EX-POST STABILITY IN LARGE GAMES
(DRAFT, COMMENTS WELCOME)

EHUD KALAI

ABSTRACT. The equilibria of strategic games with many semi-anonymous players have a strong ex-post Nash property. Even with perfect hindsight about the realized types and selected actions of all his opponents, no player has an incentive to revise his own chosen action. This is illustrated for normal form and for one shot Baysian games with statistically independent types, provided that a certain continuity condition holds.

Implications of this phenomenon include strong robustness properties of such equilibria and strong purification result for large anonymous games.

1. Introduction and Summary

In a one-shot game with many semi-anonymous players all the equilibria are approximately ex-post Nash. Even with perfect hindsight about the realized types and actions of all his opponents, no player regrets, or has an incentive to unilaterally change his own selected action. This phenomenon holds for normal form games and for one shot Bayesian games with statistically independent types, provided that the payoff functions are continuous. Moreover, the ex-post Nash property is obtained uniformly, simultaneously for all the equilibria of all the games in certain large classes, at an exponential rate in the number of players.

When restricted to normal form games, the above means that with probability close to one, the play of any mixed strategy equilibrium must produce a vector of pure strategies that is an epsilon equilibrium of the game. At an equilibrium of a Bayesian game, the vector of realized pure actions must be an epsilon equilibrium of the complete information game in which the realized vector of player types is common knowledge. Before we elaborate on the implications of this phenomena, it is helpful to view it in the context of the following example.

Simultaneously, each of two players has to choose computer I or M, and independently of each other each has 2/3 probability of being an extrovert and 1/3 probability of being an introvert. The payoff of an extrovert equals one if his choice matches the opponent's, zero otherwise, and for an introvert it is one if he mismatches the opponent's, zero otherwise. This game has two pure strategy

Date: September 21 2001.

Key words and phrases. anonymous games, Nash equilibrium, Bayesian equilibrium, ex-post Nash, purification, interchangeable equilibria, rational expectations equilibrium.

This replaces an earlier version titled "Private Information in Large Games," see Kalai (2000). The author thanks Jeff Ely, Adam Kalai, Andy Postlewaite, Mark Satterthwaite and Asher Wolinsky, for helpful conversations. Partial support for this research was made possible by the National Science Foundation, grant No. SES-9904931.

equilibrium, introverts choose M and extroverts choose I and vice versa, and a continuum of mixed strategy equilibria[1]. But non of these equilibria are ex-post Nash, or even approximately so. Under any one of these equilibria, for example, there is a significant positive probability of ending up with two introverts who chose the same computer, so that each one of then has the incentive to unilaterally change his choice after observing the other.

When the number of players is large, however, all the Nash equilibria are approximately ex-post Nash, even if the game is more complex and highly non symmetric. The n-players simultaneous-move game may have more sophisticated types, and players' payoffs may depend on opponents types in addition to their chosen computers (for example, a player may want to impress some opponents types with his choice). Moreover, different players may have (arbitrary) different payoff functions and different probability distributions by which their types are selected. Regardless of all such specifications, if conditions of semi anonymity and continuity hold, it is almost certain that even with hindsight information about all the players realized types and computer choices, no player would have a significant incentive to unilaterally change his own computer choice.

Going further, consider a family that contains, for every n=1,2,..., many possible n-person games. If the games in the family satisfy semi anonymity and continuity conditions then, uniformly at an exponential rate, all the equilibria of the games in the family become ex-post Nash as we restrict attention to games with increasing number of players.

The condition of semi anonymity requires that every player's payoff function, which naturally may depend on his own type and action, depends on his opponents only through proportions. What matters is the proportion (or the number) of the opponents that fall into the various type and choice categories, and not the identity of the players that make up these proportions[2].

The condition of continuity requires uniform equicontinuity of all payoff functions of games in the family. We discuss the significance of this restriction in the main body of the paper.

The obtained ex-post Nash property, discussed above, has important implications. First, it overcomes potential modelling pitfalls, see for example Cremer and McLean (1985) and Green and Laffont (1987). Consider real life situations in which players may revise earlier choices, for example when players rent, as opposed to buy, computers. A one-shot game model may produce equilibria that are not sustainable in the real-life situation. With hindsight information about the outcome of the one shot game, players who have better choices will simply revise, and the actual final outcome will be different from the theoretically predicted one[3]. Having the ex-post Nash property means that we need not worry about this issue.

[1]For each player each type can mix in any way, provided that in the aggregate over types he is equaly likely to choose any one of the two computers.

