash Equilibrium). *A collection of strategies* $\{\sigma_i : \mathcal{T} \to \mathcal{A}_i\}_{i=1}^n$ *forms an* η*-approximate ex-post Nash equilibrium if for every type vector* $t \in \mathcal{T}^n$*, and for every player* i *and action* $x_i \in \mathcal{A}_i$*:*

$$u'_i(\sigma_i(t_i), \sigma_{-i}(t_{-i})) \geq u'_i(x_i, \sigma_{-i}(t_{-i})) - \eta$$

That is, it forms an η*-approximate Nash equilibrium for* every *possible vector of types.*

Note that ex-post Nash equilibrium is a very strong solution concept for incomplete information games because it does not require players to know a prior distribution over types.

In a mediated game, we would like players to truthfully report their type, and then faithfully follow the suggested action of the mediator. We call this *good behavior*. Formally, the good behavior strategy is defined as $g_i(t_i) = (t_i, \mathrm{id})$ where $\mathrm{id} : \mathcal{A} \to \mathcal{A}$ is the identity function – i.e. it truthfully reports a player's type to the mediator, and applies the identity function to its suggested action.

In order to achieve this, we use the notion of *joint differential privacy* defined in Kearns et al. [20] (adapted from *differential privacy*, defined in Dwork et al. [12]), as a privacy measure for mechanisms on agents' private data (types). Intuitively, it guarantees that the output to all other agents excluding player i is insensitive to i's private type, so the mechanism protects i's private information from arbitrary coalitions of adversaries.

Definition 2 (Joint Differential Privacy [20]). *Two type profiles t and t' are i-neighbors if they differ only in the i-th component. An algorithm $\mathcal{M} : \mathcal{T}^n \to \mathcal{A}^n$ is (ε, δ)*-joint differentially private *if for every i, for every pair of i-neighbors $t, t' \in \mathcal{T}^n$, and for every subset of outputs $\mathcal{S} \subseteq \mathcal{A}^{n-1}$,*

$$\Pr[\mathcal{M}(t)_{-i} \in \mathcal{S}] \le \exp(\varepsilon) \Pr[\mathcal{M}(t')_{-i} \in \mathcal{S}] + \delta.$$

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

We here quote a theorem of Rogers and Roth [28], inspired by Kearns et al. [20] which motivates our study of private equilibrium computation.

Theorem 2 ([20, 28]). *Let M be a mechanism satisfying (ε, δ)-joint differential privacy, that on any input type profile t with probability $1 - \beta$ computes an α-approximate pure strategy Nash equilibrium of the complete information game $G(t)$ defined by type profile t. Then the "good behavior" strategy $g = (g_1, \ldots, g_n)$ forms an η-approximate ex-post Nash equilibrium of the mediated game G_M for*

$$\eta = \alpha + 2(2\varepsilon + \beta + \delta).$$

Our private equilibrium computation relies on two private algorithmic tools, sparse vector mechanism (called Sparse) and exponential mechanism (called EXP), which allows us to access agents' types in a privacy-preserving manner.

3 Private Equilibrium Computation

Let G be a d-dimensional γ-aggregative game, and $L\colon \mathcal{A}^n \to \mathbb{R}$ be a γ-Lipschitz linear loss function:

$$L(\boldsymbol{x}) = \gamma \sum_i \ell_i(x_i) \quad \text{and} \quad L(\boldsymbol{p}) = \gamma \mathop{\mathbb{E}}_{\boldsymbol{x} \sim \boldsymbol{p}} L(\boldsymbol{x}) = \gamma \sum_i \langle p_{ij}, \ell_{ij} \rangle.$$

where $0 \le \ell_i(a_j) \le 1$ for all actions $a_j \in \mathcal{A}$, and $\ell_{ij} = \ell_i(a_j)$.

Given any $\zeta \ge \gamma\sqrt{8n \log(2mn)}$, let $\mathcal{E}(\zeta)$ be the set of ζ-approximate pure strategy Nash equilibria in the game G,[6] and let

$$\mathrm{OPT}(\zeta) = \min\{L(\boldsymbol{x}) \mid \boldsymbol{x} \in \mathcal{E}(\zeta)\}.$$

We give the following main result:

[6] We show that $\mathcal{E}(\zeta)$ is non-empty for $\zeta \ge \gamma\sqrt{8n \log(2mn)}$ in the full version.

