CIS 700 Differential Privacy in Game Theory and Mechanism Design March 21, 2014

Lecture 8

Lecturer: Aaron Roth *Scribe: Aaron Roth*

Private Equilibrium Computation Yields Truthful Mediators (Part 3)

1 Introduction

Ok, last lecture on mediators! Lets briefly recall what we have shown so far:

1. We proved that in any game for which there exists a jointly-differentially private algorithm for computing an approximate Nash equilibrium, there exists a corresponding mediated game in which "good behavior" (using the mediator, truthfully reporting your type to it, and then faithfully following its suggestion) forms an ex-post Nash equilibrium, and that the resulting play forms a Nash equilibrium of the original game.

2. We instantiated this theorem for the special case of *large congestion games*: we gave a private algorithm for computing approximate Nash equilibrium.

How far can we continue to drive this agenda? One immediate roadblock that we will face is the dearth of good algorithms for computing (arbitrarily good approximations to) Nash equilibria, outside of the class of congestion games. In general, the problem of computing a(n arbitrarily close approximation to a) Nash equilibrium is known to be PPAD complete, which means that there is likely no efficient algorithm which works in the worst case for general games[1].

Instead, this lecture we will think about what the implications are for mediators which privately compute approximate *correlated* equilibria. Correlated equilibria have the benefit that there are efficient algorithms to compute (arbitrarily good approximations of) them in any n player game, and indeed, we shall be able to give algorithms that compute asymptotic correlated equilibria in arbitrary large games.

A correlated equilibrium is, informally, a joint distribution $\mathcal{D}$ on action profiles $a \in A^n$ such that when $a \sim \mathcal{D}$, a_i is always a best response for player i *even conditioned on seeing the realization* of a_i, under the assumption that i's opponents will play according to a_{-i}. Formally:

Definition 1 *A distribution $\mathcal{D} \in \Delta A^n$ is an η-approximate correlated equilibrium if for every player i and for every function $f : A \to A$:*

$$\mathrm{E}_{a\sim\mathcal{D}}[c_i(a)] \leq \mathrm{E}_{a\sim\mathcal{D}}[c_i(f(a_i), a_{-i})] + \eta$$

Note that the function f can be viewed as a hypothetical deviation that allows player i to deviate from his suggested action a_i even after he sees it. A correlated equilibrium makes all such deviations unprofitable. Note that a Nash equilibrium is just a correlated equilibrium in which $\mathcal{D}$ must be a product distribution (i.e. in that case, the marginal distribution on each player's actions a_i is just their mixed strategy).

A related, weaker solution concept we will discuss is a *coarse correlated equilibrium*. Informally, a coarse correlated equilibrium $\mathcal{D} \in \Delta A^n$ guarantees that for $a \sim \mathcal{D}$, playing a_i is a best response for player i in expectation *before* seeing the realization of a_i.

Definition 2 *A distribution $\mathcal{D} \in \Delta A^n$ is an η-approximate coarse-correlated equilibrium if for every player i and for every action $a'_i \in A$*

$$\mathrm{E}_{a\sim\mathcal{D}}[c_i(a)] \leq \mathrm{E}_{a\sim\mathcal{D}}[c_i(a'_i, a_{-i})] + \eta$$

[1] This does *not necessarily* mean that there is no jointly-differentially private algorithm for computing asymptotic Nash equilibria in arbitrary large games, since such algorithms need not be efficient! Resolving this question would be extremely interesting, because of its game theoretic applications. However, it does make it difficult to think about what such an algorithm would look like, since exponential time algorithms also tend to "touch the data" exponentially many times, which makes them difficult to make private...

We will prove a theorem about the mediator that results, given an algorithm that can privately compute a correlated equilibrium. We will state a theorem about privately computing correlated equilibria, but for simplicity, in this lecture we will only give a private algorithm for computing a coarse correlated equilibrium which is somewhat simpler.

Informally, the mediator that results from an algorithm which privately computes a correlated equilibrium is qualitatively weaker than the mediators we discussed in the last two lectures in the following respect: players will still be able to opt-out of using the mediator, and play the game on their own, and they will still be free to deviate from the mediator's suggestion once they see it. However, now we will assume that players do not have the ability to misreport their type to the mediator if they choose to use it. This is equivalent to saying that the mediator now has the power to verify types of players who opt-in to using it. Not all mediators will have this power of course, but it may be reasonable to assume in e.g. regulated markets in which misreports are punishable by law.

2 Notational Recap

Suppose $\mathcal{G}$ is a game with action set A, type set $\mathcal{T}$, and cost function $c_i(a) \equiv c(t_i, a)$.