[2]As we discuss in the paper, this is somewhat less restrictive than may appear, since information about individual opponents can be incorporated into types.

[3]The standard game theory solution would require to write a more elaborate game that allow for such revisons. This may be simple in some cases, but may be extremly difficult in others. Consider, for example, a two person match pennies game, in which the players are allowed to repeatedly revise their choices during a one hour time period past the initial selection of choices.

But going beyond revision possibilities, as reported in a companion paper, see Kalai (2002), the ex-post Nash property obtained at an exponential rate here implies that the equilibria must become *extensively robust*. They continue to be equilibria even under major modifications of the game. Players order of moves may be determined dynamically, as the game evolves, with information about types and choices partially and discriminatingly revealed at intermediary stages, commitments and revisions may be possible, cheap talk may take place and so on. Regardless of all such modifications, the equilibria of the one shot game remain Nash equilibria in all extensive modifications. This strong robustness property also implies that the equilibria may play a role similar to that of rational expectations equilibria.

We can also view the obtained ex-post Nash phenomenon as a *purification result*. For example, when restricted to normal form games the ex-post Nash property provides stronger conclusions than Schmeidler's (1973) on the role of pure strategy equilibria in large anonymous games. Working in the limit with a continuum of players, Schmeidler shows that large anonymous games have pure strategy Nash equilibria. The current paper shows, asymptotically, that in such games you cannot get away from pure strategy Nash equilibria. The play of any mixed (and of course pure) strategy equilibrium must produce profiles that are pure strategy Nash equilibrium. So every mixed strategy may be thought of as a *purification device*[4]. And extending this to Bayesian games, every equilibrium produces profiles with actions that are pure strategy Nash equilibrium of the complete information game determined by the profile of realized types.

It is also interesting to connect the findings of this paper to learning in repeated Bayesian games. A corollary in Kalai and Lehrer (1993) states that after playing a Bayesian equilibrium of a repeated game for a long time players learn to play optimally, as if they know the (unknown) realized types of their opponents. Or alternatively, after playing for a long time no player has an incentive to deviate from his ex-ante planned strategy, even if he acquires full knowledge of the realized types of his opponents. In the current paper, players play optimally, as if they know the realized types of their opponents, already in the first (and the only) period. So in the sense of Kalai and Lehrer, the players in the large game "learned," or rather already know, the relevant information before they even begin playing.

This connection suggests interesting questions on how the number of players may affect the speed of learning in a repeated Bayesian games with semi-anonymous players. On one hand, the larger the number of players the more there is to be learned about opponents. On the other hand, as indicated by the current paper, the larger the number of players the more predictable their aggregate behavior is. One possible conclusion is that the hardest repeated games of incomplete information to analyze are ones with an intermediate number of players.

Also related to Bayesian learning is the issue of learning the prior probabilities by which types of players are generated. Following a standard Harsanyi setting the results of this paper rely on the fact that these priors are known. But the strength of the conclusion shows how strong this assumption is. Jackson and Kalai (1997) offer a model where priors are learned in a repeated play. Incorporating their approach into a model with many players and learning is an interesting challenge.

[4]The existence of equilibrium is not an issue in the current model because we work with finite games

In proving the main result, we first develop bounds on the level of the ex-post Nash property that is attained at an arbitrary but fixed Bayesian equilibrium in a given game. To do so, we introduce a measure of *strategic interdependence* within a subset of possible outcomes (realized vectors of types and actions). It measures how ranking of alternatives by a player may be affected by changes in the opponents parameters within the set. Under low strategic interdependence, ex-post regret must be low, as long as the play remains within the set. This means that we can bound the level of the ex-post Nash property by identifying high probability sets of outcomes in which the strategic interdependence is low.

Our continuity assumption guarantees that strategic interdependence must be low when a set of outcomes is small. And for the case of many anonymous players, laws of large numbers show that with high probability the game is likely to be played in a small region near the expected play of an equilibrium. Put together, these two phenomena describe the intuition behind the main result.

As the above discussion suggests, being ex-post Nash is really a local property. So while the main theorem requires continuity of the payoff functions on the entire set of outcomes of the game, all that is really needed to make an equilibrium highly ex-post Nash is a large number of anonymous players and continuity near the expected play.

In an appendix, we discuss additional properties on the structure and stability properties of ex-post Nash equilibria

Both, large games and ex-post Nash equilibrium are topics of increasing popularity these days, the recent examples by Al-Najjar and Smorodinsky (2000) and Chung and Ely (2000) offer additional references. An earlier version of this paper, Kalai (2000), has further elaboration on the results reported here, as well as a more substantial discussion of earlier related literature. We should point out that Wooders,Cartwright and Selten (2002) offer, for ”pregames” in a different model, a purification result similar to the one discussed in Kalai (2000) and here.