Theorem 3. *For any $\zeta \geq \gamma\sqrt{8n\log(2mn)}$, there exists a mediator M that makes good behavior an $(\zeta + \eta)$-approximate ex-post Nash equilibrium of the mediated game G_M, and implements an approximate pure strategy Nash equilibrium $\boldsymbol{x}$ of the underlying complete information game with $L(\boldsymbol{x}) \leq \mathrm{OPT}(\zeta) + \eta$, where*

$$\eta = O\left(n^{1/3}\gamma^{2/3}\sqrt{d}\cdot \mathrm{polylog}(n, m, d)\right).$$

Recall that the quantity γ is diminishing in n; whenever $\gamma = O(1/n^{1/2+\varepsilon})$ for $\varepsilon > 0$, the approximation factor η tends towards zero as n grows large. Plugging in $\gamma = 1/n$ and $\zeta = \gamma\sqrt{8n\log(2mn)}$ recovers the bound in Theorem 1.

This result follows from instantiating Theorem 2 with an algorithm that computes an approximate equilibrium under joint differential privacy as PRESL (Private Equilibrium Selection).[7] We give here an informal description of our algorithm, absent privacy concerns, and then describe how we implement it privately, deferring the formal treatment to the full version.

The main object of interest in our algorithm is the set-valued function

$$\mathcal{V}_\xi(\widehat{s}) = \{S(\boldsymbol{p}) \mid \text{for each } i, \mathrm{Supp}(p_i) \subseteq \xi\text{-BA}_i(\widehat{s})\},$$

which maps aggregator values $\widehat{s}$ to the set of aggregator values that arise when players are randomizing between ξ-aggregative best responses to $\widehat{s}$. An approximate equilibrium will yield an aggregator $\widehat{s}$ such that $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$, so we wish to find such a fixed point for $\mathcal{V}_\xi$ (the value of ξ will be determined in the analysis, see the full version). Note that *pure strategy* Nash equilibria correspond to such fixed points, but a-priori, it is not clear that fixed points of this function (which may involve mixed strategies) are mixed strategy Nash equilibria. This is because player utility functions need not be linear in the aggregator, and so a best response to the expected value of the aggregator need not be a best response to the corresponding distribution over aggregators. However, as we will show, we can safely round such fixed points to approximate pure strategy Nash equilibria, because the aggregator will be well concentrated under rounding.

For every fixed value $\widehat{s}$, the problem of determining whether $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ is a linear program (because the aggregator is linear), and although $\mathrm{Supp}(p_i) \subseteq \xi\text{-BA}_i(\widehat{s})$ is not a convex constraint in $\widehat{s}$, the aggregative best responses are fixed for each fixed value of $\widehat{s}$. The first step of our algorithm simply searches through a discretized grid of all possible aggregators $X = \{-W, -W + \alpha, \ldots, W - \alpha\}^d$, and solves this linear program to check if some point $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$. This results in a set of aggregators S that are induced by the approximate equilibria of the game. Let p_{ij} denote the probability that player i plays the j-th action. Then

[7] In the full version of this paper, we also present details of the non-private algorithm to compute equilibrium for aggregative games.

the linear program we need to solve is as follows:

$$
\begin{aligned}
&\forall k \in [d], && \widehat{s}_k - \alpha \le \gamma \sum_{i=1}^{n} \sum_{j=1}^{m} f_{ij}^{k} p_{ij} \le \widehat{s}_k + \alpha \\
&\forall i \in [n], && \forall j \in \xi\text{-BA}_i(\widehat{s}), \qquad 0 \le p_{ij} \le 1 \\
&\forall i \in [n], && \forall j \notin \xi\text{-BA}_i(\widehat{s}), \qquad p_{ij} = 0
\end{aligned}
\tag{2}
$$

Next, we need to find a particular equilibrium (an assignment of actions to players) that optimizes our cost-objective function L. This is again a linear program (since the objective function is linear) for each $\widehat{s}$. Hence, for each fixed point $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ we simply solve this linear program, and out of all of the candidate equilibria, output the one with the lowest cost. Finally, this results in mixed strategies for each of the players, and we round this to a pure strategy Nash equilibrium by sampling from each player's mixed strategy. This does not substantially harm the quality of the equilibrium; because of the low sensitivity of the aggregator, it is well concentrated around its expectation under this rounding. The running time of this algorithm is dominated by the grid search for the aggregator fixed point $\widehat{s}$, which takes time exponential in d. Solving each linear program can be done in time polynomial in all of the game parameters.

