LEARNING AND STABILITY IN BIG UNCERTAIN GAMES

EHUD KALAI AND ERAN SHMAYA

Abstract. A big game is repeatedly played by a large population of anonymous players. The game fundamentals and the set of players change with time and there is incomplete and imperfect information about the changes and the play.

Big games admit natural myopic Markov perfect equilibria in which the proportion of unpredictable and potentially chaotic periods is limited.

Examples include social adaptation of technology and price stability in production games.

Part 1. Overview of big games

Strategic interaction in large populations is a subject of interest in economics, political science, computer science, biology and more. Indeed, the last three decades produced substantial literature that deals with large strategic games in a variety of applications such as markets [27], bargaining [22], auctions [28, 29], voting [10, 26], electronic commerce [13], market design [4, 7] and many more.

The current paper studies *big* games that are played *repeatedly* by a *large* population of annonymous players. Uncertainty about such a game is due to ongoing changes in the game fundamentals and the set of players, with incomplete and imperfect information about the play and these changes. Big uncertain games are common, yet difficult to analyse by standard Bayesian methods. However, they are amenable to Baysian analysis under behavioral assumptions, common in economics and other fields. Moreover, with the exception of a limited proportion of *learning periods*, equilibrium play of big games is *predictable* and *stable*.

Date: September 14, 2015.
Key words and phrases. Anonymous games, Nash equilibrium, Repeated games, Large games, Bayesian equilibrium, Markov equilibrium, Price taking, Rational expectations.

Examples of big games with changing fundamentals are the following: In the game of *butter consumption*, fundamental chages occured with the 'introduction of margarine,' government announcements that 'butter is unhealthy' and then that 'margarine is unhealthy,' the introduction of 'low fat' butter, etc. In *rush-hour commute games* fundamental changes occure when a 'new highway' is introduced, 'new trainlines' are introduced, 'road changes' are made, etc. In the market of *computing devices* fundamental changes included the introduction of 'mainframe computers,' 'personal computers,' 'laptop computers,' 'smart phones' and now 'smart watches.'

Predictability and stability are essential properties for the well functioning of big games. A period of play is predictable, if there is high probability that all the players' predictions about the outcome of the period, predictions made before the period started, will be highly accurate. In the rush-hour example, this may mean that the players' predictions of a day's driving times on the various roads, predictions made before they start driving, are likely to be within a few minutes of the actual realized driving times. Stability of a period is meant in a *hindsight*, *no regret*, or *ex-post* sense: There is only a small probability that once they start driving and become informed of the realized driving times of the day, as may be the case from a radio report, some player may recognize a gain of more than a few minutes by switching to a route different from the one she had chosen.

As one may anticipate, at equilibrium predictability implies stability. If a player's predictions of driving times is likely to become the real driving times, it is unlikely that she would want to change her chosen route once she observes her predicted times. Lack of stability, on the other hand, implies (potential) *chaos.* For example, unstable driving pattern means that players have the incentives to deviate from their optimally chosen routes. Such deviations may cause further deterioration in the predicted driving times, which may lead other drivers to deviate from their chosen routes, etc.

In price formation games, stability has an important economic interpretation. Consider environments in which production and consumption plans are made prior to a play of a period of a game, and prices are generated as the outcome of the period based on the realized production and consumption choices. No hindsight regret means that the choices made by the producers and consumers are essentially optimal at the realized prices. Thus the prices are competetive and the players plans were based on rational expectations.

In the type of big games studied in this paper, one may anticipate the emergence of *stability cycle*s in which equilibrium play of the game alternates between unpredictable chaotic periods and predictable stable periods.

Our general interest is predictability and stability in equilibrium play of big games, a highly complex topic for complete formal analysis. As discussed below, however, a big game can be naturally partitioned into segments during which the game fundamentals and player types are fixed. Under this partition the research may be separated into two stages: the first stage studies the equilibrium properties of the play within single segments, and the second stage studies the process of transitions from one segment to the next. Both these stages involve difficult conceptual and technical issues.

The main body of this paper, Part 2, focuses on the first stage, leaving the second stage for future research. But equilibrium properties discussed in Part 2 make the segement equilibria potential building blocks for the second stage, as we discuss briefly in Part 3 of this paper.

To study the equilibrium of a single segent in Part 2, we consider the play of a general Bayesian repeated game that starts with unknown but fixed state of fundamentals and privately known fixed player types. The types are correlated through the fundamentals, but they are independent conditional on any one of their possible states. In an example in Section 5, we study an equilibrium of a large repeated Cournot game. There, the state of fundamentals is the unknown probability distribution over production costs of a random producer, and types describe producers' realized production costs.

At an equilibrium play of a segment, the ability to perdict period outcomes is learned by observing past outcomes. Building on known facts from the literature on learning in Bayesian repeated games, in which players learn to *forecast probabilities* of period outcomes, this paper argues that when the number of players is large they actually learn to *predict* (and not just to forecast probabilities of) the outcomes. Moreover, the number of *learning periods* in which their predictions fail is bounded by a finite number K. This number K depends on the accuracy of the intial beliefs about the fundamentals prior to the start of the segment, i.e., the weight assigned by the common prior to the unknown state of fundamentals, s; and on the level of accuracy and stability satisfactory to the analysts of the system.

But, in addition to unknown initial beliefs, predictability and stability may be limited due to external factors of uncertainty. For example, even if the drivers in rush-hours learned to predict the equilibrium driving patterns and the resulting driving times, a major traffic accident on a given morning may destroy the accuracy of their prediction.

The observations above have meaningful implications regarding the number of predictable and stable periods in a segment. First, if the external uncertainty is high, the players ability to predict outcomes may be completely diminished. Assuming that the external uncertainty is low, we may further conclude that the minimal number of predictable stable periods in a segment is $L-K$, where L is the number of periods in the segment and K is the number of learning periods described above. $L-K$ increases when: (1) the players are better informed about the initial state of of the system, since it results in a smaller K; and (2) when the time of the next fundamental change becomes larger, since it increases the length of the segement L without affecting the value of K.

The discussion above seems to be consistent with informal view of two multi-segment games: the butter consumption game and the computing device game. We may expect significantly less stability in the latter due to the frequent fundamental changes, i.e., short Ls. Within the computing device game we may expect a smaller number of unpredictable periods after the introduction of smart watches than after the introduction of personal computers. This may be due to a smaller uncertainty about how the smart watches may affect our overall processes of computations and thus smaller K. Our ability to predict rush-hour commute times in big cities may be limited alltogether because of high external uncertainties about traffic accidents, unanticipated road conditions, etc.

It should be noted that the learning discussed above is 'rational learning,' done by players whose objective is to maximize expected payoff through the play of Bayesian equilibria. Irrational learning as in the case of naive best-reply dynamics for example, may never converge to stability. Consider for example a repeated rush hour commute with two parallel routes, A and B, both from starting point S to destination T, and drivers who wish to travel on the least congested route. If most drivers start on A then the best-reply dynamics would have them all alternate between $B, A, B, A, ...$ from the second day on, and the play is unstable on any day.

For a large Bayesian repeated game, or one-segment game as discussed above, we focus on a particular type of equilibrium called *imagined continuum* and show that it has the properties discussed above.

At an imagined-continuum equilibrium the players replace random variables of the population by their expected values. For example, if every player chooses independently and randomly $\mathcal{PC}$ or $\mathcal{M}$ with probabilities .5 and .5, then the players would assume that with certainty one half of the population ends up with $\mathcal{PC}$ and one half with $\mathcal{M}$. This assumption is made throughout the game whenever random choices are made independently, or independently conditional on past events. Moreover, in doing such compression to expected values, players ignore their own effect on population proportions. For example, a player who believes that the population is evenly divided between $\mathcal{PC}$ and $\mathcal{M}$, would hold this belief no matter which computer she herself chooses.[1].

One may think of the imagined-continuum assumptions as describing a natural behavioral view of big populations. This behavior may be motivated by the simplification obtained in replacing random variables by their deterministic expected values. In particular, in the games studied in this paper it eliminates many of the Bayesian calculations required of rational players. But for events such as predictability and stability in periods of the game, as game theorists we are careful to compute probabilities in the actual process: the one in

[1]See Mcafee [24] for an earlier use of this idea: in his paper a group of sellers offer competing mechanisms to buyers. While the analysis is performed on a finite set of sellers, these seller neglect their impact on the utility of buyers who don't participate in their mechanism

which n imagined-continuum players use best response strategies to generate events in the repeated game. After all, whether a player would want to deviate from a route that she has chosen depends on the actuall observed driving times, which are determined by the n drivers on the road and not by the hypothetical continuum that she imagines.

An importnt consequence of the imagined-continuum assumption is the existence of natural Markov perfect equilibria. In addition to the simplicity of computation and play, these equilibria can be adopted to multi-segment big games, as discussed in Part 3

1. Earlier literature

Kalai [17] and Deb and Kalai [9] illustrate that (hindsight) stability is obtained in large one-shot Bayesian games with independent player types. They discuss examples from economics, politics and computer science in which stability is highly desireable. In particular, in market games stability implies that Nash equilibrium prices are competetive.[2]

The applicability of the papers above is limited by two restrictive assumptions. First, unknown fundamentals in economics and other areas are in effect a device that *correlates* the types of the players. Second, in many applications games are played *repeatedly*.

The current paper expands the analysis of the papers above. It proposes a notion of predictability and extend the notion of stability to repeated games; and proceeds to show that, despite imperfect-monitoring and uncertain fundamentals, lack of predictability and stability is limited to a finite number of periods.