2. General Definitions and Notations

Two finite universal sets, $\mathcal{T}$ and $\mathcal{A}$, describe respectively *all possible player types* and *all possible player actions* that appear in games discussed in this paper. For notational efficiency, a universal set $\mathcal{C} \equiv \mathcal{T} \times \mathcal{A}$ denotes *all possible player (type-action) compositions.*

A *Bayesian game* is described by a five-tuple (N, T, τ, A, u) as follows.

$N = \{1, 2, ..., n\}$ is the set of *players.*

$T = \times_i T_i$ is the set of *type profiles* (or vectors), with each set $T_i \subseteq \mathcal{T}$ describing the *feasible types of player* i .

$\tau = (\tau_1, \tau_2, ..., \tau_n)$ is the vector of *prior probability distributions*, with $\tau_i(t_i)$ denoting the probability of player i being of type t_i ($\tau_i(t_i) \geq 0$ and $\sum_{t_i} \tau_i(t_i) = 1$).

$A = \times_i A_i$ is the set of *action profiles*, with each set $A_i \subseteq \mathcal{A}$ describing the *feasible actions of player* i.

Let $C_i \equiv T_i \times A_i$ describe the *feasible* (type-action) *compositions of player* i, and $C = \times_i C_i$ denote the set of *feasible composition profiles.* Then, the *players' utility functions* described by the vector $u = (u_1, u_2, ..., u_n)$, assuming a suitable normalization, are of the form $u_i : C \to [0, 1]$.

In addition, standard game theoretic conventions are used throughout the paper. For example, for a vector $x = (x_1, x_2, ..., x_n)$ and an element x'_i, $x_{-i} =$

$(x_1, ..., x_{i-1}, x_{i+1}, ..., x_n)$ and $(x_{i-1} : x'_i) = (x_1, ..., x_{i-1}, x'_i, x_{i+1}, ..., x_n)$. Also, while a composition profile $c = ((t_1, a_1), ..., (t_n, a_n))$, we sometimes describe it as a pair of profiles $c = (t, a)$ in the obvious way.

The Bayesian game is played as follows. In an initial stage, *independently of each other*, every player is selected to be of a certain type according to his prior probability distribution. After being privately informed of his own type, every player proceeds to select an action, possibly with the aid of a randomization devise. Following this, the players are paid, according to their individual utility functions, the payoffs computed at the realized profile of (type-action) compositions.

Accordingly, a *strategy* of player i is defined by a vector $\sigma_i = (\sigma_i(a_i \mid t_i))$ where $\sigma_i(a_i \mid t_i)$ describes the probability of player i choosing the action a_i when he is of type t_i. Together with the prior distribution over his types, a strategy of player i determines an individual distribution over player i's compositions, $\gamma_i(c_i) = \tau_i(t_i) \times \sigma_i(a_i \mid t_i)$. The profile of these distributions, $\gamma = (\gamma_1, \gamma_2, ..., \gamma_n)$, under the independence assumption, determines the overall probability distribution over *outcomes of the game*, namely composition profiles, by $\Pr(c) = \prod_i \gamma_i(c_i)$.

Using expectation and with abuse of notations, the utility functions of the players are extended to vectors of strategies by defining $u_i(\sigma) = E(u_i(c))$. As usual, a vector of strategies σ is a (Bayesian) *Nash equilibrium* if for every player i and every one of his strategies σ'_i, $u_i(\sigma) \geq u_i(\sigma_{-i} : \sigma'_i)$.

A composition profile is ex-post Nash if with full knowledge of the profile (types and selected actions of all players) no player has the incentive to unilaterally change his selected action. Alternatively, the vector of actions described by the profile is a Nash equilibrium of the complete information game determined by the profile of types. A strategy is ex-post Nash if it must lead to outcomes that have the above stability property. Note that ex-post here is in a strong sense, since it requires that the realized pure actions, not the mixed strategies, constitute a Nash equilibrium of the complete information game with all the types being known.

Another way to interpret the property of being ex-post Nash is in terms of decentralization of private information. When playing a strategy profile with this property, a player only has to know his own type in order to select his own action. Any information about the realized types and selected actions of opponents is irrelevant for the purpose of testing the optimality of his own selected action. In other words, for the purpose of acting optimally, private information is important but every player cares only about his own.

For proving asymptotic results a notion of approximate ex-post Nash is needed.