Definition 3 *A* mediator *is given by an algorithm $M : (\mathcal{T} \cup \{\perp\})^n \to (A \cup \{\perp\})^n$ mapping reported types (or $\perp$, which denotes declining to report any type) to suggested actions (or $\perp$, which denotes not suggesting any action).*

Given a game $\mathcal{G}$ and a mediator M, we can define the game $\mathcal{G}$ augmented with mediator M, denoted as $\mathcal{G}_M$, as follows. The new game has action space $A'(t_i)$ for players with type t_i:

$$A'(t_i) = \{(t', f) : t' \in \{t_i, \perp\}, f : (A \cup \{\perp\}) \to A\}$$

i.e. each player chooses whether to opt-in and report t_i, or opt-out and report $\perp$, and a function f specifying how to map the advice received by the mediator to an action a to play in the underlying game $\mathcal{G}$ (Note that picking f equal to the identity function corresponds to following the advice of the mediator, but the player can pick anything else, including a constant function which corresponds to completely disregarding the advice. Note also that if a player reports $t' = \perp$ to the mediator, then it will return the advice $\perp$, and a player will then play the action $f(\perp) \in A$, which encodes the strategy he would have used in the original game). Note also that this is a more restricted action space than we considered in previous lectures: we are now restricting players to report their true type if they opt-in to using the mediator. The new game has the following cost function:

$$c'(t_i, ((t'_1, f_1), \ldots, (t'_n, f_n))) = \mathrm{E}_{a \sim M(t')}[c(t_i, f(a))]$$

which simply corresponds to their cost in the original game when they pick actions using the mediator and their function f.

We wanted to implement "good behavior" as an ex-post equilibrium of $\mathcal{G}_M$. Recall that we defined the "good behavior" strategy $g_i = (t_i, \mathrm{id})$ where id is the identity function, which simply corresponds to opting in, and then faithfully following the suggested action of the mediator.

Informally, a *large* game is one in which the unilateral deviation of some player $j \neq i$ can have only a small effect on the cost of agent i. Note that if agent i changes his own action, his cost is allowed to change drastically – his cost is merely insensitive to the actions of single *other* players.

Definition 4 *Fix a game $\mathcal{G}$ with action set A and cost functions $c_i : A^n \to \mathbb{R}$. Say that $\mathcal{G}$ is* large *with sensitivity $\Delta_{\mathcal{G}}$ if for every pair of players $i \neq j$, every vector of actions $a \in A^n$ and every deviation a'_j for player j:*

$$|c_i(a) - c_i(a'_j, a_{-j})| \leq \Delta_{\mathcal{G}}$$

We will be interested in large games in which $\Delta_{\mathcal{G}}$ is a diminishing quantity with the number of players n.

3 Privately Computing a Correlated Equilibrium Yields a Good Mediator

We will here introduce a relaxation of differential privacy that allows not only for a multiplicative $\exp(\epsilon)$ change in probabilities of outcomes between neighboring databases, but also an additive δ change in probabilities.

Definition 5 *A mechanism $M : \mathcal{T}^n \to \mathcal{O}$ is (ϵ, δ)-differentially private if for every $t \in \mathcal{T}^n$, $t'_i \in \mathcal{T}$, and every $S \subseteq \mathcal{O}$:*

$$\Pr[M(t) \in S] \leq \exp(\epsilon) \Pr[M(t'_i, t_{-i}) \in S] + \delta$$

A mechanism $M : \mathcal{T}^n \to A^n$ is (ϵ, δ)-jointly differentially private if for every i, for every $t \in \mathcal{T}^n$, $t'_i \in \mathcal{T}$, and every $S_{-i} \subseteq A^{n-1}$:

$$\Pr[M(t)_{-i} \in S_{-i}] \leq \exp(\epsilon) \Pr[M(t'_i, t_{-i})_{-i} \in S_{-i}] + \delta$$

We now show that any mediator which computes a correlated equilibrium under the constraint of (ϵ, δ)-joint differential privacy can be used to implement "good behavior" as an ex-post Nash equilibrium of the mediated game. Note that in our previous lectures, we could have used (ϵ, δ)-differential privacy as well – but we will see that this time we need to make crucial use of it when designing our algorithm.

Theorem 6 *Let $\mathcal{G}$ be any game with costs bounded in $[0, 1]$, and let M be a mechanism satisfying (ϵ, δ)-joint differential-privacy such that on any vector of reported types t' outputs a vector of actions drawn from an η-approximate correlated equilibrium of the complete information game induced by t'. Then good behavior $g = (g_1, \ldots, g_n)$ forms an η'-approximate ex-post Nash equilibrium of $\mathcal{G}_M$ for:*

$$\eta' = 2\epsilon + \delta + \eta.$$

Proof Fix any vector of player types t. We will show that g is a pure strategy Nash equilibrium for the complete information game $\mathcal{G}_M$ with types t. We have only two possible kinds of deviations to consider. First, consider deviations of the form $b_i = (t_i, f)$, where player i continues to opt in to using the mediator (recall that he cannot lie about his type in this case), but interprets the mediator's advice using some function f other than the identity function. In this case we can compute:

$$\begin{aligned} c'(t_i, g) &= \mathrm{E}_{a \sim M(t)}[c_i(a)] \\ &\leq \mathrm{E}_{a \sim M(t)}[c_i(f(a_i), a_{-i})] + \eta \\ &= c'(t_i, (b_i, g_{-i})) + \eta \end{aligned}$$

Here the inequality follows from the fact that M computes an η-approximate correlated equilibrium.