One-shot games of complete information with a continuum of players were studied by Schmeidler [31], who showed the existence of pure strategy equilibrium.[3] The imagined continuum model is different from the standard continuum-game introduced by him in several respects. For one, in a continuum game, the actual outcomes of the game are determined by the actions of the continuum of players. In the imagined-continuum model, on the other

[2] Kalai [17] also illustrates that hindsight stability implies a stronger stability propertiy, reffered to as structural robustness, and Deb and Kalai [9] allow a continuum of actions and heterogeneous types of players. The current paper does not deal with either of these features. See also Azrieli and Shmaya [5], Babichenko [6] for other directions in the literature following Kalai's paper.

[3] See [2] for citations of follow-up literature.

hand, the outcomes are determined by the actions of the n players (who imagine that they are in a continuum). Since the games studied in this paper have a finite number of n players, the outcomes of the imagined continuum are the ones to be studied.[4]

Early reference to the currently popular Markov-perfect equilibrium is Maskin and Tirole [23]. The body of this paper elaborates on some properties of the imagined-continuum Markov equilibrium, enough to present the results on predictions and hindsight stability. We refer the reader to a companion paper, Kalai and Shmaya [20], KS for short, that studies this type of equilibria in depth and presents results on how the imagined continuum Markov equilibrium offers good assymptotic approximation to standard Nash equilibrium. The current paper uses the same model as in KS, which is described in the next section.

The "learning to predict" theorem presented in this paper relies on earlier results from the rational learning literature in repeated games as in Fudenberg and Levine [11], Kalai and Lehrer [19, 18] and Sorin [32].

A pioneering paper on large repeated game is Green [15] who studies large repeated strategic interaction in a restricted setting of complete information and with further restriction to pure strategies. Green and Sabourian [30] derive conditions under which the Nash correspondence is continuous, i.e., the equilibrium profile in the continuum game is a limit of equilibrium profiles in the games with an increasing number of players. In addition to Green's paper, the myopic property of large repeated games were studied in Al-Najjar and Smorodinsky [1].

As mentioned above, compressing computations to expected value is used in a variety of current models in economics, see for example McCaffee [24], Jehiel and Koessler [16], and Angeleton et al. [3], who study a dynamic global game with fundamental uncertainty. Also the idea of a stability cycle, done in a much different environments, can be found in Tardosh [21].

[4]See our paper [20] for the discrepency of the outcomes between continuum games and standard n person games with large n.

Part 2. Big uncertain games with unknown fixed fundamentals

2. The model

A stage game is played repeatedly in an environment with an unknown fixed state of fundamentals, s (also referred to as a state of nature), by a population of n players whose fixed privately known types t^i are statistically correlated through the state s. The environment and the game are symmetric and annonymous.

We consider first the *game skeleton* that consists of all the primitives of the game other than the number of players. Augmenting the game skeleton with a number of players n, results in a fully specified Bayesian repeated game. This organization eases the presentation of asymptotic analysis, as one can keep all the primitives of the game fixed while varying only the number of players n.

Definition 1. [game] A *game skeleton* is given by $\Gamma = (S, \theta_0, T, \tau, A, X, \chi, u)$ with the following interpretation:

- S is a finite set of possible *states of nature*; $\theta_0 \in \Delta(S)$ is an initial *prior* probability distribution over S.[5]
- T is a finite set of possible player *types*.[6] $\tau : S \to \Delta(T)$ is a stochastic *type-generating function* used initially to establish types. Conditional on the state s, $\tau_s(t)$ is the probability that a player is of type t, and it is (conditionaly) independent of the types of the opponents. The selected types remain fixed throughout the repeated game.
- A is a finite set of possible player's *actions*, available to a player in every period.
- X is a countable set of *outcomes*, and for every $s \in S$ and every $e \in \Delta(T \times A)$, $\chi_{s,e}$ is a probability distribution over X. In every period $e(t, a)$ is the empirical proportion of players in the population who are of type t and choose the action a. $\chi_{s,e}(x)$ is the

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

[6]Due to the symmetry assumption, 'T is a finite set of possible player types' should be interpreted to mean that it is so for every player in the game. This is also the case for similar statements that follow.

probability of the outcome x being realized and announced at the end of the period. We assume that the function $e \mapsto \chi_{s,e}(x)$ is continuous for every s and x.

- $u : T \times A \times X \to [0, 1]$ is a function that describes the player's *payoff*: $u(t, a, x)$ is the period payoff of a player of type t who plays a when the announced period outcome is x.[7]

Example 1 (Repeated computer choice game with correlated types)**.** As in the example of the one-shot computer-choice game from Kalai [17], let $S = T = A = \{\mathcal{PC}, \mathcal{M}\}$ denote two possible *states* of nature, two possible *types* of players and two possible *actions* to select. But now these selections are done repeatedly in discrete time priods $k = 0, 1, 2, ...$

Initially, an unknown state s is chosen randomly with equal probabilitis, $\theta_0(s = \mathcal{PC}) = \theta_0(s = \mathcal{M}) = 1/2$; and conditional on the realized state s the fixed types of the n players are drawn by an independent identical distribution: $\tau_s(t^i = s) = 0.7$ and $\tau_s(t^i = s^c) = 0.3$, where s^c is the unrealized state. Each player is privately informed of her type t^i. Both s and the vector of t^is remain fixed throughout the repeated game.

Based on player i's information at the begining of each period $k = 0, 1, ...$ she selects one of the two computers, $a^i_k = \mathcal{PC}$ or $a^i_k = \mathcal{M}$. These selections determine the empircal distribution of type action pairs, e_k, where $e_k(t, a)$ is the proprtion of players who are of type t and choose the computer a in the kth period.

At the end of each period, a random sample (with replacement) of J players is selected, and the sample proportions of $\mathcal{PC}$ users $x = x_k(\mathcal{PC}\ )$ is publicly announced $(x_k(\mathcal{M}\ ) \equiv 1 - x)$. Thus, the probability of the outcome $x = y/J$ being selected (when the state is s and the period empirical distribution is e) is determined by a Binomial probability of having y successes in J tries, with a probability of success $x_k(\mathcal{PC}\ )$.

[7]All the definitions and results hold also when X is a subset of Euclidean space, and $\chi_{s,e}$ is an absolutely continuous distribution. In this case we have to assume that the function $e \mapsto \chi_{s,e}$ is continuous when the range is equipped with the L^1-norm and that the payoff function u is a Borel function. See example in Section 5

Player i's payoff in period k is the (proportion of players that her choice matches)$^{1/3}$ plus 0.2 if she chose her computer type: $u^i_k(t^i, a^i_k, x_k) = (x_k[a^i_k])^{1/3} + 0.2\delta_{a^i_k = t^i}$. The game is infinitly repeated, and a player's overall payoff is the discounted sum of her period payoffs.

2.1. **Bayesian Markov strategies.** We study a symmetric equilibrium in which all the players use the same strategy κ. Normally, a player's strategy in the repeated game specifies a probability distribution by which the player selects an action in every period, as a function of (1) her type, (2) the observed history of past publicly announced outcomes, and (3) her own past actions. However, we are only interested in a certain class of strategies, which we call 'Bayesian markov strategies' (or Markov strategies for short). When playing a markov strategy the player does not condition her selection on her own past actions. Moreover, her selection of an action depends on the past publicly announced outcomes only through a Markovian state, which is the posterior public beliefs over the state of nature.

Definition 2. A *(Bayesian) markov strategy* is a function $\kappa : \Delta(S) \times T \to \Delta(A)$.

The interpretation is that $\kappa_{\theta,t}(a)$ is the probability that a player of type $t \in T$ selects the action $a \in A$, in periods in which the 'public belief' about the state of nature is θ. The term 'public belief' is in quotes because these beliefs are not derived from an exact Bayesian reasoning, they are derived under the imagined continuum reasonnig described in the next section.

Notice that as defined, a markov strategy κ may be used by any player regardless of the number of opponents and the repetition-payoff structure.

2.2. **Beliefs in the imagined continuum model.** By the 'public belief' at the begining of period k we mean the belief over the state of nature, held by an outside observer who (1) knows the players' strategy κ and (2) has observed the past publicly announced outcomes of the game. A main feature of our definition of markov strategies and equilibrium is that these beliefs are simplified. As a result, they are different from the correct posterior conditional distributions over the state of nature. Rather, they are updated during the game using what we call 'imagined continuum reasoning'. Under imagined continuum reasoning all uncertainty about players types and actions conditioned on the state of nature is compressed to its conditional expectations, resulting in known deterministic conditional distributions. Specifically, the *public beliefs* are defined recursively by the following process:

- The intial public belief is that the probability of every state s is $\theta_0(s)$.
- In every period that starts with a public belief θ, the imagined empirical proportion of a type-action pair (t, a) in the population is

$$d_\theta(t, a) = \tau_s(t) \cdot \kappa_{\theta,t}(a). \tag{2.1}$$

  And the posterior public belief assigned to every state s is computed by Bayes rule to be

$$\beta_{\theta,x}(s) \equiv \frac{\theta(s) \cdot \chi_{s,d_\theta}(x)}{\sum_{s' \in S} \theta(s') \cdot \chi_{s',d_\theta}(x)}. \tag{2.2}$$

However, even when a player ignores the impact of her type and actions on periods outcomes, she still has additional information for assessing probabilities of states of nature, namely her own realized type. Under imagined continuum reasoning, her type and the public outcome are conditionally independent of each other for any given state of nature. This implies that we can use Bayes formula to compute her *private belief* about the state of nature from the public belief.

