

AN ANTI-FOLK THEOREM FOR LARGE REPEATED GAMES WITH IMPERFECT MONITORING*

MALLESH M. PAI[†] AARON ROTH[‡] JONATHAN ULLMAN[§]

FEBRUARY 13, 2014

ABSTRACT

We study infinitely repeated games in settings of imperfect monitoring. We first prove a family of theorems that show that when the signals observed by the players satisfy a condition known as (ϵ, γ) -differential privacy, that the folk theorem has little bite: for values of ϵ and γ sufficiently small, for a fixed discount factor, any equilibrium of the repeated game involve players playing approximate equilibria of the stage game in every period. Next, we argue that in large games (n player games in which unilateral deviations by single players have only a small impact on the utility of other players), many monitoring settings naturally lead to signals that satisfy (ϵ, γ) -differential privacy, for ϵ and γ tending to zero as the number of players n grows large. We conclude that in such settings, the set of equilibria of the repeated game collapse to the set of equilibria of the stage game.

1 INTRODUCTION

In a repeated game, agents interact with one another repeatedly, and can base their actions on their memory of past interactions. One of the most robust features of models of infinitely repeated games is the multiplicity of equilibria. Repeatedly playing Nash equilibria of the underlying stage game is preserved as an equilibrium of the repeated game, but a much larger set is introduced as well. The so called “folk theorems” state, informally, that any individually rational payoffs for

*We gratefully acknowledge the support of NSF Grant CCF-1101389.

[†]Department of Economics, University of Pennsylvania. Email: malles@econ.upenn.edu.

[‡]Department of Computer and Information Science, University of Pennsylvania. aaroth@cis.upenn.edu Supported in part by NSF Grant CCF-1101389, and NSF CAREER grant, and a Google Research Grant.

[§]Center for Research on Computation and Society, Harvard University. jullman@seas.harvard.edu

each player (i.e. payoffs that are above their minmax value in the game) can be achieved as an equilibrium of the repeated game (if players are patient enough). The folk theorems can be interpreted as either positive or negative statements. From the perspective of an optimist (or a mechanism designer), they mean that high social-welfare behaviors (like cooperation in the Prisoner’s Dilemma) can be supported as equilibria in repeated games, even though they cannot be supported as equilibria of the stage game. From the perspective of a pessimist (or a computer scientist), they mean that the social welfare of the *worst* equilibrium can get worse (i.e. the price of anarchy increases). Regardless of the stance one takes, because of the severe multiplicity problem, Nash equilibria of repeated games lose much of their predictive power.

The simplest model of repeated games are repeated games with *perfect monitoring*. In these games, after each period, each agent observes the exact action played by each of his opponents. In several settings of interest, this assumption is unrealistic. A more natural model is repeated games with *imperfect monitoring*, in which agents receive some kind of noisy signal about the play at the last period. This signal could be an estimate of the payoff of each of a player’s actions, a noisy histogram of how many players played each action, or information that is encoded as a set of prices. The game could be one of *private monitoring*, in which each player receives his own signal privately, or *public monitoring*, in which every player observes the same signal. Folk theorems are known to hold in almost all models of repeated games, even those with imperfect public or private monitoring. See [Fudenberg, Levine, and Maskin \(1994\)](#) for a classic reference in the case of public monitoring, [Sugaya \(2011\)](#) for a recent result in the case of private monitoring, and [Mailath and Samuelson \(2006\)](#) for a textbook treatment.

In this paper, we show that these theorems will often lack bite in “large” games of imperfect monitoring. Particularly, we consider games where the observed signals satisfy the notion of “differential privacy” ([Dwork, McSherry, Nissim, and Smith, 2006](#)), which we show arises naturally in many games of imperfect monitoring with a large number of players. We show that in games with differentially private signals, the set of equilibria of the repeated games must involve play of approximate equilibria of the stage game in every period—approximate Nash equilibrium in the case of public monitoring and approximate correlated equilibrium in the case of private monitoring. Our results therefore suggests that in games with a large number of players, we should not expect to see “folk theorem equilibria.” This also adds to a long standing question in the literature on repeated games—for a given fixed δ what payoffs can be achieved by equilibria of the repeated game?¹ We implicitly provide an upper bound on this set—it is the set of payoffs that can be achieved by the play of approximate stage game equilibria in every period, with the

¹See e.g. [Mailath, Obara, and Sekiguchi \(2002\)](#) for a characterization in the case of the Prisoner’s Dilemma with perfect monitoring.

degree of approximation being bounded by our theorems.

A different way to view our results arises from a seeming discontinuity between dynamic models with a continuum of players and those with a large but finite number of players.² In a dynamic model, incentives are generated by agents punishing or rewarding others for past play. When only aggregate play is observed, then the play of one single agent does not affect the observed outcome in the continuum case, and hence individual deviations cannot be detected. With a finite number of players, a deviation may be small, but it is perfectly observable. Therefore deviations can be punished or rewarded. A small literature attempts to bridge this discontinuity by postulating a model of noisy observation. Our paper can be thought of as a “quantitative” version of their limit results for repeated games. In particular, the noisiness of our signals are parametrized by two numbers, ϵ and γ . We show that as $\epsilon + \gamma$ vanishes, we approximate the unobservability of the continuum setting, while the perfect observability setting corresponds to $\epsilon + \gamma$ large.

1.1 RELATED WORK

It has long been understood (hence the term “folk theorem”) that the intertemporal incentives generated by repeated interaction can support more outcomes than just the equilibria of the one-off interaction. The literature is too large to survey here, we refer the interested reader to the textbook by [Mailath and Samuelson \(2006\)](#) for a comprehensive and up-to-date treatment. We review concepts that are standard in the literature as needed in Section 2. Repeated games have also been studied in the computer science literature. An older literature studies repeated games when agents have limited computational resources—see e.g. [Megiddo and Wigderson \(1986\)](#), [Papadimitriou and Yannakakis \(1994\)](#) or [Abreu and Rubinstein \(1988\)](#). There has also been some recent interest in repeated games in the computer science community. These consider games of perfect monitoring, and study the complexity of computing an equilibrium of the repeated game—see e.g. [Borgs, Chayes, Immorlica, Kalai, Mirrokni, and Papadimitriou \(2010\)](#) and [Halpern, Pass, and Seeman \(2013\)](#).

The idea that noisy observation may limit the equilibria in repeated games has some precedent in the literature—notably the papers of [Green \(1980\)](#), [Sabourian \(1990\)](#) and [Al-Najjar and Smorodinsky \(2001\)](#).³ In relation to our setting, note that the first paper restricts attention to trigger strategies, the first two papers assume the stage game is anonymous, and all three consider public equilibria in a public monitoring setting. By contrast our results apply to

²We thank Nabil Al-Najjar and Preston McAfee for pointing out this connection to us.

³[Levine and Pesendorfer \(1995\)](#) study similar questions in dynamic games.

general games that satisfy the differential privacy condition, and apply also to private monitoring settings. The results in those papers are limit theorems, while we prove quantitative bounds on the relationship between the “amount” of noise and the degree of deviation from stage-Nash equilibria that can be supported. Finally, while their conditions on the monitoring structure may be hard to verify, we are able we use techniques from the literature on differential privacy (see below) to give examples of “natural” settings satisfying our conditions.