Definition 1. ***Approximate Ex-Post Nash****: Let ε and ρ be positive numbers.*

A composition profile c is ε incentive compatible for player i if for every action a'_i, $u_i(c_{-i} : (t_i, a'_i)) \leq u_i(c) + \varepsilon$.

A composition profile is ε Nash if it is ε incentive compatible for every player.

A strategy profile σ is (ε, ρ) ex-post Nash *if the probability that it yields an ε Nash composition profile is at least $1 - \rho$.*

3. Ex-Post Nash Equilibria in Large Games

In addition to the universal sets of types and actions, $\mathcal{T}$ and $\mathcal{A}$ above, we assume here that there is a universal *set of possible payoff functions $\mathcal{U}$* that consists of functions of the form $g : \mathcal{C} \times \Delta(\mathcal{C}) \longrightarrow [0, 1]$. The interpretation is that the first argument in such a function is the player's own composition and the second argument is the empirical distribution of opponents' compositions.

Moreover, we assume that the collection of functions $\mathcal{U}$ is *uniformly equicontinuous*. Recall that a collection of functions is such if for every positive ε there is a positive δ such that for every two points x, y in the functions common domain and for every function g in the collection $|g(x) - g(y)| < \varepsilon$ whenever the distance between x and y is less then δ. For example, every finite collection of continuous functions defined on the same compact domain is uniformly equicontinuous.

Definition 2. ***Empirical distribution:*** *For every composition profile c define the* empirical distribution *induced by c on the universal set of compositions $\mathcal{C}$ by $emp_c(\kappa) =$*

$$(\textit{the number of coordinates } i \textit{ with } c_i = \kappa) \ /(\textit{the number of coordinates of } c).$$

Definition 3. *The family of* ***semi-anonymous Bayesian games*** $\Gamma = \Gamma(\mathcal{A}, \mathcal{T}, \mathcal{U})$ *consist of all the Bayesian games $(N, \times T_i, \tau, \times A_i, (u_i))$ satisfying $T_i \subseteq \mathcal{T}$, $A_i \subseteq \mathcal{A}$, and where every u_i can be imbedded in some function $g \in \mathcal{U}$ so that $u_i(c) = g(c_i, emp_{c_{-i}})$.*

Theorem 1. ***Ex-post Nash in Large Games:*** *Given a family of semi-anonymous Bayesian games Γ as above and a positive ε, there are positive constants A and B, $B < 1$, with the following property. Simultaneously, all the equilibria of games in Γ with n or more players are (ε, ρ_n) ex-post Nash with $\rho_n = AB^n$.*

Before proceeding further, it may be useful to elaborate on the assumptions of semi-anonymity and continuity.

The condition of semi anonymity is less restrictive than may appear, since it only imposes anonymity within the payoff functions but without further restrictions of symmetry on the players. This means that information about named players can be incorporated into their types. Consider for example a complete information normal form game with n sellers, labeled 1,2,...,n and n buyers, labeled n+1,n+2,...,2n. Suppose that the payoff function of a seller depends on his own strategy and on the empirical distribution of the strategies of the buyers. In violation of the assumption of our model, the payoff function of this seller does not treat all the opponent anonymously, since the buyers, i.e., the players called n+1,...,2n, play a role different from the other players. But if within the buyer group all the players are anonymous for this seller, then we can overcome this problem by describing the situation by a semi anonymous Bayesian game as follows.

Allow each player to have two possible types, a seller or a buyer. Assign probability one of being a seller type to players 1,...,n, and a probability one of being a buyer type to players n+1,...,2n. Now we can write the payoff function of the above seller in the obvious way to depend on the empirical distribution of types (to depend only on the players that are of the buyer type) and actions, without having to specify player labels. Clearly, this description is possible because the model imposes no symmetry restriction on the prior distributions by which types are drawn.

Similar to the above, the model can accommodate many non symmetric games. In addition to playing different roles, as above, players may be identified as belonging to different geographical locations and to different social or professional groups.

The assumption of finitely many types, however, does restrict the generality of such descriptions.

The continuity assumption, when combined with the assumption of semi anonymity, is more restrictive than may appear. Consider for example a game with n players, each having to choose computer I or M. Player 1 is an expert, who is equally likely to be a type who prefers I or a type who prefers M. His payoff is 1 when he chooses the computer that he prefers and 0 otherwise. All the other players are of one possible type that prefers to match the choice of player 1, i.e., they are paid 1 when they match and 0 otherwise. As done above, we can describe this game as a semi-anonymous Bayesian game with three types: an expert who prefers I, an expert who prefers M, and a non expert. (Assign player 1 equal probability of being one of the first two types and to every other player probability one of being of the third type.) In this game, player 1 choosing the computer he likes and every other player randomizing with equal probability between the two computers is an equilibrium of the one shot simultaneous move game.