Formally, in every period that starts with the public belief θ, for any player of type t the private belief probability assigned to the state of nature s is

$$\theta^{(t)}(s) \equiv \frac{\theta(s) \cdot \tau_s(t)}{\sum_{s' \in S} \theta(s') \cdot \tau_{s'}(t)}. \tag{2.3}$$

2.3. **Markov perfect equilibrium.** We are now in a position to define the equilibrium concept used in this paper. Under the imagined continuum view the players ignore the impact of their own action on the outcome, and a player of type t believes the outcome is drawn from the distribution $\phi(\theta^{(t)}, \theta)$ where $\theta^{(t)}$ is given by (2.3) and $\phi : \Delta(S) \times \Delta(S) \to \Delta(X)$ is given by

$$\phi(\mu,\theta)=\sum_{s\in S}\mu(s)\chi_{s,d_\theta}, \tag{2.4}$$

where d_θ is given by (2.1). Thus, $\phi(\mu,\theta)$ is the forecast about the period outcome of an observer whose belief about the state of nature is μ when the public belief about the state of nature is θ.

Definition 3. A (symmetric, *imagined continuum) Markov (perfect) equilibrium* is given by a Markov strategy $\kappa:\Delta(S)\times T\to\Delta(A)$ such that

$$[\kappa(\theta,t)]\subseteq \operatorname{argmax}_a \sum_{x\in X}\phi(\theta^{(t)},\theta)(x)u(t,a,x)$$

for every public belief $\theta\in\Delta(S)$ about the state of nature and every type $t\in T$ where $[\kappa(\theta,t)]$ is the support of $\kappa(\theta,t)$, the private belief $\theta^{(t)}$ is given by (2.3), ϕ is given by (2.4) and d is given by (2.1).

According to the imagined continuum equilibria, each player of type t treats the public outcome as a random variable with distribution $\phi(\theta^{(t)},\theta)$, ignoring her impact on the outcome. This is a generalization of the economic 'price-taking' property in Green [14] to a *stochastic* setting and to applications other than market games. For this reason our players may be viewed as *stochastic outcome takers*. Note that imagined continuum equilibria are, by definition, myopic: At every period the players play an imagined continuum equilibrium in the one shot Bayesian game for that period.

Remark 1. In our companion paper [20] we define the notion of imagined continuum equilibrium more generally (without assuming the Markov property and myopicity) and prove that: (1) every imagined continuum equilibrium is myoptic, (2) probabilities of certain outcomes computed in the imagined game approximate the real probabilities computed in the finite large versions of the game, and (3) best responses (and Nash equilibrium) in the imagined game are uniformly ϵ best responses (and ϵ Nash equilibrium) for all sufficiently large finite versions of the game.

Notice also that under myopicity, the equilibrium that we study and the main results that follow are applicable for a variety of repetition and payoff structures: For example the game may be repeated for a finite number of periods with overall payoffs assessed by the average period payoff; or the game may be infinitely repeated with payoffs discounted by different discount paramters by different players, etc.

Another consequence of myopicity is that the set of players may change and include combinations of long-lived players, short-lived players, overlapping generations, etc., provided that the death and birth process (i) keeps the size of the population large, (ii) does not alter the state and the players' type distribution, and (iii) that players of a new generation are informed of the latest public belief about the unknown state.[8]

2.4. **The induced play path.** To compute the probability of actual events in the game, as done in the sequel, we need to describe the actual (as opposed to the beliefs that are derived from the imagined continuum reasoning) probability distribution induced over play paths when players follow a markov strategy κ.

We use bold face letters to denote random variables that assume values from corresponding sets. For example, $\mathbf{S}$ is the random variable that describes a randomly-selected state from the set of possible states S. Superscripts denote players' name, superscripts in parenthesis denote players' type and subscripts denote periods' number.

The definition below is applicable to a game with a set of n players, N, with any repetition-payoff specification. As already stated, all the players use the the same strategy κ.

Definition 4. Let κ be a markov strategy of the finite game with n players. The *random κ play-path* is a collection $\left(\mathbf{S}, \mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}_k\right)_{i \in N, k=0,1,\ldots}$ of random variables,[9]representing the state of nature, types, actions and outcomes, such that:

- The state of nature $\mathbf{S}$ is distributed according to θ_0.

[8]Games in which the number of players is large, unknown and follows a Poisson distribution were studied in Myerson [25]. By restricting ourselves to games of proportions, lack of knowledge of the number of players becomes a trevial issue in the current paper.

[9]We do not specify the probability space or the domain of these variables, but only the probability distribution over their values. The play-path is unique in distribution.

- Conditional on $\mathbf{S}$, the players types $\mathbf{T}^i$ are independent and identically distributed with the distribution $\tau_{\mathbf{S}}$.
- Conditional on the history of periods $0,\dots,k-1$, players choose period k actions $\mathbf{A}^i_k$ independently of each other. Every player $i \in N$ uses the distribution $\kappa_{\mathbf{T}^i,\ \mathbf{\Theta}_k}$, where $\mathbf{\Theta}_k$ is the public belief at the beginning of period k, given by

$$\begin{aligned}&\mathbf{\Theta}_0 = \theta_0, \text{ and}\\ &\mathbf{\Theta}_{k+1} = \beta_{\mathbf{\Theta}_k,\mathbf{X}_k}, \text{ for } k \geq 0,\end{aligned} \tag{2.5}$$

  and β is defined in (2.2).
- The outcome $\mathbf{X}_k$ of period k is drawn randomly according to the distribution $\chi_{\mathbf{S},\mathbf{e}_k}$, where

$$\mathbf{e}_k(t,a) = \#\{i \in N | \mathbf{T}^i = t, \mathbf{A}^i_k = a\}/n \tag{2.6}$$

  is the (random) *empirical type-action distribution* in period k.

In equations,

$$\begin{aligned}&\mathbb{P}\left(\mathbf{S}=s,\mathbf{T}^i=t^i\ i\in N\right)=\theta_0(s)\cdot\prod_{i\in N}\tau_s(t^i).\\ &\mathbb{P}\left(\mathbf{A}^i_k=a^i\ i\in N\,\middle|\,\mathbf{S},\mathbf{T}^i,\mathbf{A}^i_l,\mathbf{X}_l\ l<k,i\in N\right)=\prod_{i\in N}\kappa_{\mathbf{\Theta}_k,\mathbf{T}^i}(a^i)\\ &\mathbb{P}\left(\mathbf{X}_k=x\,\middle|\,\mathbf{S},\mathbf{T}^i,\mathbf{A}^i_l\ l\leq k,i\in N,\mathbf{X}_0,\dots,\mathbf{X}_{k-1}\right)=\chi_{\mathbf{S},\mathbf{e}_k}(x).\end{aligned} \tag{2.7}$$

where $\mathbf{e}_k$ is given by (2.6), and $\mathbf{\Theta}_k$ is given by (2.5).

Note that the imagined continuum reasoning enters our definition only through (2.5), which reflects the way that the outside observer and the players process information. The assumption of imagined continuum reasoning lies behind the simple form of the public beliefs process $\mathbf{\Theta}_0,\mathbf{\Theta}_1,\dots$. Two important properties are a consequence of this definition: (1) $\mathbf{\Theta}_k$ admits a recursive formula (i.e., that the outside observer and the players needs only keep track on their current belief about state of nature and not on their beliefs about players types and actions) and (2) the fact that this formula does not depends on the number of players. Both these properties do not hold about the beliefs $\mathbb{P}\left(\mathbf{S} \in \cdot \,|\mathbf{X}_0, \ldots, \mathbf{X}_{k-1}\right)$ of the game theorists who do the exact Bayesian reasoning.

3. Correct predictions

Consider a game skeleton Γ played repeatedly by n players. Let κ be a markov strategy and consider the random κ play path $\left(\mathbf{S}, \mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}_k\right)_{i \in N, k=0,1,\ldots}$.

Recall that we denote by $\boldsymbol{\Theta}_k$ the public belief about the state of nature at the beginning of period k, given by (2.5). For every type $t \in T$ let

$$\boldsymbol{\Theta}^{(t)}_k(s) = \frac{\boldsymbol{\Theta}(s) \cdot \tau_s(t)}{\sum_{s' \in S} \boldsymbol{\Theta}(s') \cdot \tau_{s'}(t)} \tag{3.1}$$

be the belief of a player of type t about the state of nature computed under the imagined continuum reasoning, as in (2.3). Also let $\boldsymbol{\phi}^{(t)}_k = \phi(\boldsymbol{\Theta}^{(t)}_k, \boldsymbol{\Theta}_k)$ be the probability distribution of the $\Delta(X)$-valued random variable that represents the forecast of a player of type t about period k-outcome, where the forecast function ϕ is given by (2.4). In this section we give conditions under which these probabilistic forecasts can be said to predict the outcome.

We assume hereafter that the space X of outcomes is equipped with a metric η. The event that *players make* (r, ϵ)*-correct predictions in period* k is given by

$$R(k, r, \epsilon) = \left\{ \boldsymbol{\phi}^{(t)}_k \left(B(\mathbf{X}_k, r)\right) > 1 - \epsilon \text{ for every } t \in T \right\} \tag{3.2}$$

where $B(\mathbf{X}_k, r) = \{x \in X | \eta(x, \mathbf{X}_k) \leq r\}$ is the closed ball of radius r around $\mathbf{X}_k$. Thus, players make (r, ϵ) correct predictions at period k if each player assigns probability at least $1 - \epsilon$ to a ball of radius r around the realized outcome $\mathbf{X}_k$, before she observes its realized value.

Definition 5. Let Γ be a game skeleton and let κ be a Markov strategy. We say that *players make asymptotically* (r, ϵ, ρ)*-correct predictions under* κ *in period* k if there exists some n_0 such that

$$\mathbb{P}(R(k, r, \epsilon)) > 1 - \rho.$$

in every n-player game with $n > n_0$.