“Differential Privacy” is a condition that a family of probability distributions (which are parameterized by n agent “reports”) may satisfy, which provides a multiplicative bound on the influence that any single report can have on the resulting distribution. It was first defined by [Dwork, McSherry, Nissim, and Smith \(2006\)](#) in the computer science literature, and has become a standard “privacy” solution concept. We do not use the condition here to deal with privacy concerns, but instead just as a useful Lipschitz condition that is commonly satisfied by many noisy families of signal distributions. The connection between differential privacy and game theory was first noted by [McSherry and Talwar \(2007\)](#) who observed that when this condition is applied to an auction mechanism, it can be used to argue approximate truthfulness. They use it to derive prior-free near revenue-optimal mechanisms for digital goods auctions. Since then, differential privacy has been used as a tool in a number of game theoretic settings, including mechanism design without money ([Nissim, Smorodinsky, and Tennenholtz, 2012](#)) and mechanism design in prior-free settings of incomplete information with mechanisms that have extremely weak powers of enforcement ([Kearns, Pai, Roth, and Ullman, 2014](#); [Rogers and Roth, 2013](#)). A related line of work seeks to use traditional tools in mechanism design in settings in which agents have preferences for “privacy”, which are quantified by measures related to differential privacy ([Ghosh and Roth, 2013](#); [Nissim, Orlandi, and Smorodinsky, 2012](#); [Xiao, 2013](#); [Chen, Chong, Kash, Moran, and Vadhan, 2013](#); [Ghosh and Ligett, 2013](#); [Nissim, Vadhan, and Xiao, 2014](#)). See [Pai and Roth \(2013\)](#) for a survey of this area and its connections to mechanism design and game theory.

2 PRELIMINARIES

We consider infinitely repeated games that are played over a series of periods. Time is discrete and is indexed by $t = 0, 1, 2, \dots$

STAGE GAME There are n players who repeatedly play a stage game G . Each player i has a finite set of actions A_i , and gets payoff $u_i : \prod_{j=1}^n A_j \rightarrow [0, 1]$. We will denote an action profile (i.e. a vector of actions, one for each player) by $\mathbf{a} \in \prod_{j=1}^n A_j$. Mixed strategies are probability

distributions over pure strategies, which we write as ΔA_i for player i . We write $\alpha_i \in \Delta A_i$ to denote a mixed strategy for player i , and α to denote a vector of mixed strategies. Payoffs for mixed strategies are, as usual, the expected payoff when the profile of actions is drawn according to the distribution specified by the mixed strategies.

All players discount the future at a rate $\delta \in [0, 1)$. Therefore an annuity that pays 1 in every period, forever, has a present discounted value of $\frac{1}{1-\delta}$. As is standard in the literature on repeated games, we normalize all payoffs by $(1 - \delta)$ —the value of the annuity of 1 forever is therefore normalized to 1, while e.g. a payoff of 1 today and nothing thereafter is normalized to $(1 - \delta)$ (because the agent would be indifferent between receiving 1 this period and nothing thereafter and an annuity of $(1 - \delta)$).

Before we describe repeated games, we begin by reviewing two approximate equilibrium concepts for the stage game.

DEFINITION 1. *An η -approximate Nash equilibrium of the stage game G is a vector $\alpha \in \prod_{j=1}^n \Delta A_j$ of mixed strategies, such that for all i and all $a_i \in A_i$:*

$$u_i(\alpha) \geq u_i(a_i, \alpha_{-i}) - \eta$$

DEFINITION 2. *An η -approximate correlated equilibrium of the stage game G is defined by a set of signals S , a distribution $\mathcal{D} \in \Delta S^n$ over n -tuples of signals, and a function $\sigma_i : S \rightarrow \Delta A_i$ mapping signals to mixed strategies for each player i . (\mathcal{D}, σ) is an η -approximate correlated equilibrium if for every signal $s \in S$ and every action $a_i \in A_i$:*

$$\mathbb{E}_{\mathbf{s} \sim \mathcal{D}}[u_i(\sigma_i(s), \sigma_{-i}(s_{-i})) | s_i = s] \geq \mathbb{E}_{\mathbf{s} \sim \mathcal{D}}[u_i(a_i, \sigma_{-i}(s_{-i})) | s_i = s] - \eta$$

MONITORING We consider models of both public and private monitoring. In a game of *public monitoring*, at the end of every period, after players have played their actions for that period, they receive feedback about the actions taken by others via a *public signal*. Formally, at the end of the period, when the profile of actions taken is \mathbf{a} , all players commonly learn a signal $s \in S$, where s is drawn independently by a distribution $P_{\mathbf{a}} \in \Delta S$.⁴

In a game of *private monitoring*, at the end of every period, after players have played their actions for that period, each player i receives a private signal $s_i \in S$. Player i does not see the signals s_j for $j \neq i$. Formally, at the end of a period in which the profile of actions taken is \mathbf{a} , a vector of signals $s \in S^n$ is drawn from a distribution $P_{\mathbf{a}} \in \Delta S^n$, and each player i observes component s_i of s . Note that the signals that different players receive can be correlated. Further,

⁴Note that this model subsumes the standard perfect monitoring model—if $S \equiv \prod A_i$ and $P_{\mathbf{a}}(s) = 1$ iff $s = \mathbf{a}$.

note that this generalizes the public monitoring setting—public monitoring is simply the case where every player’s signal is identical.

The ex-post payoff of each player i is given by a function $U_i : A_i \times S \rightarrow \mathbb{R}$ so that the signal a player receives is a sufficient statistic for the payoff he receives. We assume that

$$u_i(a_i, a_{-i}) = \sum_{s \in S} U_i(a_i, s) P_{\mathbf{a}}(s), \quad (1)$$

so that the noisy payoff does not affect the incentive structure of the game.

We consider a commonly satisfied condition on the distribution of signals that players observe—informally, that they are *noisy* in a way that does not reveal “much” about the action of any single player.

DEFINITION 3. A public monitoring signal structure is said to satisfy (ϵ, γ) -differential privacy for some $\epsilon, \gamma > 0$ if for every player i and for every $\mathbf{a} = (a_1, a_2, \dots, a_n)$, $\mathbf{a}' = (a'_i, a_{-i})$, and any event $E \subseteq S$,

$$\exp(-\epsilon) P_{\mathbf{a}'}(E) - \gamma \leq P_{\mathbf{a}}(E) \leq \exp(\epsilon) P_{\mathbf{a}'}(E) + \gamma.$$

A private monitoring signal structure satisfies (ϵ, γ) -differential privacy if for every player i and for every $\mathbf{a} = (a_1, a_2, \dots, a_n)$, $\mathbf{a}' = (a'_i, a_{-i})$, and for every event $E \subseteq S^n$,

$$\exp(-\epsilon) P_{\mathbf{a}'}(E) - \gamma \leq P_{\mathbf{a}}(E) \leq \exp(\epsilon) P_{\mathbf{a}'}(E) + \gamma.$$

REMARK 1. Differential privacy is a condition applied to study the privacy of randomized algorithms on a “database”(i.e. not generally applied to the signal structure of a repeated games), and was first defined by Dwork, McSherry, Nissim, and Smith (Dwork, McSherry, Nissim, and Smith, 2006).

REMARK 2. Due to (1), the assumption that a signalling structure satisfies (ϵ, γ) -differential privacy is nontrivial. Fixing values of ϵ, γ , there exist stage games for which no such (ϵ, γ) -differentially private signal structure can satisfy (1)). That said, several stage games of interest can be paired with differentially private signal structures. Broadly, this includes any *large* game, i.e. one in which any player, by unilaterally changing the action he plays, has a small impact on the stage-game utility of any other player. We present several examples in Section 4.

For technical reasons, the following definition is often assumed (see e.g. Mailath and Samuelson (2006)). We say that the monitoring structure is one of *no ob-*

servable deviations, i.e. that the marginal distribution of private signals have full support regardless of the action profile played. Formally:

DEFINITION 4. A repeated game of public monitoring is said to have no observable deviations if

$$\forall \mathbf{a} \in \prod A_i, s \in S : P_{\mathbf{a}}(s) > 0.$$

A repeated game of private monitoring is said to have no observable deviations if

$$\forall \mathbf{a} \in \prod A_i, s_i \in S : P_{\mathbf{a}}(s_i) > 0.$$

Intuitively, in the repeated game agents will be trying to infer what other players played from this signal—this assumption ensures that Bayesian updating is well defined after any observed signal.