The other kind of deviation we must consider is $b_i = (\perp, f)$ in which agent i opts out of using the mediator. Recall that in this case, the mediator will not make any suggestion: i.e. we will have $a_i = \perp$. Let $f(\perp) = a'_i$. In this case we have:

$$\begin{aligned} c'(t_i, g) &= \mathrm{E}_{a \sim M(t)}[c_i(a)] \\ &\leq \mathrm{E}_{a \sim M(t)}[c_i(a'_i, a_{-i})] + \eta \\ &\leq \exp(\epsilon)\mathrm{E}_{a \sim M(\perp, t_{-i})}[c_i(a'_i, a_{-i})] + \delta + \eta \\ &\leq \mathrm{E}_{a \sim M(\perp, t_{-i})}[c_i(a'_i, a_{-i})] + 2\epsilon + \delta + \eta \\ &= c'(t_i, (b_i, g_{-i})) + 2\epsilon + \delta + \eta \end{aligned}$$

■

4 Some Tools

Ok: we now know what we have to do to finish implementing our agenda – we need to give a jointly differentially private algorithm for computing correlated equilibria. Before we can do that though, we need two tools. One is a powerful composition theorem in differential privacy, and the other is an algorithm for computing correlated equilibria in general games (In fact, we will only see how to compute coarse correlated equilibria, but with more care the algorithm we give can be used as a subroutine in a slightly more complicated algorithm for computing approximate correlated equilibria...)

4.1 A Composition Theorem

Suppose we design an algorithm M by adaptively combining multiple differentially private subroutines. What can we say about the privacy properties of M? Formally, suppose that the output of M can be written as $M(t) = y_1, \ldots, y_k$ where $y_1 = Y_1(t)$, $y_2 = Y_2(t; y_1), \ldots, y_k = Y_k(t; y_1, \ldots, y_{k-1})$ where $Y_1, \ldots, Y_k$ are all ϵ-differentially private algorithms. The notation $Y_j(t; y_1, \ldots, y_{j-1})$ denotes that the mechanism Y_j is allowed to be parameterized by the realized outcomes of $Y_1, \ldots, Y_{j-1}$ – i.e. the computation can take place sequentially, and can be adaptively chosen based on the outcomes of the private computations so far. We have already claimed (and it is not hard to see) that in this situation, we can say that M is $k\epsilon$-differentially private. Here, we will see that we can make a stronger claim if we weaken our privacy solution concept to (ϵ, δ)-differential privacy.

First, some notation:

Definition 7 (Max Divergence) *The Max Divergence between two random variables Y and Z taking values from the same domain is defined to be:*

$$D_\infty(Y||Z) = \max_{S \subset Supp(Y)} \left[\ln \frac{\Pr[Y \in S]}{\Pr[Z \in S]} \right]$$

The δ-Approximate Max Divergence between Y and Z is defined to be:

$$D_\infty^\delta(Y||Z) = \max_{S \subset Supp(Y): \Pr[Y \in S] \geq \delta} \left[\ln \frac{\Pr[Y \in S] - \delta}{\Pr[Z \in S]} \right]$$

We remark that in this notation, a mechanism M is ϵ-differentially private if and only if for all $t \in \mathcal{T}^n$, and for all $t_i' \in \mathcal{T}$,

$$D_\infty(M(t)||M(t_i', t_{-i})) \leq \epsilon \text{ and } D_\infty(M(t_i', t_{-i})||M(t)) \leq \epsilon$$

Similarly, a mechanism M is (ϵ, δ)-differentially private if and only if for all $t \in \mathcal{T}^n$, and for all $t_i' \in \mathcal{T}$,

$$D_\infty^\delta(M(t)||M(t_i', t_{-i})) \leq \epsilon \text{ and } D_\infty^\delta(M(t_i', t_{-i})||M(t)) \leq \epsilon$$

We can also define an average case notion of divergence:

Definition 8 (KL-Divergence) *The KL-Divergence or Relative Entropy between two random variables Y and Z taking values from the same domain is defined to be:*

$$D(Y||Z) = \mathrm{E}_{y \sim Y} \left[\ln \frac{\Pr[Y = y]}{\Pr[Z = y]} \right]$$

If we think of differential privacy as bounding "worst case privacy loss over all possible events", we can think of KL-Divergence as bounding "average case privacy loss". Our composition argument will proceed as follows: we will argue that the composition of k ϵ-differentially private mechanisms has substantially smaller expected privacy loss than worst-case privacy loss (which is $k\epsilon$). We will then argue that with high probability (i.e. probability $\geq 1 - \delta$), the realized privacy loss is not too much larger than the expected privacy loss. For this second argument, we will need a tool from probability theory telling us how far we should expect certain kinds of random variables to deviate from their expectation:

Theorem 9 (Azuma's Inequality) *Let f be a function of m random variables $X_1, \ldots, X_m$, each X_i taking values from a set A_i such that $\mathrm{E}[f]$ is bounded. Let c_i denote the maximum effect of X_i on f – i.e. for all $a_i, a_i' \in A_i$:*

$$|\mathrm{E}[f|X_1, \ldots, X_{i-1}, X_i = a_i] - \mathrm{E}[f|X_1, \ldots, X_{i-1}, X_i = a_i']| \leq c_i$$

Then:

$$\Pr\left[f(X_1, \ldots, X_m) \geq \mathrm{E}[f] + t\right] \leq \exp\left(-\frac{2t^2}{\sum_{i=1}^m c_i^2}\right)$$

We can now state our composition theorem for differentially private algorithms:

Theorem 10 ([DRV10]) *Let $M : \mathcal{T}^n \to \mathcal{O}$ be an algorithm such that the output of M can be written as $M(t) = y_1, \ldots, y_k$ where $y_1 = Y_1(t)$, $y_2 = Y_2(t; y_1), \ldots, y_k = Y_k(t; y_1, \ldots, y_{k-1})$ where $Y_1, \ldots, Y_k$ are all ϵ-differentially private algorithms (for all values of their parameters $y_1, \ldots, y_k$). Then for any δ, M is (ϵ', δ)-differentially private for:*

$$\epsilon' = \sqrt{2k\ln(1/\delta)}\epsilon + k\epsilon(e^\epsilon - 1)$$

The next lemma, which will be important in proving Theorem 10, says that if the *maximum* privacy loss of an algorithm is bounded by $\exp(\epsilon)$, then the *expected* privacy loss is actually quite a bit lower. Together with Azuma's inequality, this will allow us to prove the composition theorem: except with small probability δ, the total privacy loss from the composition of k ϵ-differentially private mechanisms is not much more than the expected privacy loss of that composition, which scales more like $\sqrt{k}$ than like k.

Lemma 11 *Suppose that random variables Y and Z satisfy $D_\infty(Y||Z) \leq \epsilon$ and $D_\infty(Z||Y) \leq \epsilon$. Then $D(Y||Z) \leq \epsilon(e^\epsilon - 1)$.*

Proof We know that for any Y and Z it is the case that $D(Y||Z) \geq 0$ (Relative entropy is non-negative – look up the log-sum inequality!), and so it suffices to bound $D(Y||Z) + D(Z||Y)$. We get:

$$\begin{aligned}
D(Y||Z) &\leq D(Y||Z) + D(Z||Y) \\
&= \sum_y \Pr[Y = y] \cdot \left(\ln \frac{\Pr[Y = y]}{\Pr[Z = y]} + \ln \frac{\Pr[Z = y]}{\Pr[Y = y]}\right) \\
&\quad + (\Pr[Z = y] - \Pr[Y = y]) \cdot \left(\ln \frac{\Pr[Z = y]}{\Pr[Y = y]}\right) \\
&\leq \sum_y [0 + |\Pr[Z = y] - \Pr[Y = y]| \cdot \epsilon] \\
&= \epsilon \cdot \sum_y [\max\{\Pr[Y = y], \Pr[Z = y]\} - \min\{\Pr[Y = y], \Pr[Z = y]\}] \\
&\leq \epsilon \cdot \sum_y [(e^\epsilon - 1) \cdot \min\{\Pr[Y = y], \Pr[Z = y]\}] \\
&\leq \epsilon \cdot (e^\epsilon - 1).
\end{aligned}$$

■

We can now finish the proof of the theorem. The idea is that the expected privacy loss after k differentially private algorithms are run is bounded by the above lemma, and that with high probability, the total privacy loss is not much higher (By Azuma's inequality).

Proof [Proof of Theorem 10] First we will define a "bad" event which we hope will occur only with small probability. Fix any $t \in \mathcal{T}^n$ and $t'_i \in \mathcal{T}$. Write $t' = (t'_i, t_{-i})$ for shorthand. Let:

$$B = \{y \in \mathcal{O} : \Pr[M(t) = y] > e^{\epsilon'} \cdot \Pr[M(t') = y]\}.$$

We will show that $\Pr[M(t) \in B] \leq \delta$, and hence for every set S, we have

$$\Pr[M(t) \in S] \leq \Pr[M(t) \in B] + \Pr[M(t) \in (S \setminus B)] \leq \delta + e^{\epsilon'} \cdot \Pr[M(t') \in S].$$

This is equivalent to saying that $D^{\delta}_{\infty}(M(t)||M(t')) \leq \epsilon'$.