Unlike the conclusion of the theorem, however, the above equilibrium is not approximately ex-post Nash, no matter how large the number of buyers is. With high probability close to a half of the players would want to revise their choices once they observe the choice of the expert type. The difficulty is that the players payoff functions cannot be imbedded in a uniformly equicontinuous collection of payoff functions g as required by the theorem. As the number of players increases, the percentage of expert types goes to zero. Any function g must specify payoffs for compositions with zero proportions of experts. Yet arbitrarily close to such compositions, there are outcomes with payoff one and outcomes with payoff zero.

4. Bounds on the Level of Ex-Post Nash of a Given Equilibrium

Every equilibrium is (ε, ρ) ex-post Nash for sufficiently large ε and ρ. This section concentrates on a fixed Nash equilibrium of a fixed Bayesian game and develops bounds on the level of ex-post Nash stability it must have. These bounds are used later to prove the main result. As it turns out, for an equilibrium to be highly ex-post Nash, we do not need uniform continuity of the players' payoff functions, continuity in a region of the likely outcomes is sufficient. And even less, the real property that is needed is low strategic interdependence in such a region.

Definition 4. ***Strategic Interdependence:***

The strategic dependence of a player i in a set of composition profiles M, $sd_i(M)$, is defined to be

$$\max |[u_i(c^1_{-i} : (t_i, a'_i)) - u_i(c^1_{-i} : (t_i, a''_i))] - [u_i(c^2_{-i} : (t_i, a'_i)) - u_i(c^2_{-i} : (t_i, a''_i))]|$$

with the maximum taken over all actions a'_i and a''_i, all types t_i, and all composition profiles $c^1, c^2 \epsilon M$ with $t^1_i = t^2_i = t_i$.

The strategic interdependence in M is defined by $si(M) = \max_i sd_i(M)$

The interpretation is that if $sd_i(M)$ is small and the composition profiles are likely to be in M, then the gain to i in a switch from action a'_i to a''_i is almost independent of his type and the compositions of the opponents. If this is the case, any uncertainty about opponents types and actions plays only a little role in his decision of what action to choose.

Lemma 1. *A Bayesian equilibrium is (ε, ρ) ex-post Nash if for a set of composition profiles M*

$$\rho \geq \Pr(M^c) \quad \text{and } \varepsilon \geq si(M) + \max \Pr(M^c \mid c_i)/pr(M \mid c_i)$$

where the maximum is taken over all c_i's that are part of a composition profile $c\epsilon M$.

Proof. Suffices to show that at any $c\epsilon M$, no player can improve his payoff by more than ε by switching from a_i to another action. From the definition of strategic interdependence, if switching at c, from a_i to a_i', improves player i's payoff by r, then the same switch must improve his payoff by at least $r - si(M)$ at any other $\overline{c}\epsilon M$ with $\overline{t}_i = t_i$ (the improvement referred to is the following: fix the opponents' compositions and i's type to be as in $\overline{c}$, and consider the gain to his payoff as he switches from a_i to a_i'). Thus, given his type t_i and his selected action a_i, player i can improve his expected payoff by at least $[r - si(M)] \Pr(M \mid c_i) - \Pr(M^c \mid c_i)$. But since a_i was selected to be an optimal response by i, the last expression must be non-positive, which yields the desired bound. □

The above computational result illustrates that if σ generates a high probability (in the conditional senses just described) low strategic-interdependence set M, then σ is highly ex-post Nash. The following discussion illustrates that under natural restrictions on the game, and assuming a large number of players, such sets M are natural.

Starting with a strategy profile σ the induced vector of measures $\gamma = (\gamma_1, \gamma_2, ..., \gamma_n)$ may be viewed as a vector of extended distributions, each being defined over the (same) universal set of all possible player compositions $\mathcal{C}$ (as opposed to each γ_i being defined on $C_i \subset \mathcal{C}$). Simply, for any possible player composition $\kappa \in \mathcal{C}$, $\gamma_i(\kappa) = \gamma_i(c_i)$ if κ equals some $c_i \in C_i$, and $\gamma_i(\kappa) = 0$ otherwise.