We proceed to provide conditions on the game skeleton under which players make asymptotically correct predictions. For every probability distribution function ν over X let $Q_\nu : [0, \infty) \to [0, 1]$ be the *concentration function* of ν given by

$$Q_\nu(r) = 1 - \sup_D \nu(D) \tag{3.3}$$

where the supremum ranges over all closed subsets B of X with diameter $\operatorname{diam}(D) \le r$. ($\operatorname{diam}(D) = \sup_{x,x' \in B} \eta(x, x')$ where η is the metric on X). When ν is the distribution of a random variable $\mathbf{X}$, we also denote $Q_\mathbf{X} = Q_\nu$. The following are examples of concentration functions:

- If for some a, $a \le \mathbf{X} \le a + .01$ then $Q_\mathbf{X}(0.01) = 0$.
- If X is a finite set and $\mathbb{P}(\mathbf{X} = x_0) = 1 - \epsilon$ for some $x_0 \in X$ and small $\epsilon > 0$ then $Q_f(0) = \epsilon$.
- If $X = \mathbb{R}$ and $\mathbf{X}$ is a random variable with variance σ^2 then from Chebyshev's Inequality it follows that $Q_\mathbf{X}(r) \le 4\sigma^2/r^2$.
- If $X = \mathbb{R}$ and $\mathbf{X}$ is a random variable with Normal distribution with standard deviation σ then $Q_\mathbf{X}(r) = 2(1 - \Phi(r/2\sigma)) \le 2\exp(-r^2/2\sigma^2)$

For every game skeleton Γ we let $Q_\Gamma : [0, \infty) \to [0, 1]$ be given by $Q_\Gamma(r) = \sup_{s,e} Q_{\chi_{s,e}}(r)$. For example, in the *round-off case*, where outcomes are empirical distribution randomly rounded off to integral percentage, it holds that $Q_\Gamma(0.01) = 0$.

Theorem 1 (Correct predictions)**.** *Fix a game skeleton Γ. For every $\epsilon, \rho > 0$ there exists an integer K such that under every Markov strategy κ and every $r > 0$, in all but at most K periods players make $[r, Q_\Gamma(r) + \epsilon, Q_\Gamma(r) + \rho]$-asymptotically correct predictions.*

The apperance of $Q_\Gamma(r)$ in Theorem 1 is intuitively clear: Increasing concentration of the random outcome (e.g., by taking a larger sample size J in Example 1) imporves the level of predictability and stability. But if the variance is large (e.g., small sample size in the example) predictability and stability are not to be expected.

4. Stability

Consider a game skeleton Γ played repeatedly by n players. Let κ be a markov strategy and consider the random κ play path $\left(\mathbf{S}, \mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}_k\right)_{i \in N, k=0,1,\ldots}$. The event that period k is *ϵ-hindsight stable* is given by

$$H(k, \epsilon) = \left\{ u(\mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}^i) + \epsilon \geq u(\mathbf{T}^i, a, \mathbf{X}^i) \text{ for every player } i \text{ and action } a \in A \right\}.$$

This is the event that after observing the realized outcome of period k no player can improve her payoff by more than ϵ through a unilateral revision of her period-k action.

Definition 6. Let Γ be a game skeleton and let κ be a Markov strategy. We say that period k is *asymptotically (ϵ, ρ)-stable under κ* if there exists some n_0 such that

$$\mathbb{P}(H(k, \epsilon)) > 1 - \rho$$

in every n-player game with $n > n_0$.

We proceed to provide bounds on the level of hindsight stability in natural classes of large games. For this purpose, in addition to the earlier assumptions on the game skeleton, we now make an assumption about the modulus of continuity of the payoff function. Let $\omega : [0, \infty) \to [0, \infty)$ be continuous, monotone increasing function with $w(0) = 0$. We say that the utility function u admits ω as a *modulus of continuity* if $|u(t, a, x) - u(t, a, x')| \leq \omega(\eta(x, x'))$ for all $t \in T, a \in A$ and $x, x' \in X$, where η is the metric on X. The special case of *Lipschitz payoff function with constant* L is described by the function $\omega(d) = Ld$.

The following lemma says that correct predictions implies hindsight stability.

Lemma 1. *Fix $r, \epsilon > 0$. Let κ be a markov strategy of a game skeleton Γ in which the payoff function u has modulus of continuity ω, and consider the random κ-play path. For every period k, the event $R(k, r, \epsilon)$ that players make (r, ϵ)-correct predictions is contained in the event $H(k, 2\omega(r) + \epsilon/(1 - \epsilon))$ that the game is $(2\omega(r) + \epsilon/(1 - \epsilon))$-stable.*

The following theorem follows from Theorem 1 and Lemma 1

Theorem 2. *[Hindsight Stability] Fix a game skeleton* Γ *in which the payoff function* u *has modulus of continuity* ω *. Then for every* $\epsilon, \rho > 0$ *there exists an integer* K *such that in every Markov equilibrium* κ *and every* $d > 0$ *all but at most* K *periods are* $[2d + 2Q_\Gamma(\omega^{-1}(d)) + 2\epsilon, Q_\Gamma(\omega^{-1}(d)) + \rho]$ *asymptotically-stable*[10]

Next, we select $d = 2(\sigma L)^{2/3}$ and restrict ourselves to applications in which the public information X has a finite variance bounded by σ^2 and the payoff function is Lipschitz with constant L. Using Chebyshev's bound on the concentration function, we obtain the following.

Corollary 1. *Consider any game skeleton* Γ *in which the payoff function* u *is Lipschitz with constant* L *and the public signal has variance bounded by* σ^2*. Then for every* $\epsilon, \rho > 0$ *there exists a finite integer* K *such that in every Markov equilibrium* κ *all but at most* K *periods are [*$6(\sigma L)^{2/3} + 2\epsilon, (\sigma L)^{2/3} + \rho$*] asymptotically-stable.*

Under the theorems above we deduce the following examples.

Example 2 (Rouded-off empirical distribution)**.** Consider games Γ in which the reported outcome x is the realized empirical distribution of the population e, randomly rounded off up or down to the nearest percentage; games in which $Q_\Gamma(0.01) = 0$. Let r, ϵ, ρ be arbitrary small positive numbers, then there is a finite number of periods K such that:

A. In all but the K periods, under any strategy κ the players make correct predictions up to $[r, \epsilon, \rho]$.

B. If the payoff function is Lipschitz with constant $L = 1$ and κ is a Markov equilibrium, all but K periods are $(0.02 + 2\epsilon, \rho)$-stable.

5. Cournot example: Price stability

In this example of an n-person Cournot production game, the state of nature determines whether it is easy or difficult to produce a certain good, and producers are of two types, efficient and inefficient. At the begining of every period, each one of the producers chooses

[10]In case that the function ω is not invertible then $r = \omega^{-1}(d)$ is defined in any way such that $\omega(r) \leq d$.

whether or not to produce a single unit of the good. The total production determines the period's price through a random inverse demand function.

Let $S = \{\text{easy}, \text{difficult}\}$ denote the set of possible states equipped with a uniform prior $\theta_0(\text{easy}) = \theta_0(\text{difficult}) = 1/2$, and let $T = \{\text{efficient}, \text{inefficient}\}$ denote the set of player's types. Types are generated stochastically by the conditional independent identical distributions $\rho_{|s}$: $\rho(\text{efficient}|\text{easy}) = 3/4$ and $\rho(\text{inefficient}|\text{easy}) = 1/4$; symmetirically $\rho(\text{efficient}|\text{difficult}) = 1/4$ and $\rho(\text{inefficient}|\text{difficult}) = 3/4$.

A player's period production levels are described by the set of actions $A = \{0, 1\}$, and a price $x \in R$ is the outcome of every period. The period price depends entirely on the period's total production, and not on the state of nature and the types. Formally for every $s \in S$ and empirical distribution of type-action pairs $e \in \Delta(T \times A)$, $\chi_{s,e} = \text{Normal}(1/2 - r, \sigma^2)$, where $r = e(\text{easy}, 1) + e(\text{difficult}, 1)$ is the proportion of players who produce the good. One interpretation in the n player game is that there are n buyers whose demand at price x is given by $1/2 - r + \epsilon$ where $\epsilon \sim \text{Normal}(0, \sigma^2)$ is the same for all buyers. Another interpretation is that ϵ represents noisy traders who may either buy or sell the good.

The payoff functions are given by $u(t, 0, x) = 0$ for every $t \in T$ and $x \in X$, i.e., not producing results in zero payoff; and $u(t, 1, x) = x - (1/8)\delta_{t=\text{inefficient}}$, i.e., per unit production cost is zero for an efficient producer and $1/8$ for an ineficient one.

The repeated game admits the following unique imagined continuum Markov equilibrium: Let θ_k be the public belief about the state of nature at the beginning of period k, computed (according to the imagined continuum reasoning) by an outsider who observes the prices but not the players types and actions. We identify θ_k with the probability assigned to $s = \text{easy}$, so $\theta_k \in [0, 1]$. Note that if the public belief is θ_k then the belief of every efficient players is

$$\theta_k^{(\text{efficient})} = \frac{3/4 \cdot \theta_k}{3/4 \cdot \theta_k + 1/4(1 - \theta_k)} = \frac{3\theta_k}{1 + 2\theta_k}$$

and the belief of every inefficient players is

$$\theta_k^{(\text{inefficient})} = \frac{\theta_k}{3 - 2\theta_k}.$$

The equilibrium strategies in the repeated game are defined by the following:

(1) When $\theta_k \geq (7+\sqrt{33})/16 = 0.796..$ each efficient player produces with prob $p = \frac{4\theta_k+2}{8\theta_k+1}$ (thus under conpressed reasoning $\frac{4\theta_k+2}{8\theta_k+1}$ of them produce) and the inefficient players are idle. Here p is the solution to the equation

$$\theta_k^{(\text{efficient})} 3/4 \cdot p + (1-\theta_k^{(\text{efficient})}) 1/4 \cdot p = 1/2,$$

so that the efficient players expect a selling price of 0 and zero profit. In particular, when $\theta_k = 1$ a proportion $p = 2/3$ of the efficient players produce and the inefficient players are idle.

