

LARGE REPEATED GAMES WITH UNCERTAIN FUNDAMENTALS I: COMPRESSED EQUILIBRIUM

EHUD KALAI AND ERAN SHMAYA

ABSTRACT. Due to their many applications, large Bayesian games have been a subject of growing interest in game theory and related fields. But to a large extent, models (1) have been restricted to one-shot interaction, (2) are based on an assumption that player types are independent and (3) assume that the number of players is known.

The current paper develops a general theory of Bayesian repeated large games that avoids some of these difficulties. To make the analysis more robust, it develops a concept of compressed equilibrium which is applicable to a general class of large Bayesian repeated anonymous games.

JEL Codes: C71 - Cooperative Games, C72 - Noncooperative Games and C78 - Bargaining Theory

1. INTRODUCTION

Studies of large (many players) anonymous games provide general models for the analysis of interaction in economics, political science, computer science, biology and more. Early studies of large games focused on cooperative (non strategic) games, as surveyed in the book of Aumann and Shapley [4]. Our interest here is in the more recent literature that deals with large strategic games.

Indeed, the last three decades produced substantial literature that deals with large strategic games in a variety of applications that include: markets [30], bargaining [26], auctions [32], voting [14], electronic commerce [15], market design [8, 6], and more. One generally-observed phenomenon is that model predictions are more robust when the number of players is large.

However, the analysis of such games is often subject to three restrictive assumptions. First, models are of one-shot games that exclude studies of repeated interaction. Second, preferences and information are either completely known or are statistically independent across players, again excluding many important applications. Finally, it is generally assumed

Date: October 17, 2013.

Key words and phrases. Anonymous games, Nash equilibrium, Repeated games, Large games, Bayesian equilibrium, Price taking, Rational expectations.

Ehud Kalai, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston IL 60208 and Microsoft Research New England, kalai@kellogg.northwestern.edu.

Eran Shmaya, Kellogg School of Management, Northwestern University and Department of Mathematics, Tel Aviv University, e-shmaya@kellogg.northwestern.edu.

We thank Adam Kalai for suggesting Example 1 presented in the introduction and Sylvain Sorin for helpful talks about mean field games.

that the exact number of players is known to the players and other analysts. We postpone to later in this introduction the significant exception of Green [18] and follow up literature, who do study *repeated* complete-information games.

Our current research focusses on a robust theoretical model of large repeated games of incomplete information that overcomes such limitations. The current paper reports on the properties of an equilibrium that emerges out of this model. A companion paper [22] reports on learning and stability properties of this new equilibrium.

1.1. Approaches to the study of large strategic games. Our interest is in strategic games that share the following features:

- There is a large but finite number of players, n , which may be only partially known to the players and other game analysts.
- Players are of unknown types that describe their preferences and information. Their types may be statistically correlated through unknown states of nature.
- The players are anonymous: each knows his own parameters, can only observe aggregate data regarding the opponents, and his payoff depends on this observed data and his own individual parameters and choices.
- The game is repeated m times; m may be only partially known to the players and other game analysts.

Models of large games have adopted two type of approaches.

The *continuum approach* was initiated in Schmeidler [34], who studied the set of equilibria of a game with a continuum of players, \mathcal{E}^∞ . Schmeidler's main result is that one-shot continuum games with complete information have a pure strategy equilibrium.¹

The more recent *asymptotic approach*, on the other hand, studies properties of the sets of equilibria of the n -person games, \mathcal{E}^n , as n becomes infinitely large. Kalai [21] studies large one-shot games with incomplete information and shows, under an assumption of type independence, that the equilibria of such games satisfy strong robustness properties. Additional recent contributions to this growing literature include Cartwright and Wooders [11], Gradwohl and Reingold [17, 16], Azrieli and Shmaya [7], Carmona and Podczeck [10], and Azevedo and Budish [6].

In order to connect the two approaches above, recent papers study the relationship of the $\lim_{n \rightarrow \infty} \mathcal{E}^n$ to \mathcal{E}^∞ , see Carmona and Podczecky [10] and references therein.

While the approaches above are useful for the study of large games, they fall short of describing the family of large games described above due to the lack of dynamics and the restricted informational assumptions. They are also deficient in the following respects: (1) The set of equilibria \mathcal{E}^n is computed for a game in which the number of players n is known to the all the participants. If, as we assume in the current paper, n is unknown to the

¹See Khan and Sun [23] for a (somewhat outdated) survey of the substantial follow-up literature

individual participants, then \mathcal{E}^n fails to provide a proper description of the outcomes of the game. (2) If \mathcal{E}^n does not provide a proper description of the outcomes of the n -person game, then the comparisons of $\lim_{n \rightarrow \infty} \mathcal{E}^n$ with \mathcal{E}^∞ may be meaningless.

Indeed, lack of knowledge of the number of players is a difficult issue in game theory. This was pointed out by Myerson [28, 29] who viewed it as an issue of incomplete information, and offered a clever solution for the special case in which the number of players is randomly generated by a Poisson distribution.

The continuum approach suggests an alternative way of dealing with lack of knowledge of the number of players n , by simply disallowing equilibria that depend on n . In a driving game for example, all-drive-on-the-*right* and all-drive-on-the-*left* are equilibria of the continuum game and of all the finite n player games. But 'all-drive-on-the-*right* if n is even and all-drive-on-the-*left* if it is *odd*', an equilibrium of the n player game, cannot be recognized in the continuum game.

The idea of restricting players to equilibria that do not depend on n is appealing, but the continuum approach suffers from the conceptual difficulties already mentioned, and additional technical difficulties: First, unlike a standard assumption of game theory that the players' behavior is decentralized, in the a continuum game players *joint* behavior must be mathematically measurable. Second, and more importantly perhaps, is the difficulty of performing statistical analysis with a continuum of random variables, as needed when studying mixed strategies and even pure strategies in Bayesian games. In fact, there are conceptual difficulties in even defining the notion of mixed strategies in the continuum setup.

The equilibria studied in this paper also disallow behavior that depends on the number of players. But keeping the number players finite, it avoids the conceptual and technical difficulties above. Before describing the equilibrium notion, we introduce a model of an anonymous game that does not include the number of players and periods as fixed primitives.

1.2. Anonymous games of proportions. We study anonymous symmetric games with *abstract* number of players and *abstract* number of repetitions. The basic model starts with a fixed *skeleton* game that specifies the characteristics of an individual player (actions, types, etc.) and his payoff dependency on the *proportions* of types and actions of his opponents. The skeleton game is then augmented, to yield a family of standard symmetric repeated games, by specifying pairs of positive integers n and m that represent the number of players in the game and the number of times the game is repeated.

One advantages of this model is in the ease of performing asymptotic analysis, as one can vary the number of players and repetitions while holding "everything else" fixed. It also permits a definition of *skeleton repeated-game strategies* that do not depend on the number of opponents and repetitions. Such strategies enable the construction of compressed equilibrium, described below, that does not depend on the number of players and repetitions.

The concluding section of the paper shows how the assumption of player symmetry may be relaxed.

1.3. Compressed optimization. The *compressed* view combines ideas from both the asymptotic and the continuum approaches. Briefly stated, a payoff maximizing player facing a large number of conditionally-independent opponents may use *compressed probability assessments* for events in the game by making the following simplifying assumptions: (1) Similar to the price taking assumption in economics, his actions do not affect the general outcomes of the game. In addition (2), the player replaces probability distributions about the opponents' random empirical distribution of types and actions by their expected values. Under *compressed optimization* the player chooses strategies that maximize his expected payoffs relative to his compressed probability assessments.

While the above view encompasses modeling benefits of the continuum model, it bypasses the technical and conceptual difficulties discussed above. Some of the gains are illustrated by the following example.

Example 1. Optimal Charity Contribution. Each of $n + 1$ individuals, $i = 0, 1, \dots, n$, is about to make a contribution to a certain charity at possible dollar levels $x^i = 0, 1, 2, \dots, d$. Wishing to be at the 0.9 quantile of all contributors but not higher, Pl.0's VM utility function is $u(x^0, \mathbf{e}) = 1 - (\widehat{\mathbf{e}} \langle x^0 \rangle - 0.9)^2$, where \mathbf{e} describes the empirical distribution of opponent's contributions and $\widehat{\mathbf{e}}$ is its cumulative distribution, i.e., $\widehat{\mathbf{e}} \langle x^0 \rangle$ is the proportion of opponents who contribute x^0 or less.

If Pl.0 knew \mathbf{e} , he would choose x^0 to be a 0.9 quantile of it. But not knowing \mathbf{e} , he assumes that there is a vector of independent probability distributions $\bar{f} = (f^1, \dots, f^n)$ that describes the probabilities of individual contributions, i.e., $f^i(x^i)$ is the probability that contributor i contributes $\$x^i$. To choose his contribution optimally, Pl.0 should select an x^0 that maximizes:

$$\sum_{\mathbf{e}} u(x^0, \mathbf{e}) \mathcal{P}_{\bar{f}}(\mathbf{e}) = \sum_{\mathbf{e}} [1 - (\widehat{\mathbf{e}}_n \langle x^0 \rangle - 0.9)^2] \mathcal{P}_{\bar{f}}(\mathbf{e}).$$

The sum is taken over all possible empirical distributions of contributions \mathbf{e} that can be generated by n contributors, and $\mathcal{P}_{\bar{f}}(\mathbf{e})$ is the computed probabilities of obtaining the empirical distributions \mathbf{e} when the individual contributions are drawn according to \bar{f} .

Under compressed optimization, on the other hand, Pl.0 compresses the probabilities of contributions in \bar{f} to their average, $\kappa_{\bar{f}}(x) \equiv \frac{1}{n} \sum_i f^i(x)$, to obtain the expected proportion of opponents who contribute at every level x . For the cumulative distribution $\widehat{\kappa}_{\bar{f}}$ he selects an x^0 that maximize:

$$u(x^0, \kappa_{\bar{f}}) = 1 - (\widehat{\kappa}_{\bar{f}} \langle x^0 \rangle - 0.9)^2$$

In other words, he simply chooses x^0 to be a 0.9 quantile of $\kappa_{\bar{f}}$.

As discussed in this paper, under sufficient conditions the simplification in Example 1 is valid, i.e., it is asymptotically optimal for large n 's. But there may be other uncertainties that cannot be eliminated by compression to expected values. For example, suppose Pl.0 believes that there are two possible profiles of contribution probabilities, \bar{f}_1 and \bar{f}_2 ; and that based on some unknown state of nature his opponents make their choices according to the profile \bar{f}_1 , with probability θ_1 , or according to the profile \bar{f}_2 , with probability $\theta_2 = 1 - \theta_1$.

Now, the optimal contribution of Pl.0 is any x^0 that maximizes:

$$\theta_1 \sum_{\mathbf{e}} u(x^0, \mathbf{e}) \mathcal{P}_{\bar{f}_1}(\mathbf{e}) + \theta_2 \sum_{\mathbf{e}} u(x^0, \mathbf{e}) \mathcal{P}_{\bar{f}_2}(\mathbf{e}),$$

and his *compressed* optimal choice is any x_0 that maximizes:

$$\theta_1 u(x^0, \kappa_{\bar{f}_1}) + \theta_2 u(x^0, \kappa_{\bar{f}_2}).$$

In other words, conditioned on θ , he still compresses the random empirical distributions to their expected values. However, he cannot compress the uncertainty from the unknown θ .