Definition 5. ***Expected distribution****: For a vector of strategy profile σ and the induced distribution $\gamma = (\gamma_1, \gamma_2, ..., \gamma_n)$ define the* expected distribution *on the universal set of compositions $\mathcal{C}$ to be the distribution $\exp_\sigma(\kappa) = \sum_i \gamma_i(\kappa) \ / \ n$*

Recall also the definition of the **Empirical Distribution**: For every composition profile c the *empirical distribution* induced by c on the universal set of compositions $\mathcal{C}$ is $emp_c(\kappa) =$

(the number of coordinates i with $c_i = \kappa$) /(the number of coordinates of c).

For a fixed σ, the empirical distribution, being defined for every randomly selected $c \in C$, may be viewed as a $|\mathcal{C}|$ dimensional random variable. The expected distribution on $\mathcal{C}$, on the other hand, is a fixed vector of $|\mathcal{C}|$ numbers. Coordinatewise, for every $\kappa \in \mathcal{C}$, the expected value of the empirical distribution equals the expected distribution at κ. By laws of large numbers one may expect with high probability to obtain c's whose empirical distribution is close to the expected distribution.

Lemma 2. ***Chernoff-Hoeffding additive bounds:*** *Let $X_1, X_2, ..., X_n$ be a sequence of independent 0-1 random variables with $\Pr(X_i = 1) = \mu_i$. Let $\overline{X} = \sum X_i/n$ and $\overline{\mu} = \sum \mu_i/n$. Then for every $\delta > 0$,*

$$\Pr(\left|\overline{X} - \overline{\mu}\right| > \delta) \leq 2e^{-2\delta^2 n}.$$

Proof. See Theorem A.4 in Alon, Spencer and Erdos (1992) page 235. Apply it once to the variables $X_i - \mu_i$ and once to their negatives. □

Lemma 3. *Let $X_1, X_2, ..., X_n$ be a sequence of independent* 0-1 *random variables with* $\Pr(X_i = 1) = \mu_i$. *Let* $\overline{X} = \sum X_i/n$ *and* $\overline{\mu} = \sum \mu_i/n$. *Then for every* $\delta > 0$ *and every* i,

$$\Pr(\left|\overline{X} - \overline{\mu}\right| > \delta \mid X_i) \leq 2e^{-2[(n\delta-1)/(n-1)]^2(n-1)}.$$

Proof. Let $\overline{X}_{-i} = \sum_{j\neq i} X_j/(n-1)$ and $\overline{\mu}_{-i} = \sum_{j\neq i} \mu_j/(n-1)$.

$\left|\overline{X} - \overline{\mu}\right| > \delta$ implies that $\left|\overline{X}_{-i} - \overline{\mu}_{-i}\right| > \delta n/(n-1) - 1/(n-1) = (n\delta-1)/(n-1)$. The lemma follows by applying the previous lemma to $\overline{X}_{-i}$. □

Define the δ *neighborhood* of $\exp_\sigma$ by

$$nbd(\exp_\sigma, \delta) = \{c \in C : \max_\kappa |emp_c(\kappa) - \exp_\sigma(\kappa)| \leq \delta\}.$$

Lemma 4. *For any* $\delta > 0$, $\Pr[c \notin nbd(\exp_\sigma, \delta)] \leq 2\,|\mathcal{C}|\, e^{-2\delta^2 n}$.

So for any given small δ and the fixed cardinality $|\mathcal{C}|$ of the universal set of compositions, if the number of players is large there is a high probability of the empirical distribution of compositions being uniformly close to the expected distribution of compositions. Moreover, the same holds true for the conditional probabilities.

Lemma 5. *For any* $\delta > 0$, $\Pr[c \notin nbd(\exp_\sigma, \delta)|c_i] \leq 2\,|\mathcal{C}|\, e^{-2[(n\delta-1)/(n-1)]^2(n-1)}$.

Applying the previous general bounds on the level of ex-post Nash to $M = nbd(\exp_\sigma, \delta)$, one obtains the following bounds.

Theorem 2. ***Bounds on the Level of Ex-Post Nash****: For any* $\delta > 0$ *a Bayesian equilibrium* σ *is* (ε, ρ) *ex-post Nash if*

$$\begin{aligned} \varepsilon &> si[nbd(\exp_\sigma, \delta)] \\ &\quad +2\,|\mathcal{C}|\, e^{-2[(n\delta-1)/(n-1)]^2(n-1)}/[1 - 2\,|\mathcal{C}|\, e^{-2[(n\delta-1)/(n-1)]^2(n-1)}], \text{ and} \\ \rho &> 2\,|\mathcal{C}|\, e^{-2\delta^2 n}. \end{aligned}$$

5. Proof of the Main Result

We fix the family of semi-anonymous games Γ, with a collection of uniformly equicontinuous payoff functions $\mathcal{U}$ as in the statement of the theorem. It is sufficient to prove the theorem for some positive integer m. That is, given the family and a positive ε there is an m and constants A and B, with the conclusion of the theorem holding for all $n \geq m$. This is so because once you prove it for some m you can simply increase A to an A', so that the conclusion, with A' and B, is trivially satisfied for all values of $n \leq m$ (make $A'B^m > 1$), and the statement for $n > m$ is unchanged.