HISTORIES AND STRATEGIES In the public monitoring setting the public information in period t is the history of public signals, $h^t = (s^1, s^2, \dots, s^{t-1})$. The set of public histories is

$$\mathcal{H} \equiv \cup_{t=2}^{\infty} S^{t-1}.$$

A history for a player includes both the public history and the history of actions he has taken, $h_i^t = (s^1, a_i^1, s^2, a_i^2, \dots, s^{t-1}, a_i^{t-1})$. Given a private history $h_i^t = (s^1, a_i^1, s^2, a_i^2, \dots, s^{t-1}, a_i^{t-1})$, one can define the public history $h^t = (s^1, s^2, \dots, s^{t-1})$ as just the vector of public signals observed. The set of private histories for player i is

$$\mathcal{H}_i \equiv \cup_{t=2}^{\infty} (A_i \times S)^{t-1}.$$

We refer to the vector of private histories (one for each player), at period t by $\mathbf{h}^t \equiv (h_1^t, h_2^t, \dots, h_n^t)$. In the private monitoring setting, the history for a player i is the history of the actions he has taken, together with his sequence of private signals: $h_i^t = (s_i^1, a_i^1, s_i^2, a_i^2, \dots, s_i^{t-1}, a_i^{t-1})$.

A pure strategy for player i is a mapping from the set of histories of player i to the set of actions of player i :

$$\sigma_i : \mathcal{H}_i \rightarrow A_i,$$

and a mixed strategy is a mapping from histories to distributions over actions.

Strategies in a repeated game can in general depend on a player's entire history, but in a game of public monitoring, we can also consider *public* strategies which depend only on the public signals observed so far (and not on the unobserved actions played by the player).

DEFINITION 5. A strategy σ_i is *public* if it only depends on the public history \mathcal{H} , i.e. for any period t and two private histories $h_i^t, h_i^{t'}$ with the same public history h^t ,

$$\sigma_i(h_i^t) = \sigma_i(h_i^{t'}).$$

SOLUTION CONCEPTS Equilibria are defined with respect to the cumulative discounted payoff that players obtain when the play according to a profile of strategies σ . Formally, recursively define the expected payoff of strategy profile σ to player i after history h^t as:

$$V_{i,\sigma}(h^t) = (1 - \delta)u_i(\sigma(h^t)) + \delta \sum_{s \in S} V_{i,\sigma}((h^t, s))P_{\sigma(h^t)}(s).$$

DEFINITION 6 (Nash Equilibrium). A set of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ is a *Nash equilibrium* of a repeated game if for every player i and every strategy σ'_i :

$$V_{i,\sigma}(\emptyset) \geq V_{i,(\sigma_{-i}, \sigma'_i)}(\emptyset)$$

A subgame perfect Nash equilibrium requires that from any history, the continuation play forms a Nash equilibrium.

DEFINITION 7 (Subgame perfect Nash equilibrium). In a game of perfect monitoring, a set of strategies $\sigma = (\sigma_1, \dots, \sigma_n)$ is a *subgame-perfect Nash equilibrium* of a repeated game if for every player i , every finite history h^t and every strategy σ'_i :

$$V_{i,\sigma}(h^t) \geq V_{i,(\sigma_{-i}, \sigma'_i)}(h^t)$$

In the case of imperfect monitoring, histories are not common knowledge (agents privately observe the actions they take), there is no well-defined “subgame.” There is, however, one in the case where all players use public strategies. Formally, in games of public monitoring, we can define a “perfect public equilibrium.” This is a particularly structurally simple equilibrium concept in repeated games, but one in which a folk theorem is known to hold. Although our results hold more generally, we will first prove our “anti-folk-theorem” for perfect public equilibrium.

DEFINITION 8. A perfect public equilibrium (PPE) is a profile of public strategies for each player (Definition 5) such that the continuation play from every public history constitutes a Nash Equilibrium.

Fudenberg, Levine, and Maskin (1994) show that as long as the signalling structure is rich enough (satisfying a full rank condition), a Folk theorem holds in perfect public equilibria.

Finally, to consider general (not necessarily public) strategies in games of public monitoring, and games of private monitoring, there are no counterparts to subgame perfect/ perfect public equilibria. Consequently, the refinement that is normally considered is that of sequential equilibrium, which requires that for any player i and any private history h_i^t , the continuation play is a best reply to the play of others given the agent's beliefs about the private histories they observe. Formally:

DEFINITION 9. Suppose game is one of no observable deviations (Definition 4). A strategy profile σ is a sequential equilibrium if for all private histories h_i^t , $\sigma_i(h_i^t)$ is a best reply to $\mathbf{E}[\sigma_{-i}(h_{-i}^t)|h_i^t]$.

As we alluded to earlier, the assumption of no observable deviations is made to ensure that the conditional expectation $\mathbf{E}[\sigma_{-i}(h_{-i}^t)|h_i^t]$ is well defined at every private history.

3 SIGNAL PRIVACY YIELDS ANTI-FOLK THEOREMS

3.1 PUBLIC MONITORING

3.1.1 PERFECT PUBLIC EQUILIBRIA

As we said earlier, the main ideas in our construction are easiest to explain for the case of public perfect equilibria (Definition 8). The intuition is essentially encapsulated as follows: in general, the only thing that prevents an agent at some stage t from deviating from his equilibrium strategy, and playing a (stage-game) best response to the distribution over his opponent's actions is fear of punishment: if his opponents can detect this deviation, then they can change their behavior to lower his expected *future* payoff. Differential privacy provides us a simple, worst-case way of quantifying the decrease in expected future payoff that can result from a single player's one-stage unilateral deviation. If this decrease can be made small enough, then it cannot serve as an incentive to prevent any player from playing anything other than an (approximate) stage-game best response to his opponents. Thus, every day, all players must be playing an approximate equilibrium of the stage game. We formalize this intuition in the next theorem:

THEOREM 1. Fix any repeated game with discount factor δ , with public signals that satisfy (ϵ, γ) -signal privacy. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a perfect public equilibrium (Definition 8). Then for every history h^t , the distribution on actions that result, $(\sigma_1(h^t), \dots, \sigma_n(h^t))$, forms an η -approximate Nash equilibrium of the stage game, for

$$\eta = \frac{\delta}{1 - \delta}(\epsilon + \gamma).$$

PROOF. Since σ forms a perfect public equilibrium of the repeated game, for every public history h^t we have:

$$V_{i,\sigma}(h^t) \geq V_{i,(\sigma_{-i}, \sigma'_i)}(h^t)$$

Expanding this definition, and noting that in particular single stage deviations cannot be profitable, we know for every single-stage deviation a'_i :

$$\begin{aligned} & (1 - \delta)u_i(\sigma(h^t)) + \delta \sum_{s \in S} V_{i,\sigma}((h^t, s))P_{\sigma(h^t)}(s) \\ & \geq (1 - \delta)u_i(a'_i, \sigma_{-i}(h^t)) + \delta \sum_{s \in S} V_{i,\sigma}((h^t, s))P_{(a'_i, \sigma_{-i}(h^t))}(s). \end{aligned}$$

By the definition of (ϵ, γ) -signal privacy, we also know that for all s ,

$$P_{\sigma(h^t)}(s) \geq \exp(-\epsilon)P_{(a'_i, \sigma_{-i}(h^t))}(s) - \gamma$$

Since $V_{i,\cdot}$ is the normalized infinite discounted sum of numbers between 0 and 1, $V_{i,\cdot}(\cdot) \in [0, 1]$.