Making this algorithm satisfy joint differential privacy is more difficult. There are two main steps. The first is to identify the fixed point $\widehat{s} \in \mathcal{V}_\xi(\widehat{s})$ that corresponds the lowest cost equilibrium. There are exponentially in d many candidate aggregators to check, and with naive noise addition we would have to pay for this exponential factor in our accuracy bound. However, we take advantage of the fact that we only need to *output* a single aggregator – the one corresponding to the lowest objective value equilibrium – and so the *sparse vector mechanism* Sparse (described in the full version) can be brought to bear, allowing us to pay only linearly in d in the accuracy bound.

The second step is more challenging, and requires a new technique: we must actually solve the linear program corresponding to $\widehat{s}$, and output to each player the strategy they should play in equilibrium. The output strategy profile must satisfy joint differential privacy. To do this, we give a general method for solving a class of linear programs (containing in particular, LPs of the form (2)) under joint differential privacy, which may be of independent interest. This algorithm, which we call DistMW (described in the full version), is a distributed version of the classic multiplicative weights (MW) technique for solving LPs [1]. The algorithm can be analyzed by viewing each agent as controlling the variables corresponding to their own mixed strategies, and performing their multiplicative weights updates in isolation (and ensuring that their mixed strategies always fall within their best response set ξ-BA$_i(\widehat{s})$). At every round, the algorithm aggregates the current solution maintained by each player, and then identifies a coordinate in which the constraints are far from being satisfied. The algorithm uses the *exponential mechanism* EXP (described in the full version) to pick such a coordinate while maintaining the privacy of the players' actions. The identification of such a coordinate is sufficient for each player to update their own variables. Privacy then follows by combining the privacy guarantee of the exponential mechanism with a bound on the convergence time of the multiplicative weights update rule. The fact that we can solve this LP in a distributed manner to get joint differential privacy (rather than standard differential privacy) crucially depends on the fact that the sensitivity γ of the aggregator is small. The algorithm DistMW will find a set of strategies that approximately satisfy the linear program – the violation on each coordinate is bounded by

$$E = O\left(\frac{n\gamma^2}{\varepsilon} \operatorname{polylog}\left(n, m, d, \frac{1}{\beta}, \frac{1}{\delta}\right)\right)^{1/2}.$$

The algorithm PRESL has the following guarantee:

Theorem 4. *Let $\zeta \geq \gamma\sqrt{8n\log(2mn)}, \varepsilon, \delta, \beta \in (0,1)$. $\mathsf{PRESL}(t, \zeta, L, \varepsilon, \delta, \beta)$ satisfies $(2\varepsilon, \delta)$-joint differential privacy, and, with probability at least $1-\beta$, computes a $(\zeta + 12\alpha)$-approximate pure strategy equilibrium $\boldsymbol{x}$ such that $L(\boldsymbol{x}) < \mathrm{OPT}(\zeta) + 5\alpha$, where*

$$\alpha = O\left(\frac{(\sqrt{n\varepsilon} + d)\,\gamma}{\varepsilon} \operatorname{polylog}(n, m, d, 1/\beta, 1/\delta)\right).$$

We defer the full proof and technical details to the full version.

Remark 2. The running time of this algorithm is exponential in d, the dimension of the aggregative game. For games of fixed dimension (where d is constant), this yields a polynomial time algorithm. This exponential dependence on the dimension matches the best known running time for (non-privately) computing equilibrium in anonymous games by [10], which is a sub-class of aggregative games.

Theorem 3 follows by instantiating Theorem 2 with $\mathsf{PRESL}\left(t, \zeta, L, n^{1/3}\gamma^{2/3}d^{1/2}, \frac{1}{n}, \frac{1}{n}\right)$ – i.e. by setting $\varepsilon = n^{1/3}\gamma^{2/3}d^{1/2}$ and $\delta = \beta = \frac{1}{n}$.

Acknowledgments. Research supported in part by NSF grants 1253345, 1101389, CNS-1254169, US-Israel Binational Science Foundation grant 2012348, Simons Foundation grant 361105, a Google Faculty Research Award, and the Alfred P. Sloan Foundation. Research performed while the first author was visiting the University of Pennsylvania.