1.4. Compressed equilibrium. An assumption that all the players perform compressed optimization leads to a well defined and easier to compute notion of *compressed equilibrium*. The gained simplification from the substitutions of expected values is especially important for *repeated* games. Here, in addition to simplifying the assessments of payoffs, the players and analysts simplify the process of updating beliefs, as they transit from one stage of the game to the next.

These transitions are one of the issues that make the study of repeated large games more challenging than studies of one-shot large games. Since the outcome of a stage is used by the players to determine their next-stage actions, it is important that the outcome observed at the end of each stage is continuous as the number of players increases. For this reason, the model presented in this paper assumes that there is smoothing of publicly-observed period outcomes in a form of noise or other random variables. This idea is common in game theory, tracing back at least to Green's paper [18]. See also Duggan's recent proof of existence of stationary equilibrium in stochastic games with noise [13].

1.4.1. Myopicity. Another property that carries over from Green's model to the current paper is that compressed equilibria are *myopic*. Specifically, in checking whether his strategy is optimal, it suffices for a player to check it only on a period-by-period basis: Are the actions he chooses for any given period optimal for that period (based on the information he has going into it and ignoring his payoffs in future periods)? Clearly, myopicity offers a significant simplification in the analysis of equilibrium.

1.4.2. Uniform properties of compressed equilibria. Under the assumptions that the number of players or repetitions in the game is unknown, it is important to have strategic concepts

that do not depend on these parameters. Indeed, in the family of games studied in this paper, compressed equilibria possess uniformity properties in the number of players and in the number of repetitions.

With regards to the number of players, a compressed equilibrium in a game with n players is a compressed equilibrium of the same game with any number of players (even though it is a good approximation to real equilibrium only when n is large).

With regards to the number of repetitions, a strategy profile is a compressed equilibrium in a game repeated m times, if and only if it is a compressed equilibrium of the game repeated $m' \leq m$ times.

Combining and extending the above, we show the existence of compressed equilibria that are uniform in both, the number of players and the number of repetitions of the game.

1.4.3. Validation of the compressed view. Are compressed concepts useful for the analysis of repeated large anonymous games? The answer is positive in the following sense. First, a main theorem in this paper shows that assessed compressed probabilities of play paths of the game become uniformly accurate as the number of players increases. This means that compressed best-response analysis is essentially valid for all sufficiently large n 's: If strategy g is better than strategy h in the compressed sense, then it is essentially better in the real sense, and uniformly so for all sufficiently large values of n . Among several implications, every compressed equilibrium of the repeated game is an approximate equilibrium for all the versions of the game with sufficiently many players.

1.4.4. Universality of compressed equilibria. The validation argument above provides a strong rationale for playing a compressed equilibrium in a fixed repeated game when a player knows that the number of opponents is sufficiently large. However, for such strong rationale to hold, the needed "sufficiently large" number of players must increase as the number of repetitions of the game increases.

The reason for this monotonicity stems from competing implications of laws of large numbers. On the positive side, in every stage of the game laws of large numbers permit the player to replace the aggregate play of the opponents by their expected values. But no matter how many opponents the player has, there may still be a small probability that this substitution fail. This means that for any fixed number of opponents, as the game is made longer the probability that the approximation fail in at least one of the periods may increase to one.

To overcome the above difficulty, one may adopt different approaches. One possibility is to try to weaken the notion of correct approximation. For example, it may be sufficient in applications to have correct approximations in a high proportion of periods, rather than in every period.

In this paper, however, we keep the criterion of sure correct approximation for all periods of play, and make up for it by requiring more precise period approximation. In particular,

this is obtained if for larger number of repetitions we restrict our attention to games with larger number of players, as described below.

Consider a family of games G , with various values for the number of players n and for the number of repetitions m . A profile of strategies κ is said to be (asymptotically) *universal equilibrium for the family G* , if for every ε there is a critical number of players N_ε such that in any game in G with more than N_ε players, all the players are epsilon best responding (in the standard sense) by using their κ strategies. In other words, when playing a universal equilibrium, a player is assured that he is epsilon optimizing uniformly for all sufficiently large n 's, no matter which game in the family is played.

We consider families of games in which $n > m^{2+\delta}$ for arbitrarily small positive δ 's. For any such family we show that all the compressed equilibria are universal in the sense above.

1.5. Relationship to Green [18]. Green studies pure strategies in anonymous, repeated, complete-information games. He formalizes the idea of price taking in market games, and introduced a general framework of repeated games with random outcome. Green and Sabourian [33] derive conditions under which the Nash Correspondence is continuous, i.e., that equilibrium profile in the non-atomic game is a limit of equilibrium profiles in the games with increasing number of players. Advances of the current paper over Green's include the following features: incomplete information, mixed strategies, lack of knowledge of the number of opponents and repetitions; and avoidance of technical difficulties of the continuum.

1.6. Relationship to Mean Field Games. Some of the ideas presented in the current paper overlap with ideas presented in the recent literature on mean field games. This literature studies dynamic stochastic games in large populations of players, mostly in differential game settings, see the recent survey in Lions' lecture notes [24]. But there are major differences between our goals and methodology and those of the mean field literature.

Consider a large population of players in which each player controls the trajectory of his state, subject to idiosyncratic noise. Each player receives a payoff that depends on his own state and the empirical distribution of states in the population. The mean field modeling produces a pair of differential equations: a forward Kolmogorov equation that reflects the evolution of the empirical distribution of agents' states, and a backward Hamilton-Jacobi-Bellman equation that reflects the individual optimization done by the agents. The Kolmogorov equation relies on the assumption that the idiosyncratic noises of the players cancel out, and the Hamilton-Jacobi-Bellman equation relies on the assumption that individual's deviation do not affect this empirical distribution. The mean-field modeling is of course related to Schmeidler's continuum of players approach. In fact, Jovanovic and Rosenthal [19] have extended Schmeidler's one shot setup to discrete time complete-information stochastic game. See also Adlakha and Johari [?] for a recent paper in that direction.

The similarity to our compressed equilibrium approach is evident, but there is a difference in the interpretation. For us the game has a finite number of players and the compression reflects a heuristic performed by the players, whereas the mean field literature is closer in spirit to the earlier literature on games with continuum of players in that the compression is part of the definition of the game. We prefer the term ‘compressed equilibrium’ over ‘mean field games’ in order to emphasize that we are looking at an n -player games and propose an equilibrium concept for these games.

The difference in interpretation induces differences in the type of questions we ask. The focus of much of the mean field games literature has been the existence and uniqueness of solutions to the system of differential equations. In our discrete time setup existence of a solution is simple, and, as typical in game theory, uniqueness is not obtained in many interesting cases. We focus in this paper on validation of the compressed heuristic done by the players, i.e., on the extent to which their predictions agrees with what actually happens in the n -player games. The conditions needed for validation, the relationship between the number of players and the length of the game, and the introduction of smoothing noise on the outcome observed by the players, all central in the current paper, play no role in the mean field literature.

More importantly, it is not clear whether the mean field approach can incorporate an important issue in our paper: dealing with uncertainty about fundamentals that is not eliminated in the average. The issues of Bayesian updating and learning of the fundamentals, which are central to this paper, have no analogues in the current mean field games literature.

2. THE MODEL, NOTATIONS, AND CONVENTIONS

2.1. Games and strategies. Throughout the paper we restrict ourselves to a fixed game-skeleton Γ , used to construct repeated games $\Gamma^{m,n}$ that vary in the number of players n and the number of repetitions m . The game skeleton and the constructed games are all (ex-ante) symmetric in the players.

Definition 1. [game] A *game skeleton* is an eight tuple of components $\Gamma = (S, \theta, T, \tau, A, X, \chi, u)$ with the following interpretation:

- S is a finite set of possible *states of nature*; $\theta \in \Delta(S)$ is a *prior* probability distribution over S .²
- T is a finite set of possible player’s *types*; $\tau : S \rightarrow \Delta(T)$ is a stochastic *type-generating function*.
- A is a finite set of possible player’s *actions*.
- X is a Borel space of *outcomes*, and $\chi : S \times \Delta(T \times A) \rightarrow \Delta(X)$ is a Borel function that generates probability distributions over outcomes.

²As usual, $\Delta(B)$ is the set of all probability distributions over the set B .

- $u : T \times A \times X \rightarrow [0, 1]$ is a Borel function that describes player's *payoffs*.

The game $\Gamma^{m,n}$ has a set of n players, $N = \{0, 1, \dots, n-1\}$, and is played repeatedly in m stages as follows: First a state of nature $s \in S$ is randomly chosen according to θ . Then every player is informed of his type, which is drawn randomly and independently of the other players according to the distribution τ_s . Continuing inductively, at every stage $k = 0, 1, \dots, m-1$ every player chooses an action from A , and a public outcome x_k is randomly chosen according to χ_{s,e_k} , where $d_k \in \Delta(T \times A)$ is the realized empirical distribution of players type and period actions:

$$d_k[t, a] = (\text{the number of players of type } t \text{ who chose the action } a \text{ at day } k)/n.$$

All the players are informed of the realized outcome x_k before proceeding to the next stage.

We use the notations $\Gamma^{m,\cdot}, \Gamma^{\cdot,n}, \Gamma^{\cdot,\cdot}$, to denote respectively games with arbitrary number of players, with arbitrary number of repetitions and with arbitrary number of both.

This framework extends the model in the papers of Green and Sabourian in several ways. In particular, it allows for incomplete information and for the use of mixed strategies as described next.

A repeated-game strategy for $\Gamma^{\cdot,\cdot}$ (or a *strategy* for short) is a function $f : T \times (A \times X)^{<\mathbb{N}} \rightarrow \Delta(A)$.³ It describes the probabilities by which a player chooses period actions, based on his type and the history of his own past actions and observed outcomes. Specifically, viewed as a behavioral strategy, $f(t, a_0, x_0, \dots, a_{k-1}, x_{k-1})[a]$ is the probability that in the k -th period he chooses the action a , given that he is of type t , and in the previous periods he took the actions a_0, \dots, a_{k-1} and observed the outcomes x_0, \dots, x_{k-1} .

Notice that as defined, a strategy f may be used by any player in any game with any number of opponents and any number of repetitions.

A *reactive strategy* is one in which the player period choice of an action does not depend on his own past choices, i.e., $f : T \times X^{<\mathbb{N}} \rightarrow \Delta(A)$,⁴ see Kalai and Stanford [20] who introduced this concept for repeated games with complete information.

2.2. Nash Equilibrium. To define a Nash equilibrium profile of strategies $\bar{f} = (f^0, \dots, f^{n-1})$, it suffices to define the payoffs of a player under unilateral or no deviations from the profile, as done below.

However, to simplify notations below and later in the paper, we write the definitions only for one representative player, player 0, with the understanding that the same definitions apply to all the players. Players $1, \dots, n-1$ are referred to as player 0's 'opponents'.

³For any set B , $B^{<\mathbb{N}} = \bigcup_{m=0}^{\infty} B^m$.

⁴In much of what follows, the restriction to reactive strategies can be eliminated at the cost of using more complex notations.