We can therefore derive:

$$\begin{aligned} & (1 - \delta)u_i(\sigma(h^t)) \\ & \geq (1 - \delta)u_i(a'_i, \sigma_{-i}(h^t)) + \delta \sum_{s \in S} V_{i,\sigma}((h^t, s)) (P_{(a'_i, \sigma_{-i}(h^t))}(s) - P_{\sigma(h^t)}(s)) \\ & \geq (1 - \delta)u_i(a'_i, \sigma_{-i}(h^t)) + \delta \sum_{s \in S} V_{i,\sigma}((h^t, s)) (\exp(-\epsilon)P_{(a'_i, \sigma_{-i}(h^t))}(s) - \Delta_s - P_{\sigma(h^t)}(s)) \end{aligned}$$

where Δ_s is the smallest non-negative value that satisfies the inequality

$$P_{\sigma(h^t)}(s) \geq \exp(-\epsilon)P_{(a'_i, \sigma_{-i}(h^t))}(s) - \Delta_s.$$

Note that for signals s such that $P_{\sigma(h^t)}(s) \leq \exp(-\epsilon)P_{(a'_i, \sigma_{-i}(h^t))}(s)$, $\Delta_s = 0$. Let $S_+ = \{s \in S :$

$\Delta_s > 0\}$. We continue, noting that $\exp(-\epsilon) > 1 - \epsilon$ for $\epsilon < 1$:

$$\begin{aligned}
& (1 - \delta)u_i(\sigma(h^t)) \\
& \geq (1 - \delta)u_i(a'_i, \sigma_{-i}(h^t)) + \delta \sum_{s \in S} V_{i,\sigma}((h^t, s)) (-\epsilon \cdot P_{\sigma(h^t)}(s) - \Delta_s) \\
& \geq (1 - \delta)u_i(a'_i, \sigma_{-i}(h^t)) - \delta \left(\epsilon + \sum_{s \in S_+} V_{i,\sigma}((h^t, s)) \Delta_s \right) \\
& \geq (1 - \delta)u_i(a'_i, \sigma_{-i}(h^t)) - \delta(\epsilon + \gamma).
\end{aligned}$$

The last inequality follows because $\sum_{s \in S_+} \Delta_s$ is the smallest value Δ that satisfies the inequality $P_{\sigma_i(h^t)}(S_+) \geq \exp(-\epsilon)P_{(a'_i, \sigma_{-i}(h^t))}(S_+) - \Delta$, but by our guarantee of (ϵ, γ) signal privacy, we know that we must have $\Delta \leq \gamma$, and so $\sum_{s \in S_+} \Delta_s \leq \gamma$. Dividing both sides by $1 - \delta$ we get:

$$u_i(\sigma(h^t)) \geq u_i(a'_i, \sigma_{-i}(h^t)) - \frac{\delta}{1 - \delta}(\epsilon + \gamma)$$

In other words, at any history h^t , the prescribed strategy profile $\sigma(h^t)$ must form a $\frac{\delta}{1 - \delta}(\epsilon + \gamma)$ -approximate Nash equilibrium of the stage game. \square

With a bit more care, we can prove a related theorem that extends to the case of non-public strategies. The proof is similar, but with two caveats:

First, we must now consider a deviation of player i , σ'_i that not only makes a single stage game deviation to a'_i at some day t , but then continues to play using an artificial history that for every day $t' > t$, records that on day t , player i played an action drawn from $\sigma_i(h_i^t)$ instead of having played a'_i .

Second, the strategies that form the stage-game equilibrium are not exactly $\sigma_i(h_i^t)$, for with these strategies, each player might be best-responding not to the strategies of his opponent, but to the *expected* strategies of his opponents (defined by their unknown histories h_j^t , which take into account not just the publicly observed signals, but also the privately known action history), conditioned on the public part of the history.

Instead, the strategies that form the stage-game equilibrium are the randomized strategies $\sigma_i(h_i^t | s)$, where the randomness is taken over the realization of the players (privately known) action history, conditioned on the (publicly observed) signal history. We begin with some notation: Given a set of $t - 1$ public signals $s^t = (s_1, \dots, s_{t-1})$ and a vector of actions for player i , $\mathbf{a}_i^t = (a_i^1, \dots, a_i^{t-1})$, write $h_i^t(s^t, \mathbf{a}_i^t) = (a_i^1, s_1, \dots, a_i^{t-1}, s_{t-1})$ to denote the history that combines them. Given a strategy σ_i for player i mapping histories to actions, together with a vector

of public signals s^t write: $\hat{\sigma}_{i|s^t}$ to denote the probability distribution that plays each action a_i with probability equal to:

$$\Pr[a_i \sim \hat{\sigma}_{i|s^t}] = \sum_{\mathbf{a}_i^t \in A_i^{t-1}} \Pr[\mathbf{a}_i^t | s^t] \cdot \sigma_i(h_i^t(s^t, \mathbf{a}_i^t))$$

In other words, $\hat{\sigma}_{i|s^t}$ represents the distribution defined over actions played by player i , $\sigma_i(h_i^t)$ when all that is known is the public history, and the randomness is both over the choice of actions \mathbf{a}_i^t that define h_i^t , as well as the randomness in the mixed strategies $\sigma_i(h_i^t)$.

We can now state our theorem:

THEOREM 2. *Fix any repeated game with discount factor δ , with public signals that satisfy (ϵ, γ) -signal privacy and no observable deviation (Definition 4). Let $\sigma = (\sigma_1, \dots, \sigma_n)$ denote be a sequential equilibrium of the repeated game (Definition 9). Then for every public history of signals s^t , the distribution on actions played on day t , $(\hat{\sigma}_{1|s^t}, \dots, \hat{\sigma}_{n|s^t})$, forms an η -approximate Nash equilibrium of the stage game, for*

$$\eta = \frac{\delta}{1 - \delta}(\epsilon + \gamma).$$

The proof is deferred to the appendix.

3.1.2 NON-SUBGAME PERFECT EQUILIBRIA

We can prove a similar theorem even for non-subgame perfect equilibrium – except that now our claim only applies to histories that occur with nonzero probability when players play according to the specified equilibrium.

THEOREM 3. *Fix any repeated game with discount factor δ , with public signals that satisfy (ϵ, γ) -differential privacy. Let $\sigma = (\sigma_1, \dots, \sigma_n)$ denote a public equilibrium (not necessarily subgame perfect). Then for every history h^t that occurs with positive probability when players play according to σ , the distribution on actions at stage t $(\sigma_1(h^t), \dots, \sigma_n(h^t))$ forms an η -approximate Nash equilibrium of the stage game, for*

$$\eta = \frac{\delta}{1 - \delta}(\epsilon + \gamma)$$

PROOF. We will write $\Pr_\sigma[h^t]$ to denote the probability that a given public history h^t arises when players use strategies σ . For each $j \leq t$, write $h^{t, \leq j}$ to denote the sub-history of h^t consisting of

the first j periods, and $h^{t,j}$ to be the j^{th} period history. Then:

$$\Pr_{\sigma}[h^t] = \prod_{j=1}^t P_{\sigma(h^{t,\leq j})}(h^{t,j})$$

Now fix any T . We can write:

$$V_{i,\sigma}(\emptyset) = (1 - \delta) \left(\sum_{h^T \in S^T} \Pr_{\sigma}[h^T] \left(\left(\sum_{t=0}^T \delta^t u_i(\sigma(h^{T,\leq t})) \right) + \frac{\delta^{T+1}}{1 - \delta} \sum_{s \in S} V_{i,\sigma}(h^T, s) P_{\sigma(h^T)}(s) \right) \right)$$