Bibliography

[1] Sanjeev Arora, Elad Hazan, and Satyen Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing*, 8(1):121–164, 2012.

[2] Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz. Mediators in position auctions. *Games and Economic Behavior*, 67(1):2–21, 2009.

[3] Eduardo M Azevedo and Eric Budish. Strategyproofness in the large as a desideratum for market design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, page 55, 2012.

[4] Yakov Babichenko. Best-reply dynamic in large aggregative games. *SSRN abstract 2210080*, 2013.

[5] Siddharth Barman and Katrina Ligett. Finding any nontrivial coarse correlated equilibrium is hard. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 815–816, 2015.

[6] Avrim Blum, Jamie Morgenstern, Ankit Sharma, and Adam Smith. Privacy-preserving public information for sequential games. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science*, ITCS '15, pages 173–180, 2015.

[7] Xi Chen, David Durfee, and Anthi Orfanou. On the complexity of Nash equilibria in anonymous games. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, STOC '15, pages 381–390, 2015.

[8] Yiling Chen, Stephen Chong, Ian A Kash, Tal Moran, and Salil Vadhan. Truthful mechanisms for agents that value privacy. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 215–232, 2013.

[9] Rachel Cummings, Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. Privacy and truthful equilibrium selection for aggregative games. *CoRR*, abs/1407.7740, 2014.

[10] Constantinos Daskalakis and Christos H. Papadimitriou. Discretized multinomial distributions and Nash equilibria in anonymous games. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '08, pages 25–34, 2008.

[11] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4): 211–407, 2014.

[12] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Conference on Theory of Cryptography*, TCC '06, pages 265–284, 2006.

[13] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil Vadhan. On the complexity of differentially private data release: Efficient algorithms and hardness results. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*, STOC '09, pages 381–390, 2009.

[14] Cynthia Dwork, Guy N. Rothblum, and Salil Vadhan. Boosting and differential privacy. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '10, pages 51–60, 2010.

[15] Arpita Ghosh and Katrina Ligett. Privacy and coordination: Computing on databases with endogenous participation. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, EC '13, pages 543–560, 2013.

[16] Moritz Hardt and Guy N. Rothblum. A multiplicative weights mechanism for privacy-preserving data analysis. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, FOCS '10, pages 61–70, 2010.

[17] Justin Hsu, Aaron Roth, Tim Roughgarden, and Jonathan Ullman. Privately solving linear programs. In *Automata, Languages, and Programming*, volume 8572 of *Lecture Notes in Computer Science*, pages 612–624. 2014.

[18] Sampath Kannan, Jamie Morgenstern, Aaron Roth, and Zhiwei Steven Wu. Approximately stable, school optimal, and student-truthful many-to-one matchings (via differential privacy). In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 1890–1903, 2015.

[19] Michael Kearns and Yishay Mansour. Efficient Nash computation in large population games with bounded influence. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence*, UAI '02, pages 259–266, 2002.

[20] Michael Kearns, Mallesh Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, ITCS '14, pages 403–410, 2014.

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

[22] Dov Monderer and Moshe Tennenholtz. k-Implementation. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, EC '03, pages 19–28, 2003.

[23] Dov Monderer and Moshe Tennenholtz. Strong mediated equilibrium. *Artificial Intelligence*, 173(1):180–195, 2009.

[24] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[25] Kobbi Nissim, Claudio Orlandi, and Rann Smorodinsky. Privacy-aware mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 774–789, 2012.

[26] Kobbi Nissim, Rann Smorodinsky, and Moshe Tennenholtz. Approximately optimal mechanism design via differential privacy. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, ITCS '12, pages 203–213, 2012.

[27] Mallesh M Pai and Aaron Roth. Privacy and mechanism design. *SIGecom Exchanges*, 12(1):8–29, 2013.

[28] Ryan M. Rogers and Aaron Roth. Asymptotically truthful equilibrium selection in large congestion games. In *Proceedings of the 15th ACM Conference on Economics and Computation*, EC '14, pages 771–782, 2014.

[29] David Xiao. Is privacy compatible with truthfulness? In *Proceedings of the 4th Conference on Innovations in Theoretical Computer Science*, ITCS '13, pages 67–86, 2013.