(2) When $(35-\sqrt{649})/64 < \theta_k < (7+\sqrt{33})/16$, each efficient player produces with probability p and each ineficient player produces with probability q, where $0 < p, q < 1$ are the unique solution to the equations

$$\theta_k^{(\text{efficient})}(3/4 \cdot p + 1/4 \cdot q) + (1-\theta_k^{(\text{efficient})})(1/4 \cdot p + 3/4 \cdot q) = 1/2$$
$$\theta_k^{(\text{inefficient})}(3/4 \cdot p + 1/4 \cdot q) + (1-\theta_k^{(\text{inefficient})})(1/4 \cdot p + 3/4 \cdot q) = 3/8$$

so that the efficient players expect price 0 and the inefficient players expect price 1/8. For example, when $\theta = 1/2$ the strategies are $p = 11/16$ and $q = 3/16$.

(3) When $\theta \leq (35-\sqrt{649})/64 = 0.148...$ the efficient players all produced and the inefficient players produce with probability $q = (3-6\theta)/(18-16\theta)$. Here q is the solution to the equation

$$\theta_k^{(\text{inefficient})} \cdot (3/4 + 1/4 \cdot q) + (1-\theta_k^{(\text{efficient})}) \cdot (1/4 + 3/4 \cdot q) = 3/8,$$

so that the inefficient player expect price 1/8 and zero profit. In this case the efficient players expect a positive profit.

After each period the players update their beliefs using Bayes' formula:

$$\theta_{k+1} = \frac{\theta_k \cdot \exp\left(-(x_k - (3/4p_k + 1/4q_k))^2/2\right)}{\theta_k \cdot \exp\left(-(x_k - (3/4p_k + 1/4q_k))^2/2\right) + (1-\theta_k) \cdot \exp\left(-(x_k - (1/4p_k + 3/4q_k))^2/2\right)}$$

where p_k and q_k are the equilibrium strategy under θ_k and x_k is the outcome of day k.

By Theorem 2 it follows that for every $\epsilon, \rho > 0$ and every $d > 0$ every period except for a finite number is asymptotically hindsight stable at a level: $(2d + 2Q_G(d) + 2\epsilon, Q_G(d) + \rho)$. Assume for example that $\sigma = 0.01$. Choosing $d = 0.05$ we get $Q_G(d) = 0.012$. Therefore, every period except for a finitely number is asymptotically $(0.11 + 2\epsilon, 0.012 + \rho)$ stable.

Remark 2. Why is the equilibrium unique? Let θ be the outsider belief about the state of nature at the beginning of some period. Let p be the proportion of efficient players who produce at that period and q the proportion of inefficient players who produce.

Under this profile the supplied quantity that the efficient players expect is

$$\theta^{(\text{efficient})}(3/4 \cdot p + 1/4 \cdot q) + (1 - \theta^{(\text{efficient})})(1/4 \cdot p + 3/4 \cdot q)$$

and the supplied quantity that the inefficient players expect is

$$\theta^{(\text{inefficient})}(3/4 \cdot p + 1/4 \cdot q) + (1 - \theta^{(\text{inefficient})})(1/4 \cdot p + 3/4 \cdot q)$$

Assume now that (p, q) and (p', q') are two equilibrium profiles and that $q > q'$. The equilibrium condition implies that the supplied quantity that the inefficient players expect under (p, q) is weakly smaller than what they expect under (p', q'). Because $q > q'$ this implies that $p < p'$, so that the supplied quantity that the efficient players expect under (p, q) is weakly larger from under (p', q'). This is a contradiction since the difference between the expected supplied quantities of the efficient and inefficient players is monotone increasing in p and monotone decreasing in q.

Part 3. Constructing equilibria of multi-segment big games

We next discuss how the imagined-continuum equilibria of a single segment may be used for constructing equilibria for multisegment big games. Recall that by a segment we mean a sequence of consecutive periods in which the state of nature (or fundamentals) is constant. In a multisegment big game, in which the state of nature changes over time, the periods are naturally partitioned into such segments.

We consider a simple family of multisegment games, *repeated segment games*, in which the changes in fundamentals are due to exogenous random shocks. Moreover, when the fundamentals chage players learn of the change, even though they do not know the new state of fundamentals. A repeated segment game is described by a pair (Γ, ξ) in which $\Gamma = (S, \theta_0, T, \tau, A, X, \chi, u)$ is a segment game, as discussed earlier, and $1 > \xi > 0$ describes a probability of change.The extensive form of the game is described recursively as follows.

Initially and in later stages the segment game Γ is played as follows. (a) Independent of any past events, nature draws a random state of nature s according to the distribution θ_0; it also indepently draws player types by the distribution τ_s and privately informs the players of their realized types. (b) The segment game Γ is played repeatedly as described earlier but subject to the following modification: At the end of every period, independently of the history, with probability ξ the game *restarts* and with probability $1 - \xi$ it *continues*. Following the *continue* event the play continues to the next period of the segment. But following the *restrat* event the players are informed that the game is restarted, and proceed to play a new segment starting in step (a) above.

Let κ be a Bayesian Markov strategy in the one-segment game Γ, according Definition 3. Then κ induces a strategy in the multi-segment game: Initially and after every change in fundamentals, the public belief is set to be θ_0 and players use their current type when applying κ throughout the segment. Even without giving a formal definition of an equilibrium in the multi-segment game, it is intuitively clear that if κ is an equilibrium in the one segment game then the induced strategy is an equilibrium in the multi-segment game. Our results imply that in every segment there is a bounded number K of periods which are not hinsdight stable. If the probability of transition ξ is small then we get a bound $\xi \times K$ on the frequency of periods in the multisegment game which are not hinsdight stable

Moving to more general models, we may view big games as large imperfectly observed Bayesian stochastic games, where the transition probability and the new state depends on the empirical distribution of players actions and on the current state. This broader view gives rise to many questions for future research. Direct follow up of the current paper are issues regarding changes in information as we move from one segement of the game to the next. In addition, one may consider situations in which changes in fundamentals are not made public. We leave these issues to future research.

Part 4. Proofs

The main result of the paper, that asymptotic hindsight stability holds in all but finitely many chaotic learning periods, is proven in two steps.

Step one argues that the result holds in the imagined processes that describe the beliefs of the players. Building on the result of step one, step two shows that the result holds in the real process.

In step one, the intuition may be broken into two parts. First, relying on the merging literature (grain of truth is automatic in our model; see Fudenberg and Levin [11], Sorin [32] and Kalai and Lehrer [18, 19]) we argue that in an equilibrium of our model there are only a finite number of learning periods in which the forecasted probability of the period outcome is significantly different from its real probability. In other words, the players' belief about the fundamental s leads to approximately the same probability distribution over the future events as the real s. One issue we need to address in applying these results is that in our multi-player setup players with different types have different beliefs, and so may make mistake in forecasts in different periods. But we need to bound the number of periods in which some player makes a forecasting mistake. To do that we extend the previous learning result to a multi-player setup.

The second stage is based on the following reasoning. Assuming that the uncertainty in the determination of period outcomes is low, in the imagined process for every state s the period outcomes are essentially derterministic. This implies that in every nonlearning period the players learn to predict (and not just forecast the probability of) the outcomes. When the predicted period outcome (on which a player base her optimal choice of an action) is correct, she has no reason to revise her choice. Thus, in the imagined processes we have hindsight stability in all the non-chaotic periods.

In step two, to argue that hindsight stability holds in all the non-learning periods in the real process, we rely on arguments developed in our companion paper [20]. These arguments show that the probability of events in the real process are approximately the same as their counterparts in the the imagined processes. Thus the high level of hindsight stability obtained in step one apply also to the real process.

Section 6 gives a formal definition of the imagined play-path, which is the process the players have in mind when doing the 'incorrect' updating given in Section 2.2. Section 7 presents the result from our companion paper, that when the number of players is large the imagined process is not too far from the induced play-path given in Section 2.4. Section 8 presents a uniform merging result: In an environment with many player types, each starting with a different signal we provide a bound for the number of periods in which one of them changes their beliefs. Section 9 connect the dots.

6. Imagined Continuum view

In this section we describe the imagined play path, that reflects the players imagined continuum reasoning. In order to distinguish between corresponding entities in the actual play path and in the imagined play path, we denote the random variables that represent the outcomes in the imagined play by $\tilde{\mathbf{X}}_0, \tilde{\mathbf{X}}_1, \dots$, and the random variables that present public beliefs by $\tilde{\mathbf{\Theta}}_0, \tilde{\mathbf{\Theta}}_1, \dots$.

Let κ be a markov strategy. A *imagined random κ-play path* is a collection $\left(\mathbf{S}, \tilde{\mathbf{T}}, \tilde{\mathbf{X}}_0, \tilde{\mathbf{X}}_1, \dots\right)$ of random variables, representing the state of nature, type of a representative player and outcomes, such that: The state of nature $\mathbf{S}$ is distributed according to θ_0 and conditional on the history of periods $0, \dots, k-1$, the outcome $\tilde{\mathbf{X}}_k$ is drawn randomly according to probability density function $\chi_{\mathbf{S}, d_{\tilde{\mathbf{\Theta}}_k}}$ where the imagined public beliefs $\tilde{\mathbf{\Theta}}_k$ are given by

$$
\begin{aligned}
&\tilde{\mathbf{\Theta}}_0 = \theta_0, \text{ and} \\
&\tilde{\mathbf{\Theta}}_{k+1} = \beta_{\tilde{\mathbf{\Theta}}_k, \tilde{\mathbf{X}}_k}, \text{ for } k \geq 0,
\end{aligned}
\tag{6.1}
$$

β is defined in (2.2), and d_θ for every belief θ is defined in (2.1).