Consider any history h^T such that $\Pr_{\sigma}[h^T] > 0$, and consider the deviation of player i , σ'_i that is identical to σ_i , except that on history h^T player i plays a stage-game best response to his opponents. $\sigma'_i(h^T) = \arg \max_{a \in A_i} u_i(a, \sigma_{-i}(h^T)) \equiv a_i^*$. Since σ is a Nash equilibrium of the repeated game, we know that $V_{i,\sigma}(\emptyset) - V_{i,(\sigma'_i, \sigma_{-i})}(\emptyset) \geq 0$. Since $\Pr_{\sigma}[h^T] > 0$, we can divide and write this difference as:

$$\begin{aligned} 0 &\leq \frac{1}{(1 - \delta) \Pr_{\sigma}[h^T]} (V_{i,\sigma}(\emptyset) - V_{i,(\sigma'_i, \sigma_{-i})}(\emptyset)) \\ &= \left(\delta^T (u_i(\sigma(h^T)) - u_i(a_i^*, \sigma_{-i}(h^T))) + \frac{\delta^{T+1}}{1 - \delta} \sum_{s \in S} V_{i,\sigma}(h^T, s) (P_{\sigma(h^T)}(s) - P_{(a_i^*, \sigma(h^T))}(s)) \right) \\ &\leq \left(\delta^T (u_i(\sigma(h^T)) - u_i(a_i^*, \sigma_{-i}(h^T))) + \frac{\delta^{T+1}}{1 - \delta} \sum_{s \in S} (P_{\sigma(h^T)}(s) - P_{(a_i^*, \sigma(h^T))}(s)) \right) \\ &\leq \left(\delta^T (u_i(\sigma(h^T)) - u_i(a_i^*, \sigma_{-i}(h^T))) + \frac{\delta^{T+1}}{1 - \delta} (\epsilon + \gamma) \right) \end{aligned}$$

where the last inequality follows from (ϵ, γ) -differential privacy, and the fact that $\exp(-\epsilon) \geq 1 - \epsilon$. Dividing through by δ^T and rearranging, we find:

$$(u_i(a_i^*, \sigma_{-i}(h^T)) - u_i(\sigma(h^T))) \leq \frac{\delta}{1 - \delta} (\epsilon + \gamma)$$

which completes the proof. \square

3.2 PRIVATE MONITORING

Next, we consider the case of private monitoring. Our theorem here is slightly weaker: that when playing a game with differentially private signals, every equilibrium of the repeated game must at each stage play an approximate *correlated* equilibrium of the stage game. The reason is natural:

agents are no longer aware of the history that their opponents are viewing, and so can no longer consider deviations that are best responses to the distributions being played on that day by their opponents. However, they may at any stage consider deviations that base their action on the past private signals that they observe, which gives a posterior distribution on the signals that their opponents observe. Since the distribution on signals that players receive can be correlated; the appropriate stage game solution concept is correlated equilibrium.

THEOREM 4. *Fix any repeated game with discount factor δ and private signals that satisfy (ϵ, γ) -differential privacy and no observable deviations (Definition 4). Let $\sigma = (\sigma_1, \dots, \sigma_n)$ denote a sequential equilibrium of the repeated game (Definition 9). Then for every history h^t , the distribution on actions at stage $(\sigma_1(h_1^t), \dots, \sigma_n(h_n^t))$ forms an η -approximate correlated equilibrium for*

$$\eta = \frac{\delta}{1 - \delta}(\epsilon + \gamma).$$

PROOF. Fix a sequential equilibrium σ . Then, for every player i and every history h_i^t , we can define a posterior distribution $\mathcal{D}(i, h_i^t, \sigma_{-i})$ on histories h_{-i}^t observed by the other players. Fix any history $\mathbf{h}^t = (h_1^t, \dots, h_n^t)$ for the n players. For some player i , define $a^* \in \arg \max_{a_i \in A_i} \mathbb{E}_{h_{-i}^t \sim \mathcal{D}(i, h_i^t, \sigma_{-i})}[u_i(a_i, \sigma_{-i}(h_{-i}^t))]$. Define σ'_i to be the strategy that is identical to σ_i , except on history h_i^t , it plays $\sigma'_i(h_i^t) = a^*$, and then play in future periods is as if an action drawn from $\sigma_i(h_i^t)$ was played in period t .

Formally, for any period $\tau > t$, with realized history h_i^τ , the deviation involves playing $\sigma_i(h_i^{\tau'})$, where $h_i^{\tau'}$ is the same as h_i^τ except in the component corresponding to the action played at time t ; i.e. $a_i^{\tau'} \sim \sigma_i(h_i^t)$ whereas $a_i^t = a^*$ by definition. Given the equilibrium strategies σ , we can write:

$$\begin{aligned} V_{i,\sigma}(h_i^t) &= (1 - \delta) \mathbb{E}_{h_{-i}^t \sim \mathcal{D}(i, h_i^t, \sigma_{-i})}[u_i(\sigma_i(h_i^t), \sigma_{-i}(h_{-i}^t))] \\ &\quad + \delta \sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i)) P[s, a_i | h_i^t] \end{aligned}$$

Now consider σ'_i to be the deviation we described above. Note that by playing σ'_i from history h_i^t onwards, player i 's expected discounted payoff can be written as:

$$\begin{aligned} V_{i,(\sigma_{-i}, \sigma'_i)}(h_i^t) &= (1 - \delta) \mathbb{E}_{h_{-i}^t \sim \mathcal{D}(i, h_i^t, \sigma_{-i})}[u_i(a^*, \sigma_{-i}(h_{-i}^t))] \\ &\quad + \delta \sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i)) P[s | h_i^t, a^*] P[a_i | \sigma(h_i^t)] \end{aligned}$$

Since σ forms a sequential equilibrium of the repeated game for every history h_i^t we have:

$$V_{i,\sigma}(h_i^t) \geq V_{i,(\sigma_{-i},\sigma'_i)}(h_i^t)$$

Substituting the definitions of $V_{i,\sigma}(h_i^t)$ and $V_{i,(\sigma_{-i},\sigma'_i)}(h_i^t)$ into this inequality, we have:

$$\begin{aligned} & \left(\mathbb{E}_{h_{-i}^t \sim \mathcal{D}(i,h_i^t,\sigma_{-i})}[u_i(a^*, \sigma_{-i}(h_{-i}^t))] - \mathbb{E}_{h_{-i}^t \sim \mathcal{D}(i,h_i^t,\sigma_{-i})}[u_i(\sigma_i(h_i^t), \sigma_{-i}(h_{-i}^t))] \right) \\ & \leq \frac{\delta}{1-\delta} \left(\sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i))P[s, a_i|h_i^t] - \sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i))P[s|h_i^t, a^*]P[a_i|\sigma(h_i^t)] \right) \end{aligned}$$

By (ϵ, γ) -differential privacy of the private signal, and the fact that $\exp(-\epsilon) \geq 1 - \epsilon$, therefore, we have

$$\leq \frac{\delta}{1-\delta}(\epsilon + \gamma).$$

By Definition 2, therefore, the strategies σ_i involve the play of $\frac{\delta}{1-\delta}(\epsilon + \gamma)$ -approximate correlated equilibrium in every period: the analogue of the private signal observed by players in that definition is the private history h_i^t , which are correlated via the strategies σ_i and monitoring structure. \square

4 EXAMPLES OF GAMES WITH SIGNAL PRIVACY

The rest of this paper will be devoted to showing by example that such differentially private signal structures may naturally arise in large games. In particular, we will show that if the number of players n is large, ϵ and γ may naturally be small. This suggests that in large societies, inter-temporal transfers may not support “much more” than repeated play of the static equilibria of the stage game.

Gaussian (i.e. normally distributed) mean-zero noise is a natural form of perturbation that arises in many settings of impartial information (it is the limiting distribution of many independent random processes). It also happens to be a noise distribution that guarantees differential privacy. Formally, a theorem from [Dwork, Kenthapadi, McSherry, Mironov, and Naor \(2006\)](#) will be useful.⁵ First we define the (L^2) sensitivity of a function.

DEFINITION 10. *A function $f : T^n \rightarrow \mathbb{R}^d$ has L^2 sensitivity s if for any $\mathbf{t} \in T^n$, $i \leq n$ and*

⁵A proof of this theorem in the form stated here can be found in Appendix A of [Hardt and Roth \(2012\)](#).

$t' \in T$, we have that

$$\|f(\mathbf{t}) - f(t', t_{-i})\|_2 \leq s,$$

where $\|\cdot\|_2$ is the standard Euclidean norm in \mathbb{R}^d .

THEOREM 5 (Dwork, Kenthapadi, McSherry, Mironov, and Naor (2006)). *Suppose a function $f : T^n \rightarrow \mathbb{R}^d$ has L^2 sensitivity s . Then the algorithm that computes $f(\mathbf{t}) + Z$ is (ϵ, γ) -differentially private, where $Z \in \mathbb{R}^d$ is a random vector where each coordinate is drawn i.i.d. from the normal distribution $N(0, \sigma^2)$ for*

$$\sigma = \frac{s}{\epsilon} \sqrt{\log \left(\frac{1.25}{\gamma} \right)}.$$

The following corollary will be useful in our analysis.