Also below and later in the paper, we use bold face letters to denote random variables that assume values in corresponding sets. For example, \mathbf{S} is the random variable that describes a randomly-selected state from the set of possible states S . We use superscripts for players' names and subscripts for period numbers.

Let $\bar{f} = (f^0, \dots, f^{n-1})$ be a reactive strategy profile. Assume that player 0 considers playing a (not necessarily reactive) strategy g and the opponents follow the profile \bar{f} . A (random) (g, \bar{f}) -play in the n -player game is a collection $(\mathbf{S}, \mathbf{T}^i, \mathbf{A}_k^i, \mathbf{X}_k)_{i \in N, k=0,1,\dots}$ of random variables, representing the state of nature, types, actions and outcomes, such that:

- The state of nature \mathbf{S} is distributed according to θ .
- Conditional on \mathbf{S} , the players types \mathbf{T}^i are i.i.d with the distribution $\tau_{\mathbf{S}}$.
- Conditional on the history of stages $0, \dots, k-1$, players choose period k actions \mathbf{A}_k^i independently. Player 0 uses the distribution $g(\mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1})$ and each of his opponents $i \in N \setminus \{0\}$ uses the distribution $f^i(\mathbf{T}^i, \mathbf{X}_0, \dots, \mathbf{X}_{k-1})$.
- The outcome \mathbf{X}_k of day k is drawn randomly according to $\chi_{\mathbf{S}, \mathbf{d}_k}$, where

$$(2.1) \quad \mathbf{d}_k[t, a] = \#\{i \in N \mid \mathbf{T}^i = t, \mathbf{A}_k^i = a\} / n$$

is the (random) empirical type-action distribution at stage k .

In equations,

$$(2.2) \quad \begin{aligned} \mathbb{P}(\mathbf{S} = s, \mathbf{T}^i = t^i \ i \in N) &= \theta[s] \cdot \prod_{i \in N} \tau_s[t^i]. \\ \mathbb{P}(\mathbf{A}_k^i = a^i \ i \in N \mid \mathbf{S}, \mathbf{T}^i, \mathbf{A}_l^i, \mathbf{X}_l \ l < k, i \in N) &= \\ g(\mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1})[a^0] \cdot \prod_{i \in N \setminus \{0\}} f^i(\mathbf{T}^i, \mathbf{X}_0, \dots, \mathbf{X}_{k-1})[a^i] \\ \mathbb{P}(\mathbf{X}_k \in \cdot \mid \mathbf{S}, \mathbf{T}^i, \mathbf{A}_l^i \ l \leq k, i \in N, \mathbf{X}_0, \dots, \mathbf{X}_{k-1}) &= \chi_{\mathbf{S}, \mathbf{d}_k}, \end{aligned}$$

where \mathbf{d}_k is given by (2.1). Define

$$(2.3) \quad U_{m,n}^{g, \bar{f}} = \frac{1}{m} \mathbb{E} \sum_{k=0}^{m-1} u(\mathbf{T}^0, \mathbf{A}_k^0, \mathbf{X}_k)$$

to be player 0's payoff in $\Gamma^{m,n}$.

A reactive profile \bar{f} is an ϵ -equilibrium ($\epsilon \geq 0$) in $\Gamma^{m,n}$, if $U_{m,n}^{f^0, \bar{f}} + \epsilon \geq U_{m,n}^{g, \bar{f}}$ for every strategy g , and with the same restriction applied to all the other players.

3. THE COMPRESSED VIEW OF A GAME

The compressed view of a game $\Gamma^{m,n}$ is a simplified way by which players and other analysts may assess the probabilities of plays and the associated payoffs. In addition to their simplicity, concepts used under this approach are uniform in the number of players and the number of repetitions of the game. For example, as discussed later in this section, there

are natural repeated game strategies that are compressed equilibrium of all the games $\Gamma^{m,n}$ with any values of m and n .

In the sections that follow, we argue that the compressed view is effective for the purpose of analysing large games: As the number of players increases, compressed probabilities and payoffs become accurate in approximating their real counterparts. This means that for large games compressed best-response and compressed equilibrium are valid concepts in the analysis of how to play a game. Moreover, combined with their uniform applicability described in the previous paragraph, they present a robust tool for the analysis of large repeated games, especially when the number of players or the number of repetitions is not fully known.

As discussed later in the paper, the compressed view exhibits natural connections to models of non-atomic games with a continuum of players of the type studied by Schmeidler [34] for one-shot games with complete information. But a major difference is that the compressed view applies directly to n -person games, without being sidetracked to imaginary games with a continuum of players. In doing so, it avoids mathematical difficulties associated with the continuum, e.g., measurability conditions and continuum of independent random variables.

Described by the formal definitions below, the compressed view adopts the following simplifying assumptions about each player in the game: (1) In spirit similar to price-taking behavior in economic models, the player's actions have negligible effect on the probabilities of events related to current and future outcomes of the game. (2) In computing probabilities of events in the game, the player replaces random empirical distributions of types and actions with their expectations, while leaving the rest of the analysis unchanged.

3.1. compressed strategies. Most of the compressed computations are presented for symmetric profiles of strategies. But the analysis is applicable for non symmetric profiles. As we show below, every non symmetric profile can be replaced by a symmetric one that has the same probability distributions and payoffs from the view of every individual player.

Definition 2. For a profile of reactive strategies $\bar{f} = (f^0, \dots, f^{n-1})$ in Γ^n the compressed strategy of \bar{f} is defined by: $\kappa_{\bar{f}}(t, x_0, \dots, x_{k-1}) = \frac{1}{n} \sum_{i=0}^{n-1} f^i(t, x_0, \dots, x_{k-1})$.

The following direct observations are useful in the interpretation and forthcoming discussion of compressed strategies and equilibrium:

- Being the average of the probabilities of taking an action a , $\kappa_{\bar{f}}(t, x_0, \dots, x_{k-1})[a]$ is the *expected proportion of a -choosers* among the players of type t after they observe the outcomes x_0, \dots, x_{k-1} .
- $\kappa_{\bar{f}}$ itself is a reactive strategy.

Definition 3. For a reactive strategy g define the symmetric n -person version of g by $\bar{g}^n = (g, \dots, g)$.

- Every reactive strategy g is a compressed strategy. In particular, for any number of players n , $g = \kappa_{g^n}$.
- Since strategies in our model are defined unconditionally on the number of players and repetitions, any reactive strategy g is a compressed strategy in all the games Γ .
- For any reactive/compressed strategy κ we may think of the class of strategy profiles $[\kappa]$ which are equivalent to each other in the sense that they all compress to κ . Statements in the sequel that refer to κ may refer to all the profiles in $[\kappa]$.
- Despite the fact that a reactive strategy and a compressed strategy are the same mathematical object, in the discussion that follows we selectively use one term over the other to emphasize the relevant interpretations in different contexts.

3.2. Compressed probabilities and payoffs. We consider a game in which the players use reactive strategies that compress to the strategy κ , and a representative player, player 0, considers the use of some (reactive or not) strategy g . As before, we denote the state of nature by the random variable \mathbf{S} , player 0's type and period actions by \mathbf{T}^0 and by $\mathbf{A}_0^0, \mathbf{A}_1^0, \dots$, and the period outcomes by the variables $\mathbf{X}_0, \mathbf{X}_1, \dots$.

Definition 4. Let κ be a compressed strategy and g a strategy. The *compressed (g, κ) -play* is a sequence of random variables $\mathbf{S}, \mathbf{T}^0, \mathbf{A}_0, \mathbf{X}_0, \mathbf{A}_1, \mathbf{X}_1, \dots$, with the *compressed probability distribution* defined inductively by:

$$(3.1) \quad \begin{aligned} \mathbb{P}(\mathbf{S} = s, \mathbf{T}^0 = t^0) &= \theta[s] \cdot \tau_s[t^0], \\ \mathbb{P}(\mathbf{A}_k^0 = a^0 | \mathbf{S}, \mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1}) &= g(\mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1})[a^0], \\ \mathbb{P}(\mathbf{X}_k \in \cdot | \mathbf{S}, \mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1}, \mathbf{A}_k^0) &= \chi_{\mathbf{S}, \mathbf{e}_k}(\cdot), \end{aligned}$$

where $\mathbf{e}_k \in \Delta(T \times A)$ is the random empirical distribution of type and actions in period k :

$$(3.2) \quad \mathbf{e}_k[t, a] = \tau_{\mathbf{S}}[t] \kappa(t, \mathbf{X}_0, \dots, \mathbf{X}_{k-1})[a] \text{ for every } (t, a) \in T \times A.$$

Notice that in Equation 3.2 the random variables that represent the state of nature, player 0 type, and actions and the outcomes are distributed as in an n -player random (g, κ) -play. Only the random frequencies of realized opponents' types and the random frequencies of the actions chosen by the opponents in each period k are replaced by their deterministic expected values, i.e., the theoretical distributions $\tau_{\mathbf{S}}$ and $\kappa(t, \mathbf{X}_0, \dots, \mathbf{X}_{k-1})$.

Remark 1. The notion of compressed play expressed in (3.2) captures the two notions of simplification mentioned above.

First, since g is not included in the expressions $\mathbf{e}_k[t, a]$, Player 0 is a *stochastic outcome taker*: In addition to not influencing the state of nature \mathbf{S} , his actions do not affect the probabilities of $\mathbf{X}_0, \mathbf{X}_1, \dots$. This is a generalization of the economic price-taking property in Green's (1978) to a *stochastic* setup. Moreover, the application of the concept is not restricted to market games.

Second, the expression above significantly simplifies the stochastic process by replacing random empirical distributions with their expected values. In particular, the expressions for the $\mathbf{e}_k[t, a]$ s are based on the assumption that the random proportion of players with a type-action pair $[t, a]$ must equal its theoretical probabilities given by τ and κ .

Definition 5. For κ and g as above, define player 0's *compressed payoff* in the m -stage game $\Gamma^{m, \cdot}$ by

$$(3.3) \quad U_{m, \text{COMP}}^{g, \kappa} = \frac{1}{m} \mathbb{E} \sum_{k=0}^{m-1} u(\mathbf{T}^0, \mathbf{A}_k^0, \mathbf{X}_k),$$

where $\mathbf{S}, \mathbf{T}^0, \mathbf{A}_0, \mathbf{X}_0, \mathbf{A}_1, \mathbf{X}_1, \dots$ is a compressed (g, κ) -play.

The definition above is naturally extended to profiles of strategies \bar{f} by substituting $\kappa_{\bar{f}}$ for κ .

3.3. Compressed best response and equilibrium.

Definition 6. Let κ be a compressed strategy. We say that a strategy g is a compressed best response to κ if for all strategies h , $U_{m, \text{COMP}}^{g, \kappa} \geq U_{m, \text{COMP}}^{h, \kappa}$.

Since players are outcome takers and since in our model the payoff function is separable, it follows that a compressed best response is always myopic, so player 0 doesn't have to include the length of the game m and the history of the play in his period considerations. Thus, we get the following equivalent characterization of compressed best-response strategy:

Proposition 1. *Let κ be a compressed strategy. A strategy g is compressed best response to κ if and only if for every type t and for almost every sequences of outcomes x_0, \dots, x_{k-1} and actions a_0, \dots, a_{k-1} (with $k < m$) one has*

$$[g(t, a_0, x_0, \dots, a_{k-1}, x_{k-1})] \subseteq \operatorname{argmax}_{a \in A} \mathbb{E} \left(u(t, a, \mathbf{X}_k \mid \mathbf{T}^0 = t, \mathbf{X}_0 = x_0, \dots, \mathbf{X}_{k-1} = x_{k-1}) \right),$$

where $\mathbf{S}, \mathbf{T}^0, \mathbf{X}_0, \dots, \mathbf{X}_{k-1}$ is the compressed κ -play and $[g(t, a_0, x_0, \dots, a_{k-1}, x_{k-1})]$ is the support of $g(t, a_0, x_0, \dots, a_{k-1}, x_{k-1})$.