In equations,

$$
\begin{aligned}
&\mathbb{P}\left(\mathbf{S}=\cdot\right)=\theta_0\\
&\mathbb{P}\left(\tilde{\mathbf{T}}=\cdot|\mathbf{S}\right)=\tau_{\mathbf{S}}\\
&\mathbb{P}\left(\tilde{\mathbf{X}}_k=\cdot\,\middle|\,\mathbf{S},\tilde{\mathbf{T}},\tilde{\mathbf{X}}_0,\ldots,\tilde{\mathbf{X}}_{k-1}\right)=\chi_{\mathbf{S},d_{\tilde{\mathbf{\Theta}}_k}}
\end{aligned}
\tag{6.2}
$$

The difference between Equations (6.2) and the Equations (2.7) that defined the actual random play-path is that in the latter the outcome is generated from the random empirical types-actions distribution $\mathbf{e}_k$ of n players whereas in the former the outcome is generated from the conditional expectation $d_{\tilde{\mathbf{\Theta}}_k}$ of this distribution. It is for this reason that the beliefs $\tilde{\mathbf{\Theta}}_k$ are the correct conditional probabilities over the state of nature of an observer who views the outcome process $\tilde{\mathbf{X}}_0,\tilde{\mathbf{X}}_1,\ldots$ and $\tilde{\mathbf{\Theta}}_k^{(t)}$ are the correct conditional probabilities over the state of nature of a player of type $\mathbf{T}$:

$$
\begin{aligned}
\tilde{\mathbf{\Theta}}_k&=\mathbb{P}\left(\mathbf{S}=\cdot\,\middle|\,\tilde{\mathbf{X}}_0,\ldots,\tilde{\mathbf{X}}_{k-1}\right),\text{ and}\\
\tilde{\mathbf{\Theta}}_k^{(t)}&=\mathbb{P}\left(\mathbf{S}=\cdot\,\middle|\,\mathbf{T}=t,\tilde{\mathbf{X}}_0,\ldots,\tilde{\mathbf{X}}_{k-1}\right)
\end{aligned}
\tag{6.3}
$$

For every $t\in T$ where $\tilde{\mathbf{\Theta}}_k^{(t)}$ are given by

$$
\tilde{\mathbf{\Theta}}_k^{(t)}(s)=\frac{\tilde{\mathbf{\Theta}}(s)\cdot\tau_s(t)}{\sum_{s'\in S}\tilde{\mathbf{\Theta}}(s')\cdot\tau_{s'}(t)}
$$

as in (3.1). Similarly, the forecasts of the public observer and the players about the next day outcome are correct in the imagined process:

$$
\begin{aligned}
\phi(\tilde{\mathbf{\Theta}}_k,\tilde{\mathbf{\Theta}}_k)&=\mathbb{P}\left(\tilde{\mathbf{X}}_k=\cdot\,\middle|\,\tilde{\mathbf{X}}_0,\ldots,\tilde{\mathbf{X}}_{k-1}\right),\text{ and}\\
\phi(\tilde{\mathbf{\Theta}}_k^{(t)},\tilde{\mathbf{\Theta}}_k)&=\mathbb{P}\left(\tilde{\mathbf{X}}_k=\cdot\,\middle|\,\tilde{\mathbf{T}}=t,\tilde{\mathbf{X}}_0,\ldots,\tilde{\mathbf{X}}_{k-1}\right)
\end{aligned}
\tag{6.4}
$$

From (6.4) it follows that if κ is an imagined continuum equilibrium then the players choose at every round optimal actions for the imagined beliefs:

$$
[\kappa(t,\tilde{\Theta}_k)]\in\arg\max_a\mathbb{E}\left(u(t,a,\tilde{\mathbf{X}}_k)\,\middle|\,\tilde{\mathbf{T}}=t,\tilde{\mathbf{X}}_0,\ldots,\tilde{\mathbf{X}}_{k-1}\right). \tag{6.5}
$$

for period every k and every player's type $t \in T$. As mentioned in Remark 1, in our companion paper we define imagined equilibrium (not necessarily Markovian) using this property and prove that such every equilibrium is myopic.

7. Validation of the imagined view

We prove the theorem using a lemma that couples a play path in the actual game with an imagined play path, that reflects the players imagined continuum reasoning. By coupling we mean that both processes are defined on the same probability space. The coupling presented in Proposition 1 is such, when the number of players is large, the realization of the processes is, with high probability, the same. In particular, the forecasts about the outcome sequence made by the imagined continuum reasoning are not far from the correct forecasts made by an observer that performs the correct Bayesian calculation[11]. We prove Proposition 1 in our companion paper[12]. See also Carmona and Podczeck [8] and the reference therein for results of similar spirit in a static (single period) game.

Proposition 1. *Fix a game skeleton and a markov strategy κ. There exist random variables $\mathbf{S}, \mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}_k, \tilde{\mathbf{X}}_k$ for $i \in N$ and $k = 0, 1, \dots$ such that*

- $(\mathbf{S}, \mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}_0, \mathbf{X}_1, \dots)$ *is a random κ-play path of the repeated game.*
- *The outcome sequence $\mathbf{S}, \tilde{\mathbf{X}}_0, \tilde{\mathbf{X}}_1, \dots$ is a imagined random κ-play path.*
- *For every k it holds that*

$$\mathbb{P}\left(\mathbf{X}_0 = \tilde{\mathbf{X}}_0, \dots, \mathbf{X}_k = \tilde{\mathbf{X}}_k\right) > 1 - C \cdot k\sqrt{\frac{\log n}{n}}, \tag{7.1}$$

where C is a constant that depends on the game skeleton.

Let V_k be the event $\{\mathbf{X}_0 = \tilde{\mathbf{X}}_0, \dots, \mathbf{X}_k = \tilde{\mathbf{X}}_k\}$. This is the event that the continuum behavior model was validated up to day k.

[11]In a sense, the lemma claims that the incorrect, imagined continuum reasoning is validated by what the observations of the players and an outside observers. This is a similar idea to self-confirming equilibrium [12] but in self-confirming equilibrium the farecasts are fully correct.

[12]See Lemma 1 in that paper. The version of the lemma in that paper is more general than here in that it does not assume that all players play the same markov strategy and also allows arbitrary (non-markovian) deviation of a player.

8. A uniform merging theorem

The proof of Lemma 1 relies on the notion of merging. Let X be a Borel space of outcomes and let P, Q be two probability distributions over $X^{\mathbb{N}}$. For every $x = (x_0, x_1, \dots) \in X^{\mathbb{N}}$ we denote by $P_k(x_0, \dots, x_{k-1})$ and $Q_k(x_0, \dots, x_{k-1})$ the forecasts made by P and Q respectively over the next day outcome conditioned on $x_0, \dots, x_{k-1}$.

Say that Q is *δ-grain* of P if $P = \delta Q + (1-\delta)Q'$ for some distribution Q'. One can think of Q as the 'correct distribution' of some X-valued stochastic process and of P as some agent's belief. The Bayesian learning literature [Blackwell & Dubins, Kalai & Lehrer, Sorin, Fudenberg Levine] studies situations when forecasts made according to P are close in various senses to forecasts made according to Q. A seminal result in the literature is that if Q is a δ-grain of P for some $\delta > 0$ then one can bound the number of periods in which the agent makes wrong forecasts about next period outcome. We follow Sorin's paper [32].

For our purpose there is a set of agents each of them has a belief with some grain of truth. In principle, different beliefs with grain of truth may induce wrong predictions in different days. In this section we use Sorin's result to show that in our setup we can still bound the number of periods in which at least one of the agents make a wrong forecast. Assume that S is a finite set of states and for every for every $s \in S$ let $P^s \in \Delta(X)$. For a belief $\theta \in \Delta(S)$ we denote $P^\theta = \sum_{s\in S} \theta(s) P^s \in \Delta(X)$. Thus, P^θ is the belief of a player with prior θ over the states of nature. A *stochastic signal* is given by a function $\zeta : S \to [0,1]$. The interpretation is that the agent observed an outcome that has a probability $\zeta(s)$ to happen if the state of nature is s. In our setup, an agent of type t received the stochastic signal ζ that is given by $\zeta(s) = \tau_s(t)$. An agent who has some prior θ about S and receives a signal ζ updates his belief to

$$\theta^{(\zeta)}(s) = \frac{\theta(s)\cdot\zeta(s)}{\sum_{s'\in S}\theta(s')\zeta(s')}.$$

This is the same formula as (2.3) except that we use the abstract notation of stochastic signal. Finally, for $\delta > 0$ let $Z_\delta = \{\zeta : S \to [\delta, 1]\}$ be the set of stochastic signals with probability at least δ under every state.

For every k let $D_{k,\epsilon}(P,Q) \subset X^{\mathbb{N}}$ be the set of all realizations $x = (x_0, x_1, \dots) \in X^{\mathbb{N}}$ such that $\|P_k(x_0, \dots, x_{k-1}) - Q_k(x_0, \dots, x_{k-1})\| < \epsilon$.

Proposition 2. *For every $\delta, \epsilon, \rho > 0$ there exists $K = K(\delta, \epsilon, \rho)$ such that for every prior belief $\theta_0 \in \Delta(S)$ and every collection of distributions $\{P^s \in \Delta(X^{\mathbb{N}}) | s \in S\}$, in every period k except at most K of them it holds that*

$$\sum_{s \in S} \theta_0(s) P^s \left(\cap_{\zeta \in Z_\delta} D_{k,\epsilon} \left(P^s, P^{\theta_0^{(\zeta)}} \right) \right) > 1 - \rho.$$

The meaning of the condition in proposition 2 is that if the state of nature is randomized according to θ_0 then at day k, with high probability all agents who receives signal in Z_δ make simultaneously correct forecasts, as if they knew the realized state of nature.