It remains to show $\Pr[M(t) \in B] \leq \delta$. Let random variable $Y^0 = (Y^0_1, \ldots, Y^0_k)$ denote the random variable describing outcomes $y = y_1, \ldots, y_k$ drawn from $M(t)$ and let $Y^1 = (Y^1_1, \ldots, Y^1_k)$ denote the random variable describing outcomes $y = y_1, \ldots, y_k$ drawn from $M(t')$. Then, for a fixed outcome y we have:

$$\begin{aligned}
\ln\left(\frac{\Pr[M(t) = y]}{\Pr[M(t') = y]}\right) &= \ln\left(\prod_{i=1}^{k} \frac{\Pr[Y^0_i = y_i | Y^0_1 = y_1, \ldots, Y^0_{i-1} = y_{i-1}]}{\Pr[Y^1_i = y_i | Y^1_1 = y_1, \ldots, Y^1_{i-1} = y_{i-1}]}\right) \\
&= \sum_{i=1}^{k} \ln\left(\frac{\Pr[Y^0_i = y_i | Y^0_1 = y_1, \ldots, Y^0_{i-1} = y_{i-1}]}{\Pr[Y^1_i = y_i | Y^1_1 = y_1, \ldots, Y^1_{i-1} = y_{i-1}]}\right) \\
&\stackrel{\text{def}}{=} \sum_{i=1}^{k} c_i(y_1, \ldots, y_i).
\end{aligned}$$

Now for every prefix $(y_1, \ldots, y_{i-1})$ we condition on $Y^0_1 = y_1, \ldots, Y^0_{i-1} = y_{i-1}$, and analyze the expectation and maximum possible value of the random variable

$$c_i(Y^0_1, \ldots, Y^0_i) = c_i(y_1, \ldots, y_{i-1}, Y^0_i).$$

For any value y_i, we have

$$|c_i(y_1, \ldots, y_{i-1}, y_i)| = |\ln\left(\frac{\Pr[Y^0_i = y_i | y_1, \ldots, y_{i-1}]}{\Pr[Y^1_i = y_i | y_1, \ldots, y_{i-1}]}\right)| \leq \epsilon.$$

which follows from the fact that $\max(D_{\infty}(Y^0_i||Y^1_i), D_{\infty}(Y^1_i||Y^0_i)) \leq \epsilon$. By Lemma 11, we have:

$$\mathrm{E}\left[c_i(Y^0_1, \ldots, Y^0_i) | Y^0_1 = y_1, \ldots, Y^0_{i-1} = y_{i-1}\right] = D(Y_i(t; y_1, \ldots, y_{i-1})||Y_i(t'; y_1, \ldots, y_{i-1})) \leq \epsilon \cdot (e^{\epsilon} - 1).$$

Thus we can apply Azuma's Inequality (Theorem 9) to the random variables $C_i = c_i(Y^0_1, \ldots, Y^0_i)$, letting $f(C_1, \ldots, C_k) = \sum_{i=1}^{k} C_i$, and noting that that $\mathrm{E}[f] = k\epsilon \cdot (e^{\epsilon} - 1)$. Taking $t = \sqrt{2k \log(1/\delta)}\epsilon$, we find:

$$\Pr[Y^0 \in B] = \Pr[f(C_1, \ldots, C_k) \geq \mathrm{E}[f] + \sqrt{2k \log(1/\delta)}\epsilon] \leq \delta$$

as desired. ■

4.2 Computing (Coarse) Correlated Equilibria

In this section, we show a simple efficient algorithm for computing η-approximate coarse correlated equilibria in a game $\mathcal{G}$. We first that computing a *distribution* $\mathcal{D}$ that is an η-approximate coarse correlated equilibrium can be reduced to computing a *sequence of actions* $a^1, \ldots, a^T \in A^n$ that satisfy the following property:

Definition 12 *A sequence of actions $a^1, \ldots, a^T \in A^n$ has regret η if for every player i and for every action a_i':*

$$\frac{1}{T}\sum_{t=1}^{T} c_i(a^t) \leq \frac{1}{T}\sum_{t=1}^{T} c_i(a_i', a_{-i}^t) + \eta$$

It is immediate that given a sequence of actions $a^1, \ldots, a^T \in A^n$ that has regret η, the uniform distribution over $a^1, \ldots, a^T$ forms an η-approximate coarse-correlated equilibrium. Hence, we can focus on computing low regret sequences of actions.

Let us abstract away one level further, and talk about a single player having low regret with respect to an abstract series of loss functions. Suppose for every $t \in \{1, \ldots, T\}$, we can define a loss for each action by specifying $\ell^t : A \to [0, 1]$. Then we can say:

Definition 13 *A sequence of actions $a_1, \ldots, a_T \in A$ has regret η with respect to a series of loss functions $\ell^1, \ldots, \ell^T$ if for every action a_i':*

$$\frac{1}{T}\sum_{t=1}^{T} \ell^t(a_t) \leq \frac{1}{T}\sum_{t=1}^{T} \ell^t(a_i') + \eta$$

We can now give a simple algorithm that is guaranteed to generate a low regret sequence of actions with respect to an *arbitrary* sequence of loss functions.

Algorithm 1 The Multiplicative Weights Algorithm

INPUT: A, and a sequence of loss functions $\ell^1, \ldots, \ell^T$.
OUTPUT: A sequence of actions $a_1, \ldots, a_T$.