Lemma 6. ***uniform equicontinuity implies low local strategic interdependence****: For any positive* ε *there is a positive* δ *with the following property. For every game in the family and for every strategy profile* σ, $\varepsilon > si[nbd(\exp_\sigma, \delta)]$.

Proof. Recall that $si[nbd(\exp_\sigma, \delta)]$ is defined to be the maximum of the expression below, when you consider all players i, all pairs of actions a_i' and a_i'', all types t_i, and all pairs of composition profiles $c_1, c_2 \epsilon nbd(\exp_\sigma, \delta)$ having $t_1 = t_2 = t_i$:

$$|[u_i(c^1_{-i} : (t_i, a_i')) - u_i(c^1_{-i} : (t_i, a_i''))] - [u_i(c^2_{-i} : (t_i, a_i')) - u_i(c^2_{-i} : (t_i, a_i''))]|$$

But rearranging terms this expression equals

$$|[u_i(c^1_{-i} : (t_i, a_i')) - u_i(c^2_{-i} : (t_i, a_i'))] - [u_i(c^1_{-i} : (t_i, a_i'')) - u_i(c^2_{-i} : (t_i, a_i''))]|$$

This last expression can be made arbitrarily small by making each of its two terms arbitrarily small. So it suffices to show that expressions of the form $u_i(c^1_{-i} : (t_i, a'_i)) - u_i(c^2_{-i} : (t_i, a'_i))$ can be made arbitrarily small by restricting attention to c's in $nbd(\exp_\sigma, \delta)$. However, the equicontinuity assures us that by making δ sufficiently small, we can make these expressions small, simultaneously, for all the strategy profiles of all the games in Γ. $\square$

Back to the proof of the theorem, recall from the previous section, that an equilibrium σ is (ε, ρ) ex-post Nash if for some $\delta > 0$ the following two inequalities are satisfied.

$$\varepsilon > si[nbd(\exp_\sigma, \delta)] + 2\,|\mathcal{C}|\, e^{-2[(n\delta-1)/(n-1)]^2(n-1)}/[1 - 2\,|\mathcal{C}|\, e^{-2[(n\delta-1)/(n-1)]^2(n-1)}]$$

and

$$\rho > 2\,|\mathcal{C}|\, e^{-2\delta^2 n}.$$

Using the lemma above, we can choose a positive δ and an m sufficiently large so that the top inequality holds simultaneously for all strategy profiles σ of all the games with m or more players. Simply choose δ to make the first expression in the right hand side smaller than $\varepsilon/2$ and m sufficiently large to make the second expression smaller than $\varepsilon/2$.

The proof is now completed by setting $A = 2\,|\mathcal{C}|$ and $B = e^{-2\delta^2}$.

6. References.

Al-Najjar, N., and R. Smorodinsky (2000): "Large Non-Anonymous Repeated Games," *Games and Economic Behavior*, forthcoming

Alon, N., J.H. Spencer, and P. Erdos (1992): *The Probabilistic Method*, John Wiley & Sons, New York.

Chung, K.S, and J.C. Ely (2000): "Efficient and Dominance Solvable Auctions with Interdependent Valuations," Northwestern University working paper.

Cremer, J. and R.P. McLean (1985): "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," Econometrica, 53, 2.

Green, J.R., and J.J. Laffont (1987): "Posterior Implementability in a Two-person Decision Problem," *Econometrica*, 55 (1), pp. 69-94.

Jackson, M.O., and E. Kalai (1997): " Social Learning in Recurring Games," *Games and Economic Behavior*, 21, No. 1/2, 102-134.

Kalai, E., (2000), "Private Information in Large Games," The Center for Mathematical Studies in Economics and Managemental Science, Northwestern University, DP #1312

Kalai, E., (2002), "Large Robust Games," The Center for Mathematical Studies in Economics and Management Science, Northwestern University, DP #

Kalai, E., and E. Lehrer (1993): "Rational Learning Leads to Nash Equilibrium," *Econometrica* 61 (5), pp 1019-45.

Luce, D.L. and H. Raiffa (1957): *Games and Decisions*, John Wiley and Sons Inc. pp 106-108.