In particular, an outcome-taking player always have reactive best-response strategies and there is no significant loss in restricting the definition of compressed equilibrium to reactive strategies.

Definition 7. A reactive/compressed strategy κ is a *compressed equilibrium* of the m -stage repeated game $\Gamma^{m, \cdot}$, if κ is compressed best response to itself, i.e., $U_{m, \text{COMP}}^{\kappa} \equiv U_{m, \text{COMP}}^{\kappa, \kappa} \geq U_{m, \text{COMP}}^{g, \kappa}$ for every strategy g .

Notice that the notions of compressed payoffs, compressed best response, and compressed equilibrium do not depend on the number of players. In particular, in order to play compressed-optimally a player does not have to know the number of opponents and their individual strategies. All he has to know is the aggregate data expressed by the compressed strategy.

The following equivalent characterization of compressed equilibrium follows from Proposition 1.

Corollary 1. *A reactive strategy κ is compressed equilibrium of $\Gamma^{m\cdot}$ if and only if for every type t and almost every sequence of outcomes x_0, \dots, x_{k-1} (with $k < m$) one has*

$$[\kappa(t, x_0, \dots, x_{k-1})] \subseteq \operatorname{argmax}_{a \in A} \mathbb{E}(u(t, a, \mathbf{X}_k \mid \mathbf{T}^0 = t, \mathbf{X}_0 = x_0, \dots, \mathbf{X}_{k-1} = x_{k-1})),$$

where $\mathbf{S}, \mathbf{T}^0, \mathbf{X}_0, \dots, \mathbf{X}_{k-1}$ is the compressed κ -play and $[\kappa(t, x_0, \dots, x_{k-1})]$ is the support of $\kappa(t, x_0, \dots, x_{k-1})$

Corollary 1 is analogous of Theorem 5 in Green's paper [18], but it is extended for the case of incomplete information: at every stage the player's plays an equilibrium in the corresponding one-shot game with incomplete information. Green [18], Sabourian [33] and Al-Najjar and Smorodinsky [2] prove, under various condition, that this property (which Al-Najjar and smorodinsky call *myopic play*) holds approximately for every equilibrium of the repeated game. For our purpose, it is enough to prove the myopic property in the compressed view of the game, which is significantly simpler.

Again, the definitions above, applicable to a single compressed strategy κ , are applicable to any profile \bar{f} that compresses to κ . In particular, we have the following explicit definition.

Definition 8. A profile of strategies $\bar{f} = (f_1, \dots, f_n)$ is a compressed equilibrium if its compression $\kappa_{\bar{f}}$ is a compressed equilibrium.

Remark 2. Note that for every reactive profile $\bar{f} = (f_1, \dots, f_n)$ one has $[\kappa_{\bar{f}}(t, x_0, \dots, x_{k-1})] = \bigcup_i [f_i(t, x_0, \dots, x_{k-1})]$. It follows from Proposition 1 that $\bar{f} = (f_1, \dots, f_n)$ is compressed equilibrium if and only if each f_i is a best response to the compressed strategy $\kappa_{\bar{f}}$. And even more explicitly, to play optimally a player needs only to check that any action he intends to take in every period is myopically optimal relative to $\kappa_{\bar{f}}$.

Definition 9. A compressed equilibrium κ is *uniform in the number of periods*, if it is a compressed equilibrium of $\Gamma^{m\cdot}$ for every $m \geq 0$. It is uniform in the number of players if the profile $\bar{\kappa} = \kappa^n$ is a compressed equilibrium in the n -player game for every n .

This type of uniformity is reminiscent of the notion of uniform equilibrium, studied by Aumann and Maschler [5] for games with a high but unknown discount parameter. However, the uniformity presented here is stronger since there is no restriction to large m 's to parallel their restriction to high discount parameters.

Proposition 2. *There exists a compressed equilibrium uniform in both, the number of players and the number of periods .*

Before proceeding with important properties of the compressed model, we compare it to an older model of large games.

4. THE COMPRESSED VIEW VS. CONTINUUM MODEL

To highlight the relationship between the compressed view and the continuum model, it is useful to consider, as the continuum model literature generally does, a one-shot game with complete information.

Schmeidler [34] starts with a measurable space I of players equipped with a non-atomic measure λ . He defines the notion of a strategy profile as a measurable function $f : I \rightarrow \Delta(A)$ and assume that the payoff for every player depends on his own action and the aggregation $\int f(i) \lambda(di)$. Following Schmeidler, Rashid [31] and Carmona [9] study an n -player analogue of Schmeidler's non-atomic game, where every player i chooses a distribution $f(i) \in \Delta(A)$ and then the payoff of every player depends on his own action and the aggregation $\frac{1}{n} \sum_i f(i)$. Dubey, Mas-Colell and Shubik [12] build on Schmeidler's model of game with continuum of players to show that the noncooperative equilibrium outcome of a market game is a competitive equilibrium.

One fundamental difference between the papers above (abbreviated by SRC) and our approach is the following: In the SRC papers players' payoffs depends on the aggregation of the players' *mixtures*, no randomization takes place. In contrast, in our framework the payoffs depend on the aggregation of the random *realized actions*. As in standard game theory, our players choose their actions independently (conditional on past observations) according to their individual strategies. The compressed view, according to which the players assume that the aggregation of the mixtures is the realized aggregation of pure actions, only represents the way the players perceive the game.

The mathematical results in the following sections present strong justifications for this perception. On the other hand the SRC, and the Mas-Collel paper described below, take these type of results foregranted, without explicit statements and justifications.

Mas-colell [27] formulates Schmeidler's Theorem in terms joint empirical distribution of types and actions. He does not explicitly define a strategy, but every such distribution is equivalent to a strategy in the sense of this paper: the strategy is the conditional distribution of actions given types.⁵ Under this equivalence, a compressed equilibrium is an equilibrium as defined by Mas-Colell.

⁵This is similar to the equivalence between distributional strategies of Milgrom and Weber and behavioral strategies.

In addition to the above, there are the differences already mentioned: (1) We work entirely within the finite-number-of-players framework. (2) We allow for mixed strategies and incomplete information with interdependent types, and (3) we study repeated, as opposed to one-shot games.

5. VALIDATION OF THE COMPRESSED VIEW

Returning to the earlier setting, let $\bar{f} = (f^0, \dots, f^{n-1})$ be a reactive profile of strategies for the game $\Gamma^{m,n}$ with compression $\kappa (= \kappa_{\bar{f}})$. We assume that one of the players, say player 0, considers playing any (reactive or not) strategy g and that the opponents follow the profile \bar{f} . For every m denote by $\mathcal{P}_{m, \text{COMP}}^{g, \kappa}$ the joint distribution over the sequence of random outcomes of the compressed κ -play given by (3.1), and denote by $\mathcal{P}_{m, n}^{g, \bar{f}}$ the joint distribution over the sequence of random outcomes of the actual n -player (g, \bar{f}) -play given in (2.2).⁶

We say that the outcome generating function χ of Γ is Lipschitz, if for some constant L , for every s , $\|\chi(s, e) - \chi(s, e')\| \leq L \cdot \|e - e'\|_1$ where $\|\chi(s, e) - \chi(s, e')\|$ is the total variation distance between $\chi(s, e)$ and $\chi(s, e')$.⁷

Theorem 1. *Let \bar{f} , κ and g be as above, and assume that the outcome generating function χ is Lipschitz. Then*

$$\|\mathcal{P}_{m, \text{COMP}}^{g, \kappa} - \mathcal{P}_{m, n}^{g, \bar{f}}\| < C \cdot m \sqrt{\frac{\log n}{n}},$$

where C is a constant that depends on the parameters of the game skeleton (the number of actions, the number of types and the Lipschitz constant), and not on m and n .⁸

Here $\|\mathcal{P}_{m, \text{COMP}}^{g, \kappa} - \mathcal{P}_{m, n}^{g, \kappa}\|$ is the total variation distance between $\mathcal{P}_{m, \text{COMP}}^{g, \kappa}$ and $\mathcal{P}_{m, n}^{g, \kappa}$. When the distance is small, compressed forecasts of outcomes are close to the correct forecasts, which means that no statistical test can distinguish between the correct forecasts and the compressed forecasts. Thus outcome sequences that are generated by the n -player model validate the compressed perception of the game.

The following corollary states that player 0's compressed assessments of payoffs become accurate as n increases.

Corollary 2. *Let \bar{f} , κ and g be as above, then $|U_{m, \text{COMP}}^{g, \kappa} - U_{m, n}^{g, \bar{f}}| < C \cdot m \sqrt{\frac{\log n}{n}}$ for every m, n , where C is a constant that depends only on the parameters of the game skeleton.*

⁶Note however that $\mathcal{P}_{m, \text{COMP}}^{g, \kappa}$ is independent of g .

⁷The total variation distance between a pair μ, μ' of distribution over X is given by $\|\mu - \mu'\| = \sup_B |\mu(B) - \mu'(B)|$ where the supremum ranges over all Borel subsets B of X .

⁸The proof establishes the bound $C \leq 2|T| + 3|T| \cdot |A| \cdot L$. See also remark 6 for a somewhat sharper asymptotic

6. RELATING COMPRESSED EQUILIBRIA TO STANDARD CONCEPTS

Recall that for every strategy profile \bar{f} and any $\epsilon > 0$, a strategy g is an ϵ best response to \bar{f} , if $U_{m,n}^{g,\bar{f}} \leq U_{m,n}^{h,\bar{f}} + \epsilon$ for all strategies h . In a parallel manner, g is an ϵ compressed best response to $\kappa_{\bar{f}}$, if $U_{m,\text{COMP}}^{g,\kappa_{\bar{f}}} \geq U_{m,\text{COMP}}^{h,\kappa_{\bar{f}}} + \epsilon$.

The next corollary illustrates that a compressed best response in the m -times repeated game is essentially a real best response, and uniformly so for all the games with sufficiently many players.

Corollary 3. *Consider the m -times repeated game. For any $\epsilon > 0$ there is a positive integer N with the following properties:*

1. *if a strategy g is a compressed best response to a compressed strategy κ , then in all the games with more than N players g is a (real) ϵ best response to any profile \bar{f} with $\kappa_{\bar{f}} = \kappa$.*

Conversely, for such large n 's every strategy g which is (real) best response to a profile of strategies \bar{f} is an ϵ compressed best response to $\kappa_{\bar{f}}$.

Thus, for the purpose of computing essentially-optimal play, the player must only know that the number of players n is large, without having to know its value. Moreover, he only needs to know the compressed strategy, without having to know the individual strategies.