Let $\eta \leftarrow \sqrt{\frac{\log |A|}{T}}$
For each $a_i \in A$ **Let** $w_i^1 \leftarrow 1$.
for $t = 1$ to T **do**
  Let $W^t = \sum_{a_i \in A} w_i$
  Sample a_t from the distribution that puts weight $\frac{w_i^t}{W^t}$ on each action $a_i \in A$.
  Let $w_i^{t+1} \leftarrow w_i^t \cdot (1 - \eta \cdot \ell^t(a_t))$
end for
Output $a_1, \ldots, a_T$.

Theorem 14 *For any sequence of adaptively chosen loss functions $\ell^1, \ldots, \ell^T$, if actions $a_1, \ldots, a_T$ are chosen according to the multiplicative weights algorithm, then:*

$$\mathrm{E}[\frac{1}{T}\sum_{t=1}^{T} \ell^t(a_t)] \leq \frac{1}{T}\sum_{t=1}^{T} \ell^t(a_i') + 2\sqrt{\frac{\ln |A|}{T}}$$

Proof Define $L^T = \mathrm{E}[\sum_{t=1}^{T} \ell^t(a_t)]$ to be the expected cumulative loss of the algorithm, and define $F^t = \mathrm{E}_{a_t}[\ell^t(a_t)]$ to be the expected loss of the algorithm at time t. (So $L^T = \sum_{t=1}^{T} F^t$). By definition, we have:

$$F^t = \sum_{a_i \in A} \frac{w_i^t \cdot \ell^t(a_i)}{W^t}.$$

Note that we can also write:

$$W^{t+1} = W^T - \eta \cdot \sum_{a_i \in A} w_i^t \cdot \ell^t(a_i) = W^t \cdot (1 - \eta F^t)$$

Now note that $W^1 = |A|$, and so by iteratively applying the above equality, we know:

$$W^{T+1} = |A| \cdot \prod t = 1^T(1 - \eta F^t)$$

Taking logs, we find:

$$\begin{aligned}
\ln(W^{T+1}) &= \ln|A| + \sum_{t=1}^{T} \ln(1 - \eta F^t) \\
&\leq \ln|A| - \eta \sum_{t=1}^{T} F^t \\
&= \ln|A| - \eta \mathrm{E}[L^T]
\end{aligned}$$

where the inequality follows from the fact that $\ln(1 - x) \leq -x$ for $x \in [0, 1]$.

We also know that for every action $a_i \in A$:

$$\begin{aligned}
\ln(W^{T+1}) &\geq \ln(w_i^T) \\
&= \sum_{t=1}^{T} \ln(1 - \eta \ell^t(a_i)) \\
&\geq -\sum_{t=1}^{T} \eta \ell^t(a_i) - \sum_{t=1}^{T} (\eta \ell^t(a_i))^2 \\
&\geq -\eta \sum_{t=1}^{T} \ell^t(a_i') - \eta^2 T
\end{aligned}$$

where the first inequality follows from the fact that $\ln(1 - x) \geq -x - x^2$ for $x \in [0, 1]$. Combining these two inequalities, we can deduce:

$$\ln|A| - \eta \mathrm{E}[L^T] \geq -\eta \sum_{t=1}^{T} \ell^t(a_i') - \eta^2 T$$

or equivalently, for every action a_i:

$$\frac{1}{T}\mathrm{E}[L^T] \leq \frac{1}{T} \sum_{t=1}^{T} \ell^t(a_i') + \eta + \frac{\ln|A|}{\eta}$$

Plugging in $\eta = \sqrt{\frac{\ln|A|}{T}}$ proves the theorem. ■

Note that we have proven a bound only on the *expected* regret. However, we can apply Azuma's inequality here as well to get a high probability bound on the regret of the sequence of actions generated by the multiplicative weights algorithm:

Corollary 15 *With probability $1 - \beta$, the sequence of actions $a_1, \ldots, a_T$ generated by multiplicative weights satisfies:*

$$\frac{1}{T} \sum_{t=1}^{T} \ell^t(a_t) \leq \frac{1}{T} \sum_{t=1}^{T} \ell^t(a_i') + 2\sqrt{\frac{\ln|A| + \ln\frac{1}{\beta}}{T}}$$

Algorithm 2 ComputeCCE

Let $T = 16 \cdot \frac{\ln|A| + \ln\frac{2n}{\beta}}{\alpha^2}$

Let $\sigma = \Delta_{\mathcal{G}} \cdot \left(\frac{\sqrt{8T \cdot n \cdot |A| \ln(1/\delta)}}{\epsilon}\right)$

For each player i, initialize a copy of multiplicative weights $MW_i(A, \ell_i^t)$.

for $t = 1$ to T **do**

  For each player i get the t'th action a_i^t from MW_i.

  For each player i, feed to MW_i loss function defined as $\ell_i^t(a') = c_i(a', a_{-i}^t) + Z_{i,t,a'}$ for each $a' \in A$ for $Z_{i,t,a'} \sim \text{Lap}(\sigma)$.

end for

Output $a^1, \ldots, a^T$.