COROLLARY 1. *Suppose a function $f : T^n \rightarrow \mathbb{R}^d$ has L^2 sensitivity $s(n)$ (here we assume the sensitivity is a function of n , which should be thought of as diminishing in n). Then for any constant σ , the algorithm that computes $f(\mathbf{t}) + Z$ where $Z \in \mathbb{R}^d$ is a random vector where each is drawn i.i.d. from the normal distribution $N(0, \sigma^2)$ is (ϵ, γ) -differentially private for*

$$\epsilon + \gamma \in O \left(s(n) \sqrt{\log \frac{1}{s(n)}} \right)$$

In particular if the sensitivity $s(n) \in O(\frac{1}{n})$, we have $\epsilon + \gamma \in O(\frac{\sqrt{\log n}}{n})$

4.1 LARGE ANONYMOUS GAMES

A class of games that has received much attention in the literature is the class of Large Anonymous games. In these games, all players have the same set of available actions and each player's payoff depends only on the action he takes and the histogram of actions taken by others (and not on the identity of the players).⁶

Consider a repeated game that has such a large anonymous game as a stage game. A player's payoff from playing action $a \in A$, when the population distribution of plays by others is $\alpha \in \Delta A$ is $u_i(a, \alpha) \in [0, 1]$. Suppose further that in every period, a noisy histogram of the actions taken by all agents is announced, i.e. when the realized distribution of plays is $\alpha \in \Delta A$, the announced distribution is $\alpha + Z$ for $Z \in \mathbb{R}^{|A|}$, with each component of $Z \sim N(0, \sigma^2)$. Then, fixing players'

⁶Papers in this literature often refer to such games simply as Large games.

discount factor δ , by Corollary 1 and Theorem 1 that any perfect public equilibrium of the repeated game must involve each player playing an η -approximate Nash equilibrium of the stage game for

$$\eta \in O\left(\frac{\delta}{1-\delta} \frac{\sqrt{\log n}}{n}\right).$$

Note that as n tends large, $\eta \rightarrow 0$.

Note that a similar noise distribution results if $\alpha \in \Delta A$ is computed by *subsampling* player actions. It can similarly be shown that the distribution that results from subsampling (sufficiently to get a constant error rate in the action histogram) results in an (ϵ, γ) differential privacy guarantee with $(\epsilon + \gamma)$ tending to 0 as n grows large.

4.2 NOISY COURNOT GAMES

One of the classic empirical applications of repeated games is to study the possibility of collusion in repeated oligopolistic competition—see [Green and Porter \(1984\)](#) for a classic reference. We demonstrate via their model that in a large oligopoly, collusion may be impossible.

The stage game is n -firm Cournot competition. Each firm i simultaneously chooses quantity $q_i \in [0, 1]$ to produce, at a cost of $c(q_i)$. For simplicity in calculations, take $c(q_i) = q_i$. After all firms have selected quantities, a price is determined from the (realized, stochastic) demand as

$$p = \theta P\left(\frac{1}{n} \sum_{i=1}^n q_i\right),$$

where $P(\cdot)$ is a continuously differentiable decreasing demand function and θ is the realized demand shock. We assume that θ is distributed log-normally with $\mathbb{E}[\theta] = 1$. Each firm i then realizes a profit of its revenue less cost, i.e., $pq_i - c(q_i)$. Further, only this price p is observed publicly by all the firms—firms do not directly observe the quantities produced by other firms. Taking logarithms, we have that

$$\log p = \log \theta + \log P\left(\frac{1}{n} \sum_{i=1}^n q_i\right).$$

For n large, by a Taylor series expansion, we have that the sensitivity of $\log P(\cdot)$ is

$$\frac{1}{n} \sup_{x \in [0,1]} \frac{P'(x)}{P(x)}.$$

Note that fixing a function P , this quantity is diminishing at a rate of $O(1/n)$. Note also that since θ is log-normally distributed, $\log \theta$ is normally distributed. Therefore by Corollary 1 and Theorem 1, any Perfect Public equilibrium of the stage game must involve the play of approximate Nash Equilibrium of the stage game in every period for $\eta \in O\left(\frac{\delta}{1-\delta}\sqrt{\frac{\log n}{n}}\right)$.

4.3 NOISY COUNTERFACTUAL PAYOFFS

Now let us consider an example of an imperfect private monitoring setting. Fix a n player game, and assume that the game is μ -sensitive, i.e. for every player i and every player $j \neq i$, j changing the action he plays can affect player i 's payoff by at most μ *regardless* of what everyone else is playing:

$$\forall i \neq j, a_i, a_j, a'_j, a_{-ij} : |u_i(a_i, a_j, a_{-ij}) - u_i(a_i, a'_j, a_{-ij})| \leq \mu.$$

After each period, each player receives a private signal of a (noisy) estimate of the payoff she would have received in the past period for each of the actions she could have taken, i.e. $S_i = \mathbb{R}^{|A_i|}$, and when other players play a_{-i} , player i receives a signal where the component corresponding to action a is $u_i(a, a_{-i}) + N(0, \sigma^2)$. Such a signal structure has been widely studied in both the game theory and machine learning literatures. In the game theory literature, Hart and Mas-Colell (2000) and Foster and Vohra (1997) show that simple heuristic strategies (minimizing regret) played in this setting will result in the empirical distribution of play converging to the set of correlated equilibria, see e.g. Chapter 4 of Nisan, Roughgarden, Tardos, and Vazirani (2007) for a survey and the connections to machine learning.

We note that in a k -action game, the resulting signal distribution is over a vector of length $k \cdot n$ (reporting the payoff to each of the n players for each of their k actions). Since each action is individually μ sensitive, the ℓ_2 sensitivity of the entire vector of payoffs is $O(\sqrt{nk}\mu)$. Fixing players' discount factor δ , by Corollary 1 and Theorem 1 it follows that any perfect public equilibrium of the repeated game must involve each player playing an η -approximate Correlated equilibrium of the stage game for

$$\eta \in O\left(\frac{\delta}{1-\delta}\mu\sqrt{n \cdot k \log \frac{1}{\mu}}\right).$$

Note that as n tends large, if $\mu \ll \frac{1}{\sqrt{n}}$, $\eta \rightarrow 0$.

4.4 OTHER SETTINGS WITH PRIVATE SIGNALING

Finally, we remark that the three settings we have noted in which the signals are naturally differentially private (noisy anonymous games, noisy Cournot games, and noisy counterfactual payoffs) are not isolated examples. Differential privacy is a powerful measure of “influence” in part because it enjoys a strong “composition” property. A sequence of (ϵ, γ) -differentially private distributions (even chosen adaptively, as a function of the realized values of previous distributions) remain (ϵ', γ') -differentially private, for values of ϵ' and γ' depending on the number of compositions and the parameters ϵ and γ . Informally, the γ 's “add up” linearly, whereas the ϵ parameters increase only with the square root of the number of such compositions.⁷ Formally, we have the following theorem due to [Dwork, Rothblum, and Vadhan \(2010\)](#):

THEOREM 6 (Composition for differential privacy, [Dwork, Rothblum, and Vadhan \(2010\)](#)). *Let $M_1(D), \dots, M_k(D)$ be families of distributions such that each M_i is (ϵ_i, γ_i) -differentially private with $\epsilon_i \leq \epsilon'$. Then for any $\gamma \in [0, 1]$: $M(D) = (M_1(D), \dots, M_k(D))$ is $(\epsilon, \gamma + \sum_{i=1}^k \gamma_i)$ -private for*

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

$M(D)$ also satisfies (ϵ, γ) differential privacy for $\epsilon = \sum_{i=1}^k \epsilon_i$ and $\gamma = \sum_{i=1}^k \gamma_i$.⁸ Moreover, these compositional properties hold even if each distribution M_i is chosen adaptively, i.e. only after observing the realized values of distributions $M_1(D), \dots, M_{i-1}(D)$.