Corollary 4. *Let κ be a compressed equilibrium of a game $\Gamma^{m,\cdot}$, then for every game $\Gamma^{m,n}$, κ is a $\left(C \cdot m \sqrt{\frac{\log n}{n}}\right)$ -equilibrium, where C is a constant that depends only on the parameters of the game skeleton.*

Corollary 4 is a strong version of theorems that relate equilibria in games with continuum of players to approximate equilibria in games with finite number of players. There are several such theorems in the literature, mostly for one-shot games of complete information and pure strategies. See Carmona and Podczeck [10] for a recent result. The main concern of these papers is the justification of the equilibria in the game with continuum of players as limits of equilibria of the finite player games. In contrast, for us the compressed view is conceptually simpler since it stays entirely with games of finitely many players. Moreover, due to the uniform property of the best response strategies, the justification is stronger in the senses discussed above.

Closer in spirit to our motivation is a recent paper of Bodoh-Creed [8]. He analyzes equilibria in large interdependent values uniform price auction model where bidders have arbitrary preferences for multiple units using a nonatomic limit game.

Remark 3. The fact that κ is an ϵ -equilibrium implies that every player has a chance of at most $\sqrt{\epsilon}$ to reach a node in the game in which he doesn't play $\sqrt{\epsilon}$ -best response.

The notion of universal equilibrium, discussed below, is useful for games that are long and include many players. When one strategy is equilibrium for all such games, its use does not

require that the players have precise knowledge of the number of repetitions or the number of players. The theorem below is applicable to games in which the number of players is substantially larger than the number of repetitions.

Definition 10. Let κ be a repeated game strategy, and let $\widehat{\Gamma}$ be any subset of the games Γ^γ . We say that κ is a *universal (asymptotic) equilibrium of the family $\widehat{\Gamma}$* , if for every $\epsilon > 0$ there exists a natural number N such that for all the games $\Gamma^{m,n} \in \widehat{\Gamma}$ with $n > N$ and for every strategy g ,

$$(6.1) \quad U_{m,n}^{g,\kappa} \leq U_{m,n}^{\kappa,\kappa} + \epsilon$$

Theorem 2. [?] *Let $\widehat{\Gamma}$ be a family of games $\Gamma^{m,n}$ such that $m < \delta(n)$, where $\delta : \mathbb{N} \rightarrow \mathbb{N}$ is some function satisfying $\delta(n) \cdot \sqrt{\frac{\log n}{n}} \xrightarrow{n \rightarrow \infty} 0$. Let κ be any uniform compressed equilibrium as in Definition 7, Then κ is universal equilibrium of $\widehat{\Gamma}$.*

Since Proposition 2 establishes the existence of uniform compressed equilibrium, Theorem ?? is an existence theorem for universal equilibrium on the class of games it describes.

7. ILLUSTRATIVE EXAMPLE: REPEATED COMPUTER CHOICE WITH UNKNOWN FUNDAMENTALS

7.1. The game. Let $S = T = A = \{\mathcal{PC}, \mathcal{M}\}$ denote two possible *states* of nature, two possible *types* of players and two possible *actions* (selections) that a player may choose at every stage of a repeated computer-choice game. For example the state $s = \mathcal{M}$ means that computer \mathcal{M} has better features and thus it is more likely to be favored by the players. Player i is of type $t^i = \mathcal{M}$ means that she prefers computer \mathcal{M} , i.e., everything else being equal she receives a higher payoff when she chooses $a^i = \mathcal{M}$ than when she chooses $a^i = \mathcal{PC}$.

Initially, an unknown state s is chosen randomly with equal probabilities, $\theta(s = \mathcal{PC}) = \theta(s = \mathcal{M}) = 1/2$; and conditional on the realized state s the types of n fixed players are drawn by iid: $\tau_s(t^i = s) = 0.7$ and $\tau_s(t^i \neq s) = 0.3$. Each player is informed of her type t^i . Both s and the vector of t^i 's remain fixed throughout the repeated game.

Based on player i 's information at the beginning of each period k , she selects (with the possible use of a randomization device) one of the two computers, $a_k^i = \mathcal{PC}$ or $a_k^i = \mathcal{M}$. These simultaneous choices determine an empirical distribution $e_k[t, a]$, i.e., the proportion of players who are of type t and choose computer a in period k .

At the end of each period, a random sample (with replacement) of J players is selected, and the sample proportions $x_k = [x_k(\mathcal{PC}), x_k(\mathcal{M})]$ is publically announced. Player i 's payoff in period k is $u_k^i(t^i, a_k^i, x_k) = (x_k[a_k^i])^{1/3} + 0.2\delta_{a_k^i=t^i}$. In other words, her period payoff increases in a non-linear way with the announced proportion of players who match her choice plus an extra bonus of 0.2 if she chooses her favorite computer.

The game is repeated m times, and a player's overall payoff is the average of her m period payoffs.

7.2. Compressed computations. To illustrate compressed computations of a player in this example, consider the symmetric strategy profile in which every player repeatedly chooses her favorite computer in every period.

Updating beliefs about the state of nature:

Prior to the start of the game but after learning her type the player updates her .50-.50 prior probabilities to be $\theta_0(\mathcal{PC}) = .70$ and $\theta_0(\mathcal{M}) = .30$, if she is the \mathcal{PC} type; but $\theta_0(\mathcal{PC}) = .30$ and $\theta_0(\mathcal{M}) = .70$, if she is the \mathcal{M} type.

Proceeding inductively with $k = 0, 1, 2, \dots$, let $\theta_k(\mathcal{PC})$ and $\theta_k(\mathcal{M})$ denote period k th (random) prior probabilities of the two states and let $j_k(\mathcal{PC})$ denote the random number of the \mathcal{PC} users observed among the J users sampled in period k (i.e., $j_k(\mathcal{PC}) = x_k(\mathcal{PC}) \cdot J$). Then,

$$\theta_{k+1}(\mathcal{PC}) = \theta_k(\mathcal{PC})\beta(j_k(\mathcal{PC}); J, 0.7) / [\theta_k(\mathcal{PC})\beta(j_k(\mathcal{PC}); J, 0.7) + \theta_k(\mathcal{M})\beta(j_k(\mathcal{PC}); J, 0.3)]$$

and

$$\theta_{k+1}(\mathcal{M}) = 1 - \theta_{k+1}(\mathcal{PC}) .$$

In the expression above, β denotes Binomial coefficients: For example, $\beta(j_k(\mathcal{PC}); J, 0.7)$ is the probability of observing $j_k(\mathcal{PC})$ \mathcal{PC} users in the sample of J users when the probability of \mathcal{PC} use by a random player is 0.7.

Computing the period expected payoffs.

The compressed k th period expected payoffs are also relatively simple to compute. For example for a \mathcal{PC} - type player with prior probabilities θ_k they are:

$$\sum_{j=0, \dots, J} [\theta_k(\mathcal{PC})\beta(j; J, 0.7) + \theta_k(\mathcal{M})\beta(j; J, 0.3)](j/J)^{1/3} + 0.2 \text{ when she chooses } \mathcal{PC}, \text{ and}$$

$$\sum_{j=0, \dots, J} [\theta_k(\mathcal{PC})\beta(j; J, 0.3) + \theta_k(\mathcal{M})\beta(j; J, 0.7)](j/J)^{1/3} \text{ when she chooses } \mathcal{M}.$$

Remark 4. The computations above are significantly simpler than standard Bayesian computations due to the assumptions of the compressed view..

1. Assuming that the population proportions of realized types, conditional on the value of the unknown θ , are deterministic and known, (.70 to .30 or .30 to .70) means that the players only have to update their beliefs about the value of the unknown state θ . This is significantly simpler than full non-compressed computation in which in addition to the above the players would have to start with and update beliefs about the realized unknown population-type-distribution.

2. In addition to the gained ease in computation, not having to update beliefs about the population-type-distribution means that a player does not have to know the populations size, as would be the case under full non-compressed updating.

3. If we did not compress the population-type-distribution to the expected values of .7 and .3, even the computations in the four remaining expressions above would be significantly more

demanding: Each of the values of .7 and .3 in these expressions would have to be replaced by random variables that assume all the possible values of population proportions. This would make the expression quite complex and would again require knowledge of the population size.

4. The assumption that the player's own actions do not affect the probability of publically observed outcomes means that in updating her beliefs she may ignore her own past actions (which may be chosen randomly in more general cases). She also doesn't have to form beliefs about previous actions of her opponents since she compresses them to their expected values.

5. The observation just made means that under compressed view all the players face the same aggregate opponent strategy and thus they all use the same updating formulas. In particular, all players of the same type have the same prior probability distribution at the beginning of every period.

7.3. A compressed equilibrium with coordinated learning. To construct an equilibrium of the repeated game, define α to be the value of $\theta_k(\mathcal{PC})$ that equates the two period's expected payoff in the expressions above. In other words, facing opponents who choose their favorite computers, a \mathcal{PC} type player would prefer to choose \mathcal{PC} in periods with $\theta_k(\mathcal{PC}) > \alpha$ but would prefer to choose \mathcal{M} in periods with $\theta_k(\mathcal{PC}) < \alpha$. It is relatively easy to compute α , for example, if the sample size $J = 10$, $\alpha = .08$.

Notice that α must be smaller than .5, it is independent of k , and that in this symmetric example we would obtain the same value for α if we compute the same probability for switching from computer \mathcal{M} to \mathcal{PC} by a player of type \mathcal{M} . Since all the players of the same type carry with them the same posterior probabilities they would all prefer to switch to the opposite computer at the same period

We define the *compressed equilibrium with coordinated learning* by the following symmetric strategy: Choose your favorite computers initially, and continue to do so as long as the posterior on the state of your favorite computer $\theta_k(t^i) \geq \alpha$. If at some k $\theta_k(t^i) < \alpha$, then from this k on choose the other computer.

Remark 5. Using the properties of the compressed concepts, we observe the following:

(1) Due to the ease of computing α and updating the posterior beliefs, the equilibrium strategy is easy to play.

(2) Using the myopic characterization in Corollary 1, it is easy to see that the strategy profile above is a compressed equilibrium. All we have to show is that every player is compressed optimizing in each period, which is true by the definition of α .

(3) The equilibrium play will be of the following type: For some time k each player will be choosing her favorite computer up to time k , but they will all settle on one computer that will be chosen repeatedly from time k on.

(4) For sufficiently large number of players n , the above strategy will be an ϵ -Nash equilibrium of the repeated game.

Several questions may be of interest. How long will it take them to settle on one computer? Will they settle on the right computer (the one determined by the realized state)?, etc. The answer to such questions depend on the size of the sample J .

One may notice here that under the compressed approach these questions can be asked meaningfully. They could not be asked in the continuum model, because it does not provide the needed probability structure, and they would be difficult (impossible without knowledge of n) to compute in the asymptotic model.

8. COUNTER-EXAMPLES

The examples below illustrate the importance of two issues that arise in our approximation results: (1) The need to bound the number of game repetitions m relative to the number of players n (or the need to have larger number of players for longer games). (2) The need to smooth the outcome functions using noise, such as the sampling in the illustrative example above.