The following claim, which is a generalization of Cauchy-Schwartz Inequality $|\mathrm{Cov}(\mathbf{X}, \mathbf{Y})| \leq \sqrt{\mathrm{Var}(X)\mathrm{Var}(Y)}$ to random variables that assume values in a Banach space, will be used in the proof of Proposition 2. (The case we are interested in is where the random variable $\boldsymbol{\phi}$ is an agent's forecast, which assumes values in $\Delta(X)$, viewed as a subspace of the Banach space of all signed measure over X equipped with the bounded variation norm).

Claim 1. *Let $\boldsymbol{\phi}$ be a random variable which assumes values in some Banach space V and let $\boldsymbol{\zeta}$ be a real valued random variable, both bounded. Define $Cov(\boldsymbol{\zeta}, \boldsymbol{\phi}) = \mathbb{E}\boldsymbol{\zeta}\boldsymbol{\phi} - \mathbb{E}\boldsymbol{\zeta}\mathbb{E}\boldsymbol{\phi} \in V$. Then $\|Cov(\boldsymbol{\zeta}, \boldsymbol{\phi})\| \leq \sqrt{\mathrm{Var}(\boldsymbol{\zeta}) \cdot \mathbb{E}\|\boldsymbol{\phi} - \mathbb{E}\boldsymbol{\phi}\|^2}$.*

Proof. From the linearity of the expectation we get that

$$\mathrm{Cov}(\boldsymbol{\zeta}, \boldsymbol{\phi}) = \mathbb{E}\left(\boldsymbol{\zeta} - \mathbb{E}\boldsymbol{\zeta}\right)\left(\boldsymbol{\phi} - \mathbb{E}\boldsymbol{\phi}\right)$$

Therefore, it holds that

$$\|\mathrm{Cov}(\boldsymbol{\zeta}, \boldsymbol{\phi})\| = \mathbb{E}\|\left(\boldsymbol{\zeta} - \mathbb{E}\boldsymbol{\zeta}\right)\left(\boldsymbol{\phi} - \mathbb{E}\boldsymbol{\phi}\right)\| =$$
$$\mathbb{E}\left(|\boldsymbol{\zeta} - \mathbb{E}\boldsymbol{\zeta}| \cdot \|\boldsymbol{\phi} - \mathbb{E}\boldsymbol{\phi}\|\right) \leq \sqrt{\mathrm{Var}(\boldsymbol{\zeta}) \cdot \mathbb{E}\|\boldsymbol{\phi} - \mathbb{E}\boldsymbol{\phi}\|^2},$$

where the first inequality follows from Jensen's inequality and convexity of the norm, the equality from properties of the norm and the third from Cauchy Schwartz Inequality. □

Proof of Proposition 2. Consider a probability space equipped with random variables $\mathbf{S}, \mathbf{X}_0, \mathbf{X}_1, \ldots$ such that

$$\mathbb{P}(\mathbf{S} = s) = \theta_0 \text{ for every} s \in S, \text{ and}$$

$$\mathbb{P}(\mathbf{X}_0 = \cdot, \mathbf{X}_1 = \cdot, \ldots | \mathbf{S} = s) = P^s$$

for every $s \in S$. Let $\mathcal{F}_k$ be the sigma-algebra that is generated by $\mathbf{X}_0, \ldots, \mathbf{X}_{k-1}$ and let $\boldsymbol{\phi}_k = \mathbb{P}(X_k = \cdot | \mathbf{S}, \mathbf{X}_0, \ldots, \mathbf{X}_{k-1})$ be the $\Delta(X)$-valued random variable that represents the prediction about $\mathbf{X}_k$ of an agent that knows the state of nature and observed previous outcomes. Let $\gamma = \epsilon^2 \rho \delta / 4$ and $K = 2|S| \log(|S|)/\gamma = 8|S| \log(|S|)/(\epsilon^2 \rho \delta)$. It follows from standard arguments that for every period k except at most K of them it holds that

$$\mathbb{E} \left\| \boldsymbol{\phi}_k - \mathbb{E}(\boldsymbol{\phi}_k | \mathcal{F}_k) \right\|^2 < \gamma.$$

We call days k on which this inequality holds *good days*. It follows that on a good day k there exists an $\mathcal{F}_k$-measurable event G_k such that $\mathbb{P}(G_k) > 1 - \rho$ and

$$\mathbb{E} \left(\left\| \boldsymbol{\phi}_k - \mathbb{E}(\boldsymbol{\phi}_k | \mathcal{F}_k) \right\|^2 \middle| \mathcal{F}_k \right) < \gamma / \rho \tag{8.1}$$

on G_k. Note that in the last inequality $\mathbb{E}(\boldsymbol{\phi}_k | \mathcal{F}_k)$ is the prediction of about $\mathbf{X}_k$ of an agent who doesn't observe the state of nature (but knows that it is distributed according to θ_0) and observe previous outcomes. More generally, let $\zeta \in \Delta(S)$ and let $\boldsymbol{\zeta} = \zeta(\mathbf{S})$. Then the prediction about $\mathbf{X}_k$ of an agent who receives the signal ζ is given by

$$\boldsymbol{\phi}_k^{(\zeta)} = \mathbb{E}(\boldsymbol{\zeta} \boldsymbol{\phi}_k | \mathcal{F}_k) / \mathbb{E}(\boldsymbol{\zeta} | \mathcal{F}_k). \tag{8.2}$$

By the concavity of the square root function and Jensen's inequality we get from (8.1) that

$$\mathbb{E} \left(\left\| \boldsymbol{\phi}_k - \mathbb{E}(\boldsymbol{\phi}_k | \mathcal{F}_k) \right\| \middle| \mathcal{F}_k \right) < \sqrt{\gamma / \rho} \tag{8.3}$$

on G_k. Now let $\zeta \in Z_\delta$. Then from (8.2) we get that

$$\boldsymbol{\phi}_k^{(\zeta)} - \mathbb{E}(\boldsymbol{\phi}_k|\mathcal{F}_k) = \frac{\mathrm{Cov}\left(\boldsymbol{\zeta}, \boldsymbol{\phi}_k \,|\mathcal{F}_k\right)}{\mathbb{E}(\boldsymbol{\zeta}|\mathcal{F}_k)}$$

which implies by Claim 1 (conditioned on $\mathcal{F}_k$) that

$$\|\boldsymbol{\phi}_k^{(\zeta)} - \mathbb{E}(\boldsymbol{\phi}_k|\mathcal{F}_k)\| = \frac{\|\mathrm{Cov}\left(\zeta, \boldsymbol{\phi}_k \,|\mathcal{F}_k\right)\|}{|\mathbb{E}(\boldsymbol{\zeta}|\mathcal{F}_k)|} \leq \frac{\sqrt{\mathrm{Var}(\boldsymbol{\zeta}|\mathcal{F}_k)}}{|\mathbb{E}(\boldsymbol{\zeta}|\mathcal{F}_k)|}\sqrt{\mathbb{E}\left(\|\boldsymbol{\phi}_k - \mathbb{E}(\boldsymbol{\phi}_k|\mathcal{F}_k)\|^2\middle|\mathcal{F}_k\right)} < \sqrt{\frac{\gamma}{\delta\rho}}$$

on G_k where the last inequality follows from the fact that $\delta \leq \boldsymbol{\zeta} \leq 1$ and (8.1). From the last equation and (8.3) it follows that

$$\|\boldsymbol{\phi}_k^{(\zeta)} - \boldsymbol{\phi}_k\| < \sqrt{\frac{\gamma}{\rho}}(1 + 1/\sqrt{\delta}) < \epsilon$$

on G_k. Therefore

$$\sum_{s\in S}\theta_0(s)P^s\left(\cap_{\zeta\in Z_\delta}D_{k,\epsilon}\left(P^s, P^{\theta_0^{(\zeta)}}\right)\right) = \mathbb{P}\left(\cap_{\zeta\in Z_\delta}\{\|\boldsymbol{\phi}_k^{(\zeta)} - \boldsymbol{\phi}_k\| < \epsilon\}\right) \geq \mathbb{P}(G_k) > 1-\rho$$

□

9. Proof of Theorem 1

Claim 2. *Let* $\mathbf{X}$ *be an* X*-valued random variable and let* $\nu \in \Delta(X)$ *be the distribution of* $\mathbf{X}$*. Then for every* $r > 0$ *it holds that*

$$\mathbb{P}\big(\nu(B(\mathbf{X}, r)) \geq 1 - Q_\nu(r)\big) \geq 1 - Q_\nu(r)$$

where $Q_\nu(r)$ *is the concentration function of* ν *given by* (3.3) *and* $B(\mathbf{X}, r)$ *is the ball of radius* r *around* $\mathbf{X}$*.*

Proof. Let D be a subset of X such that diameter$(D) \leq d$ and $\nu(D) = 1 - Q_\nu(r)$. Then the event $\mathbf{X} \in D$ implies the event that $D \subseteq B(\mathbf{X}, r)$, which is implies the event $\nu(B(\mathbf{X}, r)) \geq 1 - Q_\nu(r)$ Therefore

$$\mathbb{P}\left(\nu(B(\mathbf{X}, r)) \geq 1 - Q_\nu(r)\right) \geq \mathbb{P}(\mathbf{X} \in D) = \nu(D) = 1 - Q_\nu(r).$$

□

Proof of Theorem 1. Consider the coupling $\mathbf{S}, \mathbf{T}^i, \mathbf{A}^i_k, \mathbf{X}_k, \tilde{\mathbf{X}}_k$ of the real and imagined κ-play paths given in Proposition 1. We prove that players make asymptotically correct predictions in the imagined play-path and then use (7.1) to deduce that they make correct predictions in the real play path.