Most immediately, this means that our results hold for any combination of the signalling schemes discussed so far: for example, agents could observe *both* a noisy price, *and* a noisy histogram of actions that were played, and the resulting composite signal would remain differentially private. More generally, a large number of simple noisy “primitives” guarantee differential privacy: we have already seen that the natural operation of adding Gaussian noise to a low sensitivity vector is differentially private. However, many other distributions guarantee differential privacy as well, including the “Laplace” distribution (which is a symmetric exponential distribution, and gives an even stronger guarantee of $(\epsilon, 0)$ -differential privacy compared to Gaussian noise) ([Dwork, McSherry, Nissim, and Smith, 2006](#)). Other simple operations such as subsampling k out of n player actions guarantee $(0, k/n)$ -differential privacy.

⁷It is this property that makes differential privacy a more appealing notion of influence than mere statistical distance (which corresponds to $(0, \gamma)$ -differential privacy). If we were to talk only about statistical distance, then the degradation of the statistical distance parameter would always be linear in the number of compositions, whereas the degradation in the ϵ parameter in differential privacy can be much slower.

⁸For large values of k , this simpler statement is substantially weaker than the more sophisticated composition statement. However, it can be stronger for small values of k .

Moreover, differential privacy does not degrade with post-processing. The following fact is a simple consequence of the definition:

THEOREM 7 (Post-processing). *Let $M(D)$ be an (ϵ, γ) -differentially private family of distributions over some range \mathcal{O} , and let $f : \mathcal{O} \rightarrow \mathcal{O}'$ be any randomized mapping. Then $f(M(D))$ is an (ϵ, γ) -differentially private family of distributions over \mathcal{O}' .*

What this means is that any signal distribution that can be described as a differentially private distribution (resulting, from e.g. the composition theorem above), followed by post-processing (for example, we could have a noisy estimate of the fraction of players playing each action, *truncated and renormalized* to form a distribution) remains differentially private.

As a result of the composition theorem and the post-processing theorem, differentially private signaling distributions are not difficult to come by in natural settings. Most noisy settings can be seen to satisfy (ϵ, γ) -privacy for some finite values of ϵ and γ . Our theorems apply when these parameters ϵ and γ tend to zero as n grows large – and as we have seen, this is generally the case when the noise rate is fixed (independent of n), and the game is large.

5 CONCLUSIONS AND OPEN QUESTIONS

We have shown that in large games with (public or private) imperfect monitoring, it can naturally arise that the signaling structure satisfies the constraint of “differential privacy”, in which the privacy parameters ϵ and γ tend to 0 as the size n of the game grows large. Moreover, we have shown that whenever this is the case the equilibria of the repeated game collapse to the set of equilibria of the stage game in large populations. This conclusion holds for a broad class of equilibria of the repeated game, that e.g. need not necessarily even be subgame perfect. With public monitoring, the equilibrium set collapses to the set of *Nash* equilibria of the stage game, and with private monitoring, the equilibrium set collapses to the set of *correlated* equilibria of the stage game.

This “anti-folk-theorem” for large, noisy repeated games has several interpretations. On the one hand, it means that cooperative, high welfare folk-theorem equilibria which cannot be supported as equilibria of the stage game are also impossible to support in repeated games from this natural class. On the other hand, it also eliminates low-welfare equilibria of the repeated game that are not supported in the stage game, and so in general, *improves* the *price of anarchy* when computed over the set of Nash equilibria of the repeated game. Our theorem can also be seen as an argument that equilibria of the repeated game can be more predictive in large noisy games, than in general repeated games. By removing the bite of the folk theorem, the multiplicity of

possible equilibria of the repeated game can be greatly reduced, in many cases to a single (or small number of) Nash equilibria of the stage game.

There remain many interesting questions along this line of inquiry. For example, is there an even more general condition than differential privacy that can yield the same results? One natural candidate is *joint* differential privacy, introduced by [Kearns, Pai, Roth, and Ullman \(2014\)](#), which informally would state that in a game of private monitoring, a unilateral change in action by a single player will not substantially affect the distribution over signals observed by *other* players (but in violation of the standard differential privacy constraint, can arbitrarily affect that player’s own signal distribution).

Similarly, there is the question of whether our theorems are qualitatively tight. Certainly they are in public monitoring settings – it is always possible to support a sequence of stage game Nash equilibria as an equilibrium of the repeated game, and we have shown that this is essentially the *only* class of equilibria that can be supported in the games that we study. However, in the private monitoring setting, we have shown only that equilibria of the repeated game must consist of a series of (approximate) correlated equilibria of the stage game. But is it the case that for every stage game and every sequence of correlated equilibria of that stage game, there exists a private signaling structure that does indeed support the sequence of correlated equilibria as an equilibrium of the repeated game?

REFERENCES

- ABREU, D., AND A. RUBINSTEIN (1988): “The structure of Nash equilibrium in repeated games with finite automata,” *Econometrica: Journal of the Econometric Society*, pp. 1259–1281.
- AL-NAJJAR, N. I., AND R. SMORODINSKY (2001): “Large nonanonymous repeated games,” *Games and Economic Behavior*, 37(1), 26–39.
- BORGS, C., J. CHAYES, N. IMMORLICA, A. T. KALAI, V. MIRROKNI, AND C. PAPADIMITRIOU (2010): “The myth of the folk theorem,” *Games and Economic Behavior*, 70(1), 34–43.
- CHEN, Y., S. CHONG, I. A. KASH, T. MORAN, AND S. VADHAN (2013): “Truthful mechanisms for agents that value privacy,” in *Proceedings of the fourteenth ACM conference on Electronic commerce*, pp. 215–232. ACM.
- DWORK, C., K. KENTHAPADI, F. MCSHERRY, I. MIRONOV, AND M. NAOR (2006): “Our Data, Ourselves: Privacy Via Distributed Noise Generation,” in *EUROCRYPT*, pp. 486–503.

- DWORK, C., F. MCSHERRY, K. NISSIM, AND A. SMITH (2006): “Calibrating Noise to Sensitivity in Private Data Analysis,” in *TCC '06*, pp. 265–284.
- DWORK, C., G. N. ROTHBLUM, AND S. P. VADHAN (2010): “Boosting and Differential Privacy,” in *FOCS*, pp. 51–60.
- FOSTER, D., AND R. VOHRA (1997): “Calibrated learning and correlated equilibrium,” *Games and Economic Behavior*, 21(1-2), 40–55.
- FUDENBERG, D., D. LEVINE, AND E. MASKIN (1994): “The folk theorem with imperfect public information,” *Econometrica: Journal of the Econometric Society*, pp. 997–1039.
- GHOSH, A., AND K. LIGETT (2013): “Privacy and coordination: computing on databases with endogenous participation,” in *ACM Conference on Electronic Commerce*, pp. 543–560.
- GHOSH, A., AND A. ROTH (2013): “Selling privacy at auction,” *Games and Economic Behavior*, Preliminary Version appeared in the Proceedings of the Twelfth ACM Conference on Electronic Commerce (EC 2011).
- GREEN, E. J. (1980): “Noncooperative price taking in large dynamic markets,” *Journal of Economic Theory*, 22(2), 155–182.
- GREEN, E. J., AND R. H. PORTER (1984): “Noncooperative collusion under imperfect price information,” *Econometrica: Journal of the Econometric Society*, pp. 87–100.
- HALPERN, J., R. PASS, AND L. SEEMAN (2013): “The Truth Behind the Myth of the Folk Theorem,” .
- HARDT, M., AND A. ROTH (2012): “Beating randomized response on incoherent matrices,” in *Proceedings of the 44th symposium on Theory of Computing*, pp. 1255–1268. ACM.
- HART, S., AND A. MAS-COLELL (2000): “A simple adaptive procedure leading to correlated equilibrium,” *Econometrica*, 68(5), 1127–1150.
- KEARNS, M., M. PAI, A. ROTH, AND J. ULLMAN (2014): “Mechanism design in large games: Incentives and privacy,” in *Proceedings of the 5th conference on Innovations in theoretical computer science*, pp. 403–410. ACM.
- LEVINE, D., AND W. PESENDORFER (1995): “When are agents negligible?,” *American Economic Review*, pp. 1160–1170.