Both issues can be illustrated even for games of complete information. But to keep the presentation short here, we stay with our illustrative example above. The first example shows the role of the bound on the length of the game, m in Theorem 1

Example 2. Fix a small positive number ϵ . For every $i \geq 1$ let $m_i \in \mathbb{N}$ be sufficiently large so that when tossing a fair coin m_i times, the probability that the proportion of successes being "short by at least $1/i$ " is at most $\frac{\epsilon}{2^i}$, i.e., $\mathcal{P}[\text{number of successes} < (1/2 - 1/i)m_i] < \frac{\epsilon}{2^i}$. Now define inductively, for $i = 1, 2, \dots$, the corresponding blocks of periods by: $M_0 = 0$, $M_i = M_{i-1} + m_i$; and $B_i = \{M_{i-1} + 1, M_{i-1} + 2, \dots, M_{i-1} + m_i\}$. In other words, we partition the periods of play into successive blocks B_1, B_2, \dots , with each B_i consisting of m_i periods that start after period M_{i-1} and end in period M_i .

In the illustrative example, assume that the sample size $J = 2$ and define a period to be *balanced* if its sampled outcome has one \mathcal{PC} user and one \mathcal{M} user. Consider the following strategy: At every stage k the player plays $(1/2, 1/2)$ unless there exists some block of periods B_i ending at period $M_i < k$ in which the proportion of balanced periods is "short by at least $1/i$," i.e.,

$$(8.1) \quad \#\{j \in B_i : j \text{ is balanced}\} < (1/2 - 1/i)m_i.$$

If there exists a block B_i in which (8.1) holds then the player plays \mathcal{PC} .

In words, the player considers the proportion of balanced periods in each block B_i of m_i periods. He continues to play $(1/2, 1/2)$ as long as the proportion of balanced periods in

each block B_i is not short by more than $1/i$. If at some block B_i of m_i periods the proportion of balanced periods is short by more than $1/i$, the player changes to playing \mathcal{PC} forever.

Consider now the symmetric profile in which all the players follow the strategy above. According to the compressed view, when the players randomize $(1/2, 1/2)$ the probability that a period is balanced is $1/2$. Thus, given the choice of the numbers m_i , under the compressed view the $\mathcal{P}_{comp}[(8.1) \text{ happens for some } i] < \epsilon$.

But in the real game with n players, if the players randomize $(1/2, 1/2)$ the probability that a period is balanced is $1/2 - 1/n$, since there is a probability of $1/n$ that the same player is chosen twice in the two-player sample (chosen with replacement). Thus for large i s the probability that (8.1) happens is close to one and the real $\mathcal{P}[(8.1) \text{ happens for some } i] = 1$.

As seen above the event that from some point on the players play \mathcal{PC} forever has compressed probability of at most ϵ but has real probability 1. So in this example large number of periods leads to severe discrepancies between the compressed and real probabilities.

The next example shows that Theorem 1 fails without the assumption of norm continuity of the outcome generating function.

Example 3. Modify the illustrative example to have the set of possible outcomes $X = [0, 1]$, and assume that the proportion of players who choose \mathcal{M} , $e_k(\mathcal{M}) \in [0, 1]$, is reported precisely (with no sampling or any other noise) at the end of each period k . This corresponds to the outcome generating function $\chi(e) = \delta_e$ that puts probability one on the true empirical distribution, a function which is not continuous in the total variation distance on $\Delta(X)$.

Consider the following play: in each period k the players randomize $(1/2, 1/2)$ provided that the reported proportions of \mathcal{M} users in all previous periods are precisely $1/2$. But following histories that include a reported proportion different from $1/2$ they choose \mathcal{PC} .

The event that the players continue to randomize $(1/2, 1/2)$ forever has compressed probability one. But the real probability that the actual (and reported) proportion of \mathcal{M} users in period 0 is precisely $1/2$ is close to zero (especially for large n 's). Thus the event that they choose \mathcal{PC} from period two on has compressed probability zero and real probability close to one.

9. EXTENSIONS AND VARIATIONS

9.1. Extensions. Some restrictions in this paper were done to keep the simplicity of presentation and notation. Three of these restrictions may be relaxed as follows.

9.1.1. Asymmetric games, allowing for player roles . Symmetry of players may be relaxed by letting players be of different roles: male and females, sellers and buyers, members of different biological species, etc.

This can be done by expanding the notion of game skeleton (Definition 1) to include a finite set of player *roles* R and a distribution $\rho \in \Delta(R)$, which represents the proportion of players in each role. The distribution must consist of rational numbers and be combined with restrictions on the possible number of players n that guarantee integer number of players in each role. For example, we may consider games with even number of players $n = 2h$ with $\rho[m] = \rho[f] = 1/2$, so that we have h *males* and h *females*.

In addition to type-independence conditional on the state of nature, assumed in the current model, player types would be independent, conditional on the state of nature *and* their roles. The rest of the definitions and results of this paper carry through.

9.1.2. *Discounting vs finite-long games.* Much of the theory presented here may be extended from finitely repeated games to infinitely repeated games with discounting, in which statements related to the length of the game are replaced by statement about the future discount factor.

For example, Corollary 4 has a natural analogue for discounted games: A compressed equilibrium is an approximate equilibrium when the number of players is large and the discount factor is bounded away from 1.

9.1.3. *Weighted players.* The symmetry of players' impact on the period outcomes in the games $\Gamma^{m,n}$ can be relaxed, as long as the impact of the players diminished as n increases. This may be relevant in biological games, congestion games, market games and other applications where players's size needs not be the same.

Fix a skeleton game Γ . For every m, n and every distribution of *player weights* $w \in \Delta(\{0, 1, \dots, n-1\})$, let $\Gamma^{m,w}$ be similar to the game $\Gamma^{m,n}$ defined in Section 1, but subject to the following modifications: (1) the average used in computing the empirical distribution e_k of players types and actions should be weighted by w , and (2) the average used to compute the compressed strategy should also be weighted by w . Under these modification similar results to ours holds if the $\|\cdot\|_\infty$ -norm of w is sufficiently small.

9.2. **Additional directions of research.** The mathematics involved in the main results reported in this paper require specific assumptions, such as noise in the period reported outcomes and limitations on the length of the game. From a technical point of view, further investigation of these and alternative models may be important.

9.2.1. *Alternatives to the noise model.* As Example 3 shows, Theorem 1 relies on the introduction of noise. The implication of noise is that even if the players' behavior is not continuous in the actual history, it is continuous in the "noisy history."

An analogue to theorem 1 without noise can be proved if we assume that the players strategies are continuous in the history of actual empirical choices. However, additional

restrictions on the game should be satisfied so that compressed equilibrium in continuous strategies exist.

9.2.2. Removing the limitations on m . As Example 2 shows, the number of players required in Theorem 1 grows as the number of periods grow. This is the case because every additional period may add probability of failure of the law of large number in the empirical distribution of actions. When players behavior is assumed to depend on the entire past history, a failure of the compressed view in one period may produce failure in all subsequent periods.

But the coordinated learning equilibrium in our illustrative example, was not subject to this difficulty. It provided a good approximation of the probabilities of play regardless of the number of repetitions. This uniform good approximation is due to two features of the equilibrium: (1) the infinite future play depends only on the outcomes of an initial segment of period, and (2) no new learning takes place beyond this initial segment.

One may try to construct a general class of equilibria that exhibit such features. Specifically, restrict attentions to strategies that depend only on a fixed number of initial periods, say periods $0, \dots, K$, then construct an analogue of Theorem 1 in which for large enough n 's (that depend on K and ϵ , and not on the total number of game repetitions m) the future play depends only on the outcomes of the first K periods.

It may be required, however, that all the relevant learning that takes place during the play of the game should happen with high probability in the first K periods.

10. PROOFS

10.1. Preliminaries.

10.1.1. Notations. For every Borel space Z we denote by $\Delta(Z)$ the space of probability distributions over Z . We let $\|\mu - \mu'\|$ denote the total variation distance between the distributions $\mu, \mu' \in \Delta(Z)$. If Z is finite then we view $\Delta(Z)$ as a subset of \mathbb{R}^Z , in which case $\|\mu - \mu'\| = \|\mu - \mu'\|_1/2$ where $\|\cdot\|_1$ denotes the L^1 norm. In the case of finite Z we also denote by $\|\mu - \mu'\|_\infty$ the L^∞ distance given by $\|\mu - \mu'\|_\infty = \max_{z \in Z} |\mu[z] - \mu'[z]|$ and we denote by $[\mu]$ the support of μ .

10.1.2. Concentration of Empirical Distribution. We use the following propositions.

Proposition 3. *Let Z be a finite set. Then for every Z -valued independent random variables $\mathbf{Z}_1, \dots, \mathbf{Z}_n$*

$$\mathbb{P}(\|\mathbf{m} - \mathbb{E}\mathbf{m}\|_\infty > \eta) < 2|Z| \cdot e^{-2\eta^2 n},$$

where $\mathbf{m} \in \Delta(Z)$ is the random empirical distribution of $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ given by $\mathbf{m}[z] = \frac{1}{n} \#\{i | Z_i = z\}$.

Proof. It follows from Chernoff Bound [3, Corollary A.1.7] that

$$\mathbb{P}(\|\mathbf{m}[z] - \mathbb{E}\mathbf{m}[z]\| > \eta) < 2e^{-2\eta^2 n}$$

for every $z \in Z$. Therefore

$$\mathbb{P}(\|\mathbf{m} - \mathbb{E}\mathbf{m}\|_\infty > \eta) \leq \sum_{z \in Z} \mathbb{P}(\|\mathbf{m}[z] - \mathbb{E}\mathbf{m}[z]\| > \eta) \leq 2|Z|e^{-2\eta^2 n}.$$

□

Proposition 4. *Let Z be a finite set. Then for every Z -valued independent random variables $\mathbf{Z}_1, \dots, \mathbf{Z}_n$*

$$\mathbb{E}\|\mathbf{m} - \mathbb{E}\mathbf{m}\|_1 < |Z|/\sqrt{n},$$

where $\mathbf{m} \in \Delta(Z)$ is the random empirical distribution of $\mathbf{Z}_1, \dots, \mathbf{Z}_n$ given by $\mathbf{m}[z] = \frac{1}{n} \#\{i | Z_i = z\}$.

Proof. For every $z \in Z$ the random variable $\mathbf{m}[z]$ is the average of n independent variables with values in $\{0, 1\}$. Therefore

$$\mathbb{E}|\mathbf{m}[z] - \mathbb{E}\mathbf{m}[z]| \leq \sqrt{\text{Var}(\mathbf{m}[z])} < 1/\sqrt{n}.$$

It follows that

$$\mathbb{E}\|\mathbf{m} - \mathbb{E}\mathbf{m}\|_1 < \sum_{z \in Z} \mathbb{E}|\mathbf{m}[z] - \mathbb{E}\mathbf{m}[z]| < |Z|/\sqrt{n}.$$

□

10.1.3. *Coupling.* Let X be a Borel space and μ and ν two probability distributions over X . A *coupling* of μ and ν is a pair (\mathbf{X}, \mathbf{Y}) such that \mathbf{X} is distributed μ and \mathbf{Y} is distributed ν . The following lemma shows how coupling of random variables is useful in order to study distance between distributions.