5 Privately Computing Equilibria

We can now describe a jointly differentially private algorithm for computing a coarse-correlated equilibrium in any large game $\mathcal{G}$ with sensitivity $\Delta_{\mathcal{G}}$.

Theorem 16 *ComputeCCE is (ϵ, δ)-jointly differentially private.*

Proof [Sketch] The idea will be to argue that the mechanism's output visible to all players $j \neq i$, $a_{-i}^1, \ldots, a_{-i}^T$, is (ϵ, δ)-differentially private in the *actions* $a_i^1, \ldots, a_i^T$ played over the course of the algorithm by MW_i. Since the mechanism depends on agent i's type only through these actions, this is sufficient to prove (ϵ, δ)-joint differential privacy with respect to player types.

Note that the distribution on actions $a_{-i}^1, \ldots, a_{-i}^T$ is fixed as a deterministic function of the loss vectors $\{\ell_{-i}^t(a')\}_{t \in [T], a' \in A}$, so by the fact that post-processing does not degrade differential privacy guarantees, it is sufficient to argue the privacy of these loss vectors.

Recall that each $j \neq i$, $a' \in A$, and $t \in [T]$, $\ell_j^t(a') = c_i(a', a_{-i}^t) + Z_i^t$ where $Z_{i,t,a'} \sim \text{Lap}\left(\Delta_{\mathcal{G}} \cdot \left(\frac{\sqrt{8T \cdot n \cdot |A| \ln(1/\delta)}}{\epsilon}\right)\right)$. Since c_j has sensitivity $\Delta_{\mathcal{G}}$ to the action of player i, by the guarantee of the Laplace mechanism, the computation of each $\ell_j^t(a')$ satisfies $\epsilon' = \frac{\epsilon}{\sqrt{8T \cdot n \cdot |A| \ln(1/\delta)}}$-differential privacy. We can therefore view the release of the vector $\{\ell_j^t(a')\}_{j \neq i, t \in [T], t' \in [T], a' \in A}$ as the composition of $T \cdot |A| \cdot (n-1) \leq T \cdot |A| \cdot n$ ϵ'-differentially private mechanisms.

But by Theorem 10, the composition of $T \cdot |A| \cdot n$ ϵ'-differentially private algorithms is (ϵ, δ) differentially private, which is what we wanted. ■

Finally, we can analyze the quality of the coarse correlated equilibrium that ComputeCCE produces.

Theorem 17 *With probability $1 - \beta$, ComputeCCE returns a sequence $a^1, \ldots, a^T \in A^n$ that has at most α regret when setting:*

$$\alpha = \left(\frac{\Delta_{\mathcal{G}} \sqrt{192n \cdot |A| \cdot \ln(1/\delta) \cdot \ln\left(\frac{2|A| \cdot n}{\beta}\right)}}{\epsilon}\right) = \tilde{O}\left(\frac{\Delta_{\mathcal{G}} \cdot \sqrt{n \cdot |A|}}{\epsilon}\right)$$

Proof We wish to bound:

$$\circledast \quad \stackrel{\text{def}}{=} \quad \max_{i \in [n], a_i' \in A} \frac{1}{T} \sum_{t=1}^{T} \left(c_i(a^t) - c_i(a_i', a_{-i}^t)\right)$$

$$\leq \underbrace{\max_{i\in[n],a_i'\in A}\left|\frac{1}{T}\sum_{t=1}^{T}\left(\ell_i^t(a_i^t)-\ell_i^t(a_i')\right)\right|}_{\circledast\circledast}+\underbrace{\max_{i\in[n],a_i'\in A}\left|\frac{1}{T}\sum_{t=1}^{T}Z_{i,t,a_i'}\right|}_{\circledast\circledast\circledast}$$

We bound these two terms separately. $\circledast\circledast$ is exactly the quantity bounded by Corollary 15. It tell us that for each player i except with probability $\beta/2n$:

$$\max_{a_i'\in A}\left|\frac{1}{T}\sum_{t=1}^{T}\left(\ell_i^t(a_i^t)-\ell_i^t(a_i')\right)\right|\leq 2\sqrt{\frac{\ln|A|+\ln\frac{2n}{\beta}}{T}}=\alpha/2$$

Hence, except with probability $\beta/2$, this condition holds for all n players i and we have:

$$\circledast\circledast\leq\frac{\alpha}{2}$$

It remains to bound $\circledast\circledast\circledast$. We recall a theorem we saw several lectures ago: If $Z_1,\ldots,Z_T\sim\text{Lap}(k/\epsilon)$, and $Y=\sum_{t=1}^{T}Z_i$ then:

$$\Pr[|Y|\geq c\cdot k]\leq\exp\left(\frac{c^2\epsilon^2}{6T}\right)$$

Applying this bound, we find that except with probability $\beta/2$,

$$\circledast\circledast\circledast\leq\frac{\alpha}{2}.$$

In combination, we have shown that with probability $1-\beta$, the sequence $a^1,\ldots,a^T\in A^n$ has regret at most $\alpha=\tilde{O}\left(\frac{\Delta_{\mathcal{G}}\cdot\sqrt{n\cdot|A|}}{\epsilon}\right)$, which is what we wanted. ■