We first prove that in the imagined game players make correct forecasts, as if they knew the state of nature. We use Proposition 2 where P^s is the joint distribution of $\tilde{\mathbf{X}}_0, \tilde{\mathbf{X}}_1, \dots$ conditioned on $\mathbf{S} = s$ for every state s. Let $K = K(\delta, \epsilon, \rho)$ as in Proposition 2 where $\delta = \min \tau_s(t)$ and the minimum ranges over all states s and all types t such that $\tau_s(t) > 0$. Then it follows from Proposition 2 that

$$\mathbb{P}\left(\left\|\mathbb{P}(\tilde{\mathbf{X}}_k = \cdot|\tilde{T} = t, \tilde{\mathbf{X}}_0, \dots, \tilde{\mathbf{X}}_{k-1}) - \mathbb{P}(\tilde{\mathbf{X}}_k = \cdot|\tilde{\mathbf{S}}, \tilde{\mathbf{X}}_0, \dots, \tilde{\mathbf{X}}_{k-1})\right\| < \epsilon \text{ for every } t \in T\right) > 1-\rho$$

for all days except at most K of them. From the last equation, (3.1) and (6.2) we get that on all good days

$$\mathbb{P}\left(\left\|\tilde{\boldsymbol{\phi}}_k^{(t)} - \chi_{\mathbf{S}, d_{\tilde{\boldsymbol{\Theta}}_k}}\right\| < \epsilon \text{ for every } t \in T\right) > 1 - \rho \tag{9.1}$$

where $\tilde{\boldsymbol{\phi}}_k^{(t)} = \phi(\tilde{\boldsymbol{\Theta}}_k^{(t)}, \tilde{\boldsymbol{\Theta}}_k)$ is the $\Delta(X)$-valued random variable that represents the forecast of a player of type t about period k-outcome computed in the imagined play-path.

From Claim 2 conditioned on $\tilde{\mathbf{X}}_0, \dots, \tilde{\mathbf{X}}_{k-1}$ and (6.2) it follows that

$$\mathbb{P}\left(\chi_{\mathbf{S}, d_{\tilde{\boldsymbol{\Theta}}_k}}(B(\tilde{\mathbf{X}}_k, r)) \geq 1 - Q_\Gamma(r)\right) \geq 1 - Q_\Gamma(d).$$

From the last equation and (9.1) we get that

$$\mathbb{P}\left(\tilde{\boldsymbol{\phi}}_k^{(t)}(B(\tilde{\mathbf{X}}_k, r)) > 1 - Q_\Gamma(r) - \epsilon\right) > 1 - Q_\Gamma(r) - \rho$$

From Theorem 1 it follows that

$$\mathbb{P}\left(\boldsymbol{\phi}_k^{(t)}(B(\mathbf{X}_k, r)) > 1 - Q_\Gamma(r) - \epsilon\right) > 1 - Q_\Gamma(r) - \rho$$

for sufficiently large n, as desired. □

10. Proof of Theorem 2

Proof of Lemma 1. Let $\boldsymbol{\phi}_k^{(t)}$ be the $\Delta(X)$-valued random variable that represents the forecast of a player of type t about period k-outcome computed under the imagined reasoning. On $R(k,r,\epsilon)$ it holds that

$$(1-\epsilon)\left(u(t,a,\mathbf{X}_k)-\omega(r)\right) \le \sum_x u(t,a,x) \;\; \boldsymbol{\phi}_k^{(t)}(x) \le (1-\epsilon)\left(u(t,a,\mathbf{X}_k)+\omega(r)\right)+\epsilon \tag{10.1}$$

for every type t and action a. Therefore on $R(k,r,\epsilon)$ it holds that

$$u(t,b,\mathbf{X}_k) \le \frac{1}{1-\epsilon}\sum_x u(t,b,x) \;\boldsymbol{\phi}_k^{(t)}(x) + \omega(r) \le \\ \frac{1}{1-\epsilon}\sum_x u(t,a,x) \;\boldsymbol{\phi}_k^{(t)}(x) + \omega(r) \le u(t,a,\mathbf{X}_k) + 2\omega(r) + \epsilon/(1-\epsilon)$$

for every $a \in [\kappa(t,\boldsymbol{\Theta}_k)]$ and $b \in A$, where the first inequality follows from (10.1), the second from the equilibrium condition $\sum_x u(t,b,x) \;\boldsymbol{\phi}_k^{(t)}(x) \le \sum_x u(t,a,x) \;\boldsymbol{\phi}_k^{(t)}(x)$ and the third from (10.1). □

Proof of Theorem 2. Let $r=\omega^{-1}(d)$ so that $\omega(r)\le d$. By Theorem 1 there exists an integer K such that under every Markov strategy κ and every $r>0$ all but at most K periods it holds that

$$\mathbb{P}(R(k,r,Q_\Gamma(r)+\epsilon)) > 1-(Q_\Gamma(r)+\rho).$$

By Lemma 1

$$\{R(k,r,Q_\Gamma(r)+\epsilon))\} \subseteq \{H(k,2\omega(r)+(Q_\Gamma(r)-\epsilon)/(1-Q_\Gamma(r)+\epsilon)\} \subseteq \{H(k,2\omega(r)+2Q_\Gamma(r)+2\epsilon)$$

It follows that

$$\mathbb{P}\big(H(k,2\omega(r)+2Q_\Gamma(r)+2\epsilon)\big) > 1-(Q_\Gamma(r)+\rho).$$

as desired. □

References

[1] Nabil I. Al-Najjar and Rann Smorodinsky. Pivotal players and the characterization of influence. *J. Econom. Theory*, 92(2):318–342, 2000.

[2] 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.

[3] George-Marios Angeletos, Christian Hellwig, and Alessandro Pavan. Dynamic global games of regime change: Learning, multiplicity, and the timing of attacks. *Econometrica*, 75(3):711–756, 2007.

[4] Eduardo Azevedo and Eric Budish. Strategyproofness in the large. Technical report, Working paper, 2012.

[5] Yaron Azrieli and Eran Shmaya. Lipschitz games. *Mathematics of operation research*, forthcoming.

[6] Yakov Babichenko. Best-reply dynamics in large binary-choice anonymous games. *Games and Economic Behavior*, 81:130–144, 2013.

[7] Aaron Bodoh-Creed. Efficiency and information aggregation in large uniform-price auctions. *Journal of Economic Theory*, 148(6):2436–2466, 2013.

[8] Guilherme Carmona and Konrad Podczeck. Ex-post stability of bayes–nash equilibria of large games. *Games and Economic Behavior*, 74(1):418–430, 2012.

[9] Joyee Deb and Ehud Kalai. Stability in large bayesian games with heterogeneous players. *Working paper*, 2010.

[10] Timothy Feddersen and Wolfgang Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica: Journal of the Econometric Society*, pages 1029–1058, 1997.

[11] Drew Fudenberg and David K Levine. The theory of learning in games. 1998.

[12] Drew Fudenberg and David K Levine. Self-confirming equilibrium. *Econometrica: Journal of the Econometric Society*, pages 523–545, 1993.

[13] Ronen Gradwohl and Omer Reingold. Fault tolerance in large games. *Games and Economic Behavior*, 2013.

[14] Edward J. Green. Non-cooperative price taking in large dynamic markets. preprint.

[15] Edward J. Green. Continuum and finite-player noncooperative models of competition. *Econometrica*, 52(4):975–993, 1984.

[16] Philippe Jehiel and Frederic Koessler. Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62(2):533–557, 2008.

[17] Ehud Kalai. Large robust games. *Econometrica*, 72(6):1631–1665, 2004.

[18] Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1019–1045, 1993.

[19] Ehud Kalai and Ehud Lehrer. Subjective games and equilibria. *Games and Economic Behavior*, 8(1):123–163, 1995.

[20] Ehud Kalai and Eran Shmaya. Markov equilibrium in big uncertain games. *Working paper*, 2014.

[21] Thodoris Lykouris, Vasilis Syrgkanis, and Eva Tardos. Learning and efficiency in games with dynamic population. *arXiv preprint arXiv:1505.00391*, 2015.

[22] George J. Mailath and Andrew Postlewaite. Asymmetric information bargaining problems with many agents. *The Review of Economic Studies*, 57(3):351–367, 1990.

[23] Eric Maskin and Jean Tirole. Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory*, 100(2):191–219, 2001.

[24] R Preston McAfee. Mechanism design by competing sellers. *Econometrica: Journal of the econometric society*, pages 1281–1312, 1993.

[25] Roger B. Myerson. Population uncertainty and Poisson games. *Internat. J. Game Theory*, 27(3):375–392, 1998.

[26] Roger B. Myerson. Large Poisson games. *J. Econom. Theory*, 94(1):7–45, 2000.

[27] William Novshek and Hugo Sonnenschein. Cournot and Walras equilibrium. *J. Econom. Theory*, 19(2):223–266, 1978.

[28] Wojciech Olszewski and Ron Siegel. Large contests. Technical report, Mimeo, 2013.

[29] 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.

[30] Hamid Sabourian. Anonymous repeated games with a large number of players and random outcomes. *J. Econom. Theory*, 51(1):92–110, 1990.

[31] David Schmeidler. Equilibrium points of nonatomic games. *J. Statist. Phys.*, 7:295–300, 1973.

[32] Sylvain Sorin. Merging, reputation, and repeated games with incomplete information. *Games and Economic Behavior*, 29(1):274–308, 1999.

Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston IL 60208

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

School of Mathematical Sciences, Tel Aviv University and Kellogg School of Management, Northwestern University

E-mail address: e-shmaya@kellogg.northwestern.edu