Proposition 5. [25, Theorem 5.2] *Let X be a Borel space and $\mu, \nu \in \Delta(X)$. Then*

- (1) *If (\mathbf{X}, \mathbf{Y}) is a coupling of (μ, ν) then $\mathbb{P}(\mathbf{X} \neq \mathbf{Y}) \geq \|\mu - \nu\|$.*
- (2) *There exists an optimal coupling (\mathbf{X}, \mathbf{Y}) of (μ, ν) such that $\mathbb{P}(\mathbf{X} \neq \mathbf{Y}) = \|\mu - \nu\|$.*

10.2. **Proof of Proposition 2.** A *one-shot strategy* in Γ is given by a function $f : T \rightarrow \Delta(A)$. In the one-shot compressed setup, if players play according to f then the state of nature \mathbf{S} , a player 0's type \mathbf{T} , and the outcome \mathbf{Y} are random variables with joint distribution

$$(10.1) \quad \mathbb{P}(\mathbf{S} = s, \mathbf{T} = t, \mathbf{Y} \in B) = \theta(s) \cdot \tau_s[t] \cdot \chi_{d_f(s)}(B),$$

where $d_f(s) \in \Delta(T \times A)$, the *empirical distribution of players type-action under f when the state of nature is s* , is given by

$$(10.2) \quad d_f(s)[t, a] = \tau_s[t] \cdot f(t)[a].$$

We call a triple $(\mathbf{S}, \mathbf{T}, \mathbf{X})$ with distribution (10.1) a *random one-shot f -play*.

A one-shot strategy f is a *compressed one-shot equilibrium* if

$$[f(t)] \subseteq \operatorname{argmax}_{a \in A} \mathbb{E}(u(t, a, \mathbf{X} | \mathbf{T} = t))$$

for every type t , where $\mathbf{S}, \mathbf{T}, \mathbf{X}$ is a random one-shot f -play.

Lemma 1. *Every one-shot game admits a compressed one-shot equilibrium*

Proof. Since T is finite, the set of one-shot strategies is compact and convex and the best response correspondence has closed graph and convex values. Therefore it has a fixed point by Kakutani's fixed point theorem. \square

Proof of Proposition 2. We construct by induction the compressed strategy κ and the compressed κ -play process $\mathbf{S}, \mathbf{X}_0, \mathbf{X}_1, \dots$ so that the condition of Corollary 1 is satisfied, as follows: Let \mathbf{S} be S -valued random variable with distribution θ . For every k let $\kappa(t, x_0, \dots, x_{k-1})$ be such that $f(t) = \kappa(t, x_0, \dots, x_{k-1})$ is a one-shot equilibrium in the game with prior $\mathbb{P}(\mathbf{S} = \cdot | \mathbf{X}_0 = x_0, \dots, \mathbf{X}_{k-1} = x_{k-1})$ (The function κ can be selected to be measurable by von-Neumann's Selection Theorem) and let \mathbf{X}_k be a random variable such that

$$\mathbb{P}(\mathbf{X}_k \in \cdot | \mathbf{S} = s, \mathbf{X}_0 = x_0, \dots, \mathbf{X}_{k-1} = x_{k-1}) = \chi_{d_f(s)},$$

where $d_f(s) \in \Delta(T \times A)$ is given by (10.2). \square

10.3. Proof of Theorem 1. We prove the theorem using a lemma that couples a play in the actual game with a compressed play. In order to distinguish between corresponding entities in the actual play and in the compressed play, we denote the random variables that represent player 0 actions in the compressed play by $\mathbf{B}_0^0, \mathbf{B}_1^0, \dots$ and the random variables that represent the outcomes in the compressed play by $\mathbf{Y}_0, \mathbf{Y}_1, \dots$.

Lemma 2. *Let Γ be a game skeleton. For every strategy g and a reactive strategy profile $\bar{f} = (f_1, \dots, f_n)$ of Γ with compression $\kappa = \kappa_{\bar{f}}$ there exist random variables $\mathbf{S}, \mathbf{T}^i, \mathbf{A}_k^i, \mathbf{X}_k, \mathbf{Y}_k$ for $i \in N$ and $k = 0, 1, \dots$ such that*

- $(\mathbf{S}, \mathbf{T}^0, \mathbf{B}_0^0, \mathbf{Y}_0, \mathbf{B}_1^0, \mathbf{Y}_1, \dots)$ is a compressed random (g, κ) -play.
- $(\mathbf{S}, \mathbf{T}^i, \mathbf{A}_k^i, \mathbf{X}_0, \mathbf{X}_1, \dots)$ is a random (g, \bar{f}) -play of $\Gamma^{\cdot n}$.
- For every m it holds that

$$(10.3) \quad \mathbb{P}(\mathbf{A}_0^0 = \mathbf{B}_0^0, \mathbf{X}_0 = \mathbf{Y}_0, \dots, \mathbf{A}_m^0 = \mathbf{B}_m^0, \mathbf{X}_m = \mathbf{Y}_m) >$$

$$1 - (2|T| + 3|T| \cdot L \cdot |A|) \cdot m \sqrt{\frac{\log n}{n}}.$$

Proof. We couple a compressed (g, κ) -play $(\mathbf{S}, \mathbf{T}^0, \mathbf{B}_k^0, \mathbf{Y}_k)$ and a (g, \bar{f}) -play $(\mathbf{S}, \mathbf{T}^i, \mathbf{A}_k^i, \mathbf{X}_k)$ as follows:

- The state of nature \mathbf{S} is distributed according to θ .
- Conditioned on \mathbf{S} , the players' types \mathbf{T}^i for $i \in N$ are i.i.d $\tau_{\mathbf{S}}$.
- Conditioned on the history of stages $0, \dots, k-1$, the players in the actual play choose their day k actions independently, where player 0 mixes according to g and opponent $i \in N \setminus \{0\}$ mixes according f_i . Player 0's action B_k^0 in the compressed play is optimally coupled with his action A_k^0 in the actual play according to Proposition 5.
- The outcomes \mathbf{X}_k and \mathbf{Y}_k of day k in the actual and compressed plays are randomized according to the corresponding empirical distribution of day k types and actions. Moreover, these outcomes are optimally coupled according to Proposition 5.

In equations,

$$\mathbb{P}(\mathbf{S} = s, \mathbf{T}^i = t^i \ i \in N) = \theta[s] \cdot \prod_i \tau_s[t^i]$$

$$\mathbb{P}(\mathbf{A}_k^i = a^i \ i \in N \mid \mathcal{F}_k) = g(\mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1})[a^0] \cdot \prod_{i \in N \setminus \{0\}} f^i(\mathbf{T}^i, \mathbf{X}_0, \dots, \mathbf{X}_{k-1})[a^i]$$

$$\mathbb{P}(\mathbf{B}_k^0 = a \mid \mathcal{F}_k) = g(\mathbf{T}^i, \mathbf{B}_0^0, \mathbf{Y}_0, \dots, \mathbf{B}_{k-1}^0, \mathbf{Y}_{k-1})[a]$$

$$\mathbb{P}(\mathbf{B}_k^0 \neq \mathbf{A}_k^0 \mid \mathcal{F}_k) = \|g(\mathbf{T}^i, \mathbf{B}_0^0, \mathbf{Y}_0, \dots, \mathbf{B}_{k-1}^0, \mathbf{Y}_{k-1}) - g(\mathbf{T}^0, \mathbf{A}_0^0, \mathbf{X}_0, \dots, \mathbf{A}_{k-1}^0, \mathbf{X}_{k-1})\|$$

$$\mathbb{P}(\mathbf{X}_k \in \cdot \mid \tilde{\mathcal{F}}_k) = \chi_{\mathbf{S}, \mathbf{d}_k}$$

$$\mathbb{P}(\mathbf{Y}_k \in \cdot \mid \tilde{\mathcal{F}}_k) = \chi_{\mathbf{S}, \mathbf{e}_k}$$

$$\mathbb{P}(\mathbf{X}_k \neq \mathbf{Y}_k \mid \tilde{\mathcal{F}}_k) = \|\chi_{\mathbf{S}, \mathbf{d}_k} - \chi_{\mathbf{S}, \mathbf{e}_k}\|$$

where \mathcal{F}_k is the *sigma*-algebra generated by $\mathbf{S}, \mathbf{T}^i, \mathbf{A}_l^i, \mathbf{B}_l^0, \mathbf{Y}_l, \mathbf{X}_l$ for $i \in N$ and $l < k$, which represents the k -th day history, $\tilde{\mathcal{F}}_k$ is the σ -algebra generated by \mathcal{F}_k and the k -th day actions \mathbf{A}_k^i and \mathbf{B}_k^0 , and \mathbf{d}_k and \mathbf{e}_k are the random empirical distributions of day k type and actions in the actual and finite plays, given by

$$\mathbf{d}_k[t, a] = \#\{i \in N \mid \mathbf{T}^i = t, \mathbf{A}_k^i = a\}/n, \text{ and}$$

$$\mathbf{e}_k[t, a] = \tau_{\mathbf{S}}[t] \kappa(t, \mathbf{Y}_0, \dots, \mathbf{Y}_{k-1})[a]$$

for every $(t, a) \in T \times A$.

The construction renders $(\mathbf{S}, \mathbf{T}^0, \mathbf{B}_k^0, \mathbf{Y}_k)$ a compressed (g, κ) - play and $(\mathbf{S}, \mathbf{T}^i, \mathbf{A}_k^i, \mathbf{X}_k)$ a (g, f) - play in $\Gamma^{,n}$.

Let G_0 be the event that the empirical distribution of types is η -concentrated around its expectation, i.e.,

$$G_0 = \{\|\mathbf{m} - \tau_{\mathbf{S}}\|_{\infty} < \eta\},$$

where $\mathbf{m} \in \Delta(T)$ is the random empirical distribution of $\mathbf{T}_1, \dots, \mathbf{T}_n$ given by

$$(10.4) \quad \mathbf{m}[t] = \frac{1}{n} \#\{i \mid T^i = t\}.$$

and η is TBD. Since $\mathbf{T}_1, \dots, \mathbf{T}_n$ are independent conditioned on \mathbf{S} it follows from Proposition 3 that

$$(10.5) \quad \mathbb{P}(G_0) > 1 - 2|T|e^{-2\eta^2 n}.$$

For every $k = 0, 1, 2, \dots$ consider the event

$$G_k = G_0 \cap \{\mathbf{Y}_0 = \mathbf{X}_0, \dots, \mathbf{Y}_{k-1} = \mathbf{X}_{k-1}\}.$$

Note that from the coupling condition of \mathbf{B}_k^0 and \mathbf{A}_k^0 it follows by induction that $\mathbf{B}_k^0 = \mathbf{A}_k^0$ on G_k and therefore

$$G_m \subseteq \{\mathbf{A}_0^0 = \mathbf{B}_0^0, \mathbf{X}_0 = \mathbf{Y}_0, \dots, \mathbf{A}_m^0 = \mathbf{B}_m^0, \mathbf{X}_m = \mathbf{Y}_m\}.$$

Thus, to prove the lemma it is sufficient to establish the bound from below on $\mathbb{P}(G_m)$. To do that we establish a bound on $\mathbb{P}(G_{k+1} | G_k)$ for every k .