So we have just seen how to compute coarse correlated equilibria under the constraint of joint differential privacy. Using essentially the same argument, but with a more sophisticated base-algorithm that generates sequences of actions that correspond to correlated equilibria rather than coarse correlated equilibria, we could have shown the following theorem:

Theorem 18 ([KPRU14]) *There is an (ϵ,δ)-jointly differentially private mechanism that given any large game with sensitivity $\Delta_{\mathcal{G}}$ and action space A, with probability $1-\beta$ computes an α-approximate correlated equilibrium for:*

$$\alpha=\tilde{O}\left(\frac{\Delta_{\mathcal{G}}|A|^{3/2}\sqrt{n\log(1/\delta)}\log(1/\beta)}{\epsilon}\right)$$

Using somewhat more sophisticated techniques to compute the losses (rather than just adding Laplace noise), we can also prove the following theorem, which has an exponentially better dependence on the number of actions $|A|$, at the expense of incurring a logarithmic dependence on the size of the typespace.

Theorem 19 ([KPRU14]) *There is an (ϵ,δ)-jointly differentially private mechanism that given any large game with sensitivity $\Delta_{\mathcal{G}}$, action space A, and typespace $\mathcal{T}$, with probability $1-\beta$ computes an α-approximate correlated equilibrium for:*

$$\alpha=\tilde{O}\left(\frac{\Delta_{\mathcal{G}}\sqrt{n}\log(|\mathcal{T}|)^{3/2}\log(|A|/\beta)\log(1/\delta)}{\epsilon}\right)$$

6 Game Theoretic Implications

All that remains is to combine our theorem statements now, and state a theorem about mediated games. What we have proven is the following:

Theorem 20 ([KPRU14]) *In every large game $\mathcal{G}$ with sensitivity $\Delta_{\mathcal{G}}$, action space A, and typespace $\mathcal{T}$ there exist mediators M_1 and M_2 such that good behavior forms an η_1-approximate ex-post Nash equilibrium in the mediated game $\mathcal{G}_{M_1}$ and an η_2-approximate ex-post Nash equilibrium in the mediated game $\mathcal{G}_{M_2}$ for:*

$$\eta_1 = \tilde{O}\left(\sqrt{\Delta_{\mathcal{G}}} n^{1/4} |A|^{3/4}\right)$$

and

$$\eta_2 = \tilde{O}\left(\sqrt{\Delta_{\mathcal{G}}} n^{1/4} \sqrt{\log(|A|)} \log(|\mathcal{T}|)^{3/4}\right)$$

When $\Delta_{\mathcal{G}} = O(1/n)$ this gives $\eta_1 = \tilde{O}\left(\frac{|A|^{3/4}}{n^{1/4}}\right)$ and $\eta_2 = \tilde{O}\left(\frac{\sqrt{\log|A|}\log(\mathcal{T})^{3/4}}{n^{1/4}}\right)$

Moreover, when players play according to this ex-post Nash equilibrium, the induced play is an approximate correlated equilibrium of the full information game.

A couple of remarks about these theorems, in comparison to the theorems we proved in previous lectures:

1. The mediator here requires the ability to verify types, or alternately, in the mediated games we discuss here, players do not have the ability to misreport their type if they use the mediator (but they can opt out of using it, and don't have to follow its advice). This is a stronger assumption, but you can imagine settings (e.g. in regulated financial markets) where type verifiability can be enforced.

2. In exchange for this stronger mediator, we have not had to assume *anything* about the game, other than the largeness condition on player utilities. This is an extremely mild assumption compared to "large market" assumptions commonly made in economics: we do not need any assumptions on how players types are realized, for example.

3. We also don't need the market to be that large for these results to kick in. Using mediator 2, for example, we get asymptotic truthfulness even when the size of the game n is exponentially smaller than either the typespace or the actionspace of the game. Using mediator 1, we have no dependence at all on the size of the typespace (but now incur a polynomial dependence on the size of A).

Bibliographic Information The composition theorem in this lecture is from Dwork, Rothblum, and Vadhan [DRV10]. The results about mediators in large games is from Kearns, Pai, Roth, and Ullman [KPRU14].

References

[DRV10] Cynthia Dwork, Guy N. Rothblum, and Salil P. Vadhan. Boosting and differential privacy. In *FOCS*, pages 51–60. IEEE Computer Society, 2010.

[KPRU14] Michael Kearns, Mallesh M Pai, Aaron Roth, and Jonathan Ullman. Mechanism design in large games: Incentives and privacy. *Proceedings of the annual 5th Innovations in Theoretical Computer Science (ITCS) conference*, 2014.