- MAILATH, G. J., I. OBARA, AND T. SEKIGUCHI (2002): “The maximum efficient equilibrium payoff in the repeated prisoners’ dilemma,” *Games and Economic Behavior*, 40(1), 99–122.
- MAILATH, G. J., AND L. SAMUELSON (2006): *Repeated games and reputations: long-run relationships*. Oxford University Press.
- MCSHERRY, F., AND K. TALWAR (2007): “Mechanism Design via Differential Privacy,” in *FOCS*, pp. 94–103.
- MEGIDDO, N., AND A. WIGDERSON (1986): “On play by means of computing machines: preliminary version,” in *Proceedings of the 1986 Conference on Theoretical aspects of reasoning about knowledge*, pp. 259–274. Morgan Kaufmann Publishers Inc.
- NISAN, N., T. ROUGHGARDEN, E. TARDOS, AND V. VAZIRANI (2007): *Algorithmic game theory*. Cambridge Univ Press.
- NISSIM, K., C. ORLANDI, AND R. SMORODINSKY (2012): “Privacy-aware mechanism design,” in *Proceedings of the 13th ACM Conference on Electronic Commerce*, pp. 774–789. ACM.
- NISSIM, K., R. SMORODINSKY, AND M. TENNENHOLTZ (2012): “Approximately optimal mechanism design via differential privacy,” in *ITCS*, pp. 203–213.
- NISSIM, K., S. VADHAN, AND D. XIAO (2014): “Redrawing the boundaries on purchasing data from privacy-sensitive individuals,” in *Proceedings of the 5th conference on Innovations in theoretical computer science*, pp. 411–422. ACM.
- PAI, M., AND A. ROTH (2013): “Privacy and Mechanism Design,” *Sigecom Exchanges*, pp. 8–29.
- PAPADIMITRIOU, C. H., AND M. YANNAKAKIS (1994): “On complexity as bounded rationality,” in *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pp. 726–733. ACM.
- ROGERS, R., AND A. ROTH (2013): “Asymptotically Truthful Equilibrium Selection in Large Congestion Games,” *arXiv preprint arXiv:1311.2625*.
- SABOURIAN, H. (1990): “Anonymous repeated games with a large number of players and random outcomes,” *Journal of Economic Theory*, 51(1), 92–110.

SUGAYA, T. (2011): “Folk Theorem in repeated games with private monitoring,” *Economic Theory Center Working Paper*, (011-2011).

XIAO, D. (2013): “Is privacy compatible with truthfulness?,” in *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pp. 67–86. ACM.

APPENDIX

of Theorem 2. Fix a sequential equilibrium $\sigma = (\sigma_1, \dots, \sigma_n)$. For every player i and every history $h_i^t = h_i^t(a_i^t, s^t)$, we can define a posterior distribution $\mathcal{D}(i, h_i^t, \sigma_{-i})$ on histories h_{-i}^t observed by the other players. We note several things: first, since the public signals s^t are fixed, $\mathcal{D}(i, h_i^t, \sigma_{-i})$ defines a distribution only over the actions a_j^t played by each of the other players $j \neq i$. Second, since the actions a_j^t played by the other players are a function only of their own past actions and the public signals s^t , they are conditionally independent of actions a_i^t fixing the public history s^t . Hence, we can write $\mathcal{D}(i, h_i^t, \sigma_{-i}) \equiv \mathcal{D}(i, s^t, \sigma_{-i})$. Finally, equivalently note that the distribution over actions a_j^t, a_k^t of any pair of players $j \neq k \neq i$ are conditionally independent fixing the public signals s^t , and so we can decompose $\mathcal{D}(i, s^t, \sigma_{-i})$ as a product distribution:

$$\mathcal{D}(i, s^t, \sigma_{-i}) = \prod_{j \neq i} \mathcal{D}_j(i, s^t, \sigma_{-i})$$

where $\mathcal{D}_j(i, s^t, \sigma_{-i})$ represents the marginal distribution on histories h_j^t of player j fixing s^t .

Fix any set of public signals s^t , and private histories (h_1^t, \dots, h_n^t) for the n players consistent with s^t . For any player i , define

$$\begin{aligned} a^* &= \arg \max_{a_i \in A_i} \mathbb{E}_{h_{-i}^t \sim \prod_{j \neq i} \mathcal{D}_j(i, s^t, \sigma_{-i})} [u_i(a_i, \sigma_{-i}(h_{-i}^t))] \\ &= \arg \max_{a_i \in A_i} \mathbb{E}_{a_{-i} \sim (\hat{\sigma}_{1|s^t}, \dots, \hat{\sigma}_{n|s^t})} [u_i(a_i, a_{-i})] \end{aligned}$$

Define σ_i^t to be the strategy that is identical to σ_i , except on history h_i^t , it plays $\sigma_i^t(h_i^t) = a_i^*$, and then play in future periods is as if an action drawn from $\sigma_i(h_i^t)$ was played in period t .

Formally, for any period $\tau > t$, with realized history h_i^τ , the deviation involves playing $\sigma_i(h_i^{\tau'})$, where $h_i^{\tau'}$ is the same as h_i^τ except in the component corresponding to the action played at time t ; i.e. $a_i^{\tau'} \sim \sigma_i(h_i^t)$ whereas $a_i^\tau = a^*$ by definition. Given the equilibrium strategies σ , we

can write:

$$V_{i,\sigma}(h_i^t) = (1 - \delta) \mathbb{E}_{a_{-i} \sim (\hat{\sigma}_{1|st}, \dots, \hat{\sigma}_{n|st})} [u_i(\sigma_i(h_i^t), a_{-i})] \\ + \delta \sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i)) P[s, a_i | h_i^t]$$

Now consider σ'_i to be the deviation we described above. Note that by playing σ'_i from history h_i^t onwards, player i 's expected discounted payoff can be written as:

$$V_{i,(\sigma_{-i}, \sigma'_i)}(h_i^t) = (1 - \delta) \mathbb{E}_{a_{-i} \sim (\hat{\sigma}_{1|st}, \dots, \hat{\sigma}_{n|st})} [u_i(a^*, a_{-i})] \\ + \delta \sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i)) P[s | h_i^t, a^*] P[a_i | \sigma(h_i^t)]$$

Since σ forms a sequential equilibrium of the repeated game, for every history h_i^t we have:

$$V_{i,\sigma}(h_i^t) \geq V_{i,(\sigma_{-i}, \sigma'_i)}(h_i^t)$$

Substituting the definitions of $V_{i,\sigma}(h_i^t)$ and $V_{i,(\sigma_{-i}, \sigma'_i)}(h_i^t)$ into this inequality, we have:

$$\left(\mathbb{E}_{a_{-i} \sim (\hat{\sigma}_{1|st}, \dots, \hat{\sigma}_{n|st})} [u_i(a^*, a_{-i})] - \mathbb{E}_{a_{-i} \sim (\hat{\sigma}_{1|st}, \dots, \hat{\sigma}_{n|st})} [u_i(\sigma_i(h_i^t), a_{-i})] \right) \\ \leq \frac{\delta}{1 - \delta} \left(\sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i)) P[s, a_i | h_i^t] - \sum_{s \in S, a_i \in A_i} V_{i,\sigma}((h_i^t, s, a_i)) P[s | h_i^t, a^*] P[a_i | \sigma(h_i^t)] \right)$$

By (ϵ, γ) -differential privacy of the private signal, and the fact that $\exp(-\epsilon) \geq 1 - \epsilon$, we therefore have that this is at most

$$\leq \frac{\delta}{1 - \delta} (\epsilon + \gamma).$$

which completes the proof. □