Conditioned on \mathcal{F}_k the type of opponent $i \in N \setminus \{0\}$ is \mathbf{T}^i and his day k action is randomized according to $f^i(\mathbf{T}^i, \mathbf{Y}_0, \dots, \mathbf{Y}_{k-1})$. The empirical distribution \mathbf{d}_k is induced by the type-action of all the opponents and the contribution of player 0. Therefore

$$(10.6) \quad \left| \mathbb{E}(\mathbf{d}_k[t, a] | \mathcal{F}_k) - \mathbf{m}[t] \cdot \kappa(t, \mathbf{Y}_0, \dots, \mathbf{Y}_{k-1})[a] \right| \leq 1/n$$

for every $(t, a) \in T \times A$ where $\mathbf{m} \in \Delta(T)$ is the random empirical distribution of types given by (10.4). It follows that on G_k

$$(10.7) \quad \begin{aligned} & \left\| \mathbb{E}(\mathbf{d}_k | \mathcal{F}_k) - \mathbf{e}_k \right\|_1 = \sum_{t,a} \left| \mathbb{E}(\mathbf{d}_k[t, a] | \mathcal{F}_k) - \mathbf{e}_k[t, a] \right| \\ & \leq \sum_{t,a} \left(\left| \mathbf{m}[t] \cdot \kappa(t, \mathbf{Y}_0, \dots, \mathbf{Y}_{k-1})[a] - \tau_{\mathbf{S}}[t] \cdot \kappa(t, \mathbf{Y}_0, \dots, \mathbf{Y}_{k-1})[a] \right| + 1/n \right) \\ & \leq \left\| \mathbf{m} - \tau_{\mathbf{S}} \right\|_1 + |T| \cdot |A| \cdot 1/n < |T| \cdot \eta + |T| \cdot |A|/n, \end{aligned}$$

where the first inequality follows from (10.6) and (3.2) and the last inequality from the fact that $\left\| \mathbf{m} - \tau_{\mathbf{S}} \right\|_1 \leq |T| \cdot \left\| \mathbf{m} - \tau_{\mathbf{S}} \right\|_\infty < \eta$ on G_k since $G_k \subseteq G_0$.

Since the players type-actions are independent conditioned on \mathcal{F}_k it follows from Proposition 4 that

$$\mathbb{E} \left(\left\| \mathbf{d}_k - \mathbb{E}(\mathbf{d}_k | \mathcal{F}_k) \right\|_1 \middle| \mathcal{F}_k \right) < |A| \cdot |T| / \sqrt{n}.$$

From the last inequality and (10.7) it follows that

$$\mathbb{E} \left(\left\| \mathbf{d}_k - \mathbf{e}_k \right\|_1 \middle| \mathcal{F}_k \right) < |A| \cdot |T| / \sqrt{n} + |T| \cdot \eta + |T| \cdot |A|/n$$

On G_k . From the coupling condition of \mathbf{X}_k and \mathbf{Y}_k , the Lipschitz condition on χ and the last inequality it follows that on G_k

$$\begin{aligned} \mathbb{P}\left(\mathbf{Y}_k \neq \mathbf{X}_k \mid \mathcal{F}_k\right) &= \mathbb{E}\left(\mathbb{P}\left(\mathbf{Y}_k \neq \mathbf{X}_k \mid \tilde{\mathcal{F}}_k\right) \mid \mathcal{F}_k\right) = \\ \mathbb{E}\left(\|\chi_{\mathbf{s}, \mathbf{e}_k} - \chi_{\mathbf{s}, \mathbf{d}_k}\| \mid \mathcal{F}_k\right) &\leq L \cdot \mathbb{E}\left(\|\mathbf{d}_k - \mathbf{e}_k\|_1 \mid \mathcal{F}_k\right) \leq L \cdot |A| \cdot |T| / \sqrt{n} + L|T|\eta + L|T| \cdot |A|/n. \end{aligned}$$

Since G_k is \mathcal{F}_k -measurable it follows from the last inequality that

$$\mathbb{P}\left(G_{k+1} \mid G_k\right) = 1 - \mathbb{P}\left(\mathbf{Y}_k \neq \mathbf{X}_k \mid G_k\right) \geq 1 - L \cdot |A| \cdot |T| / \sqrt{n} + L|T|\eta + L|T| \cdot |A|/n.$$

From the last inequality and (10.5) it follows that

$$(10.8) \quad \mathbb{P}(G_m) \geq 1 - 2|T|e^{-2\eta^2 n} - m \cdot L \cdot |T| \cdot (|A|/\sqrt{n} + \eta + |A|/n).$$

The last inequality holds for every η . Choosing $\eta = \sqrt{\frac{\log n}{n}}$ we get that

$$\begin{aligned} \mathbb{P}(G_m) &\geq 1 - 2|T|/n^2 - m \cdot L \cdot |T| \cdot \left(|A|/\sqrt{n} + \sqrt{\frac{\log n}{n}} + 1/n\right) \\ &> 1 - (2|T| + 3|T| \cdot L \cdot |A|)m \sqrt{\frac{\log n}{n}}, \end{aligned}$$

as desired □

Remark 6. By choosing $\eta = C \cdot \sqrt{n}$ in (10.8) we get that we can replace the conclusion of Lemma 2 with the assertion that for every $\epsilon > 0$

$$\mathbb{P}\left(\mathbf{A}_0^0 = \mathbf{B}_0^0, \mathbf{X}_0 = \mathbf{Y}_0, \dots, \mathbf{A}_m^0 = \mathbf{B}_m^0, \mathbf{X}_m = \mathbf{Y}_m\right) > \epsilon + Cm/\sqrt{n},$$

where C is a constant that depends on ϵ and the parameters of the game.

Proof of Theorem 1. Let $C = 2|T| + 3|T| \cdot L \cdot |A|$. Lemma 2 establishes a coupling of the compressed random outcomes $(\mathbf{Y}_0, \dots, \mathbf{Y}_{m-1})$ and the random outcomes $(\mathbf{X}_0, \dots, \mathbf{X}_{m-1})$ in Γ^n such that

$$\mathbb{P}(\mathbf{X}_0 = \mathbf{Y}_0, \dots, \mathbf{X}_{k-1} = \mathbf{Y}_{k-1}) > 1 - C \cdot m \sqrt{\frac{\log n}{n}}.$$

Therefore it follows from Proposition 5 that $\|\mathcal{P}_{m, \mathcal{CO}MP}^{g, F} - \mathcal{P}_{m, n}^{g, F}\| < C \cdot m \sqrt{\frac{\log n}{n}}$ as desired. □

REFERENCES

- [1] Sachin Adlakha and Johari Ramesh. Mean field equilibrium in dynamic games with strategic complementarities. (2010).
- [2] Nabil I. Al-Najjar and Rann Smorodinsky. Pivotal players and the characterization of influence. *J. Econom. Theory*, 92(2):318–342, 2000.

- [3] Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., Hoboken, NJ, third edition, 2008. With an appendix on the life and work of Paul Erdős.
- [4] R. J. Aumann and L. S. Shapley. *Values of non-atomic games*. Princeton University Press, Princeton, N.J., 1974. A Rand Corporation Research Study.
- [5] R.J. Aumann. *Repeated games with incomplete information*. MIT press, 1995.
- [6] E. Azevedo and E. Budish. Strategyproofness in the large. Technical report, Working paper, 2012.
- [7] Yaron Azrieli and Eran Shmaya. Lipschitz games. *Mathematics of operation research*, forthcoming.
- [8] A. Bodoh-Creed. Approximation of large games with applications to uniform price auctions. *Auctions, Market Mechanisms, and Their Applications*, pages 54–54, 2012.
- [9] Guilherme Carmona. On the purification of Nash equilibria of large games. *Econom. Lett.*, 85(2):215–219, 2004.
- [10] Guilherme Carmona and Konrad Podczeck. Ex-post stability of Bayes-Nash equilibria of large games. *Games and Economic Behavior*, 74(1):418–430, 2012.
- [11] Edward Cartwright and Myrna Wooders. On equilibrium in pure strategies in games with many players. *Internat. J. Game Theory*, 38(1):137–153, 2009.
- [12] P. Dubey, A. Mas-Colell, and M. Shubik. Efficiency properties of strategies market games: An axiomatic approach. *Journal of Economic Theory*, 22(2):339–362, 1980.
- [13] Duggan, John. Noisy stochastic games. *Econometrica*, 80:2017–2045, (2012).
- [14] T. Feddersen and W. Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica: Journal of the Econometric Society*, pages 1029–1058, 1997.
- [15] R. Gradwohl and O. Reingold. Fault tolerance in large games. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 274–283. ACM, 2008.
- [16] R. Gradwohl, O. Reingold, A. Yadin, and A. Yehudayoff. Players’ effects under limited independence. *Mathematics of Operations Research*, 34(4):971–980, 2009.
- [17] Ronen Gradwohl and Omer Reingold. Partial exposure in large games. *Games and Economic Behavior*, 68(2):602–613, 2010.
- [18] Edward J. Green. Non-cooperative price taking in large dynamic markets. preprint.
- [19] Boyan Jovanovic and Robert W. Rosenthal. Anonymous sequential games *Journal of Mathematical Economics*, pages 77–87, 1988.
- [20] E. Kalai and W. Stanford. Finite rationality and interpersonal complexity in repeated games. *Econometrica: Journal of the Econometric Society*, pages 397–410, 1988.
- [21] Ehud Kalai. Large robust games. *Econometrica*, 72(6):1631–1665, 2004.
- [22] Ehud Kalai and Eran Shmaya. Large repeated games with incomplete information. preprint.
- [23] M. Ali Khan and Yeneng Sun. Non-cooperative games with many players. In Robert J. Aumann and Sergiu Hart, editors, *Handbook of game theory with economic applications*, volume 3, chapter 46, pages 1761–1808. Elsevier, Amsterdam, The Netherlands, 2002.
- [24] Pierre-Louis Lions. Notes on Mean Field Games, 2012.
- [25] Torgny Lindvall. *Lectures on the Coupling Method*. Wiley, New York, 1992.
- [26] G.J. Mailath and A. Postlewaite. Asymmetric information bargaining problems with many agents. *The Review of Economic Studies*, 57(3):351–367, 1990.
- [27] Andreu Mas-Colell. On a theorem of Schmeidler. *J. Math. Econom.*, 13(3):201–206, 1984.
- [28] Roger B. Myerson. Population uncertainty and Poisson games. *Internat. J. Game Theory*, 27(3):375–392, 1998.

- [29] Roger B. Myerson. Large Poisson games. *J. Econom. Theory*, 94(1):7–45, 2000.
- [30] William Novshek and Hugo Sonnenschein. Cournot and Walras equilibrium. *J. Econom. Theory*, 19(2):223–266, 1978.
- [31] Salim Rashid. Equilibrium points of nonatomic games: Asymptotic results. *Econom. Lett.*, 12(1):7–10, 1983.
- [32] A. Rustichini, M.A. Satterthwaite, and S.R. Williams. Convergence to efficiency in a simple market with incomplete information. *Econometrica: Journal of the Econometric Society*, pages 1041–1063, 1994.
- [33] Hamid Sabourian. Anonymous repeated games with a large number of players and random outcomes. *J. Econom. Theory*, 51(1):92–110, 1990.
- [34] David Schmeidler. Equilibrium points of nonatomic games. *J. Statist. Phys.*, 7:295–300, 1973.

KELLOGG SCHOOL OF MANAGEMENT, NORTHWESTERN UNIVERSITY, 2001 SHERIDAN ROAD, EVANSTON
IL 60208 AND MICROSOFT RESEARCH NEW ENGLAND

E-mail address: kalai@kellogg.northwestern.edu