Minehart, D., and S. Scotchmer (1999): "Ex Post Regret and the Decentralized Sharing of Information," *Games and Economic Behavior*, 27, pp 114-131

Schmeidler, D., (1973): "Equilibrium Points of Nonatomic Games," *Journal of Statistical Physics* 17 (4), pp 295-300.

Wooders, M., E. Cartwright and R. Selten, (2002) "Social Conformity and Equilibrium in Pure Strategies in Games with Many Players," Warwick Discussion Paper.

7. Appendix: Further Properties of Ex-Post Nash Equilibria

We first discuss some properties of (full) ex-post Nash equilibrium, before commenting on the possible extensions of these properties to approximate ex-post Nash equilibrium.

7.1. **Structure.**

	.34	.60	.06	0	0
.27	**6, 7**	**5, 7**	**7, 7**	1, 2	1, 3
.38	**6, 5**	**5, 5**	**7, 5**	9, 2	5, 1
.35	**6, 8**	**5, 8**	**7, 8**	5, 4	4, 3
0	2, 2	2, 4	0, 1	9, 8	7, 7
0	3, 7	1, 9	1, 9	1, 1	9, 9

As already stated, in normal form games every pure strategy equilibrium is ex-post Nash. The payoff table above offers a typical example of ex-post Nash equilibrium in mixed strategies. Since every one of the nine balded entries is a pure strategy Nash equilibrium, and since these are the support of the indicated mixed strategy profile, it is clear that the mixed strategy equilibrium is ex-post Nash.

Several additional properties must hold. Recall first the notion of Nash interchangeability for strategies of normal form games. Two equilibria $a = (a_1, ..., a_n)$ and $b = (b_1, ..., b_n)$ are *interchangeable*, see Luce and Raiffa (1957), if every strategy profile $m = (m_1, ..., m_n)$ which is a coordinatewise selection from a and b, i.e. every $m_i = a_i$ or b_i, is also a Nash equilibrium. This property guarantees that for the sake of choosing his own best reply a player is not concerned with which equilibrium a or b is the one being played. A best reply in one is automatically a best reply to the other and to any coordinatewise selection from the two. Since the support of a mixed strategy equilibrium has a product structure it must be that a profile of strategies is an ex-post Nash equilibrium if and only if its support consists of interchangeable Nash equilibria.

This observation generalizes to Bayesian equilibria with independent types. Consider a set of Nash composition profiles S, and recall that every one of its elements $c = (c_1, ..., c_n)$ can be viewed as a pair of ordered profiles (t, a) describing respectively the types and actions of the n players. Define the elements of S to be *interchangeable* if every coordinatewise selection of such elements m, i.e., every $m_i = c_i$ for some $c \in S$, is Nash (recall this means that $m = (t, a)$ with a being a Nash equilibrium of the complete information game induced by t). Similar to the case of normal form games, interchangeability in a Bayesian game means that a player is not concerned with which outcome in the set S prevails. If the player is of type t_i and chooses a_i as best response to some opponents' composition profile in an interchangeable S, then this choice is automatically best responding to any other opponents' composition profile from S.

Proposition 1. ***Interchangeability***. *A strategy profile of a Bayesian game is ex-post Nash if and only if its support consists of Nash composition profiles that are interchangeable.*

Proof. This follows from the fact that both the strategies and the prior over types have a product structure (the latter under the independence assumption made throughout this paper). □

A mixed strategy ex-post Nash equilibrium may also be thought of as having weak version of strategic dominance property. Note that any one of the strategies in the support of the row chooser above, i.e. rows 1-3, is weakly dominant if you restrict the column chooser to strategies in her support, i.e., columns 1-3. Similar dominance holds for strategies of the column chooser when you restrict the row chooser to his equilibrium strategies. This is also true at Bayesian equilibrium.

Proposition 2. ***Restricted Dominance.*** *For every $c_i = (t_i, a_i)$ which is a coordinate of a composition profile c in the support of an ex-post Nash equilibrium, a_i must be a dominant strategy for type t_i, as one restrict the outcomes of the game to have any c'_{-i} from the support of the other players' equilibrium distribution.*

The structure described above shows that being ex-post Nash implies ex-ante Nash.

Proposition 3. ***Ex-Post implies Ex-Ante Nash.*** *If a strategy profile is ex-post Nash it must be a Nash equilibrium of the game.*

Proof. A Bayesian equilibrium may be characterized by having every player i using actions a_i that have highest expected value relative to every one of his realized types t_i. The structural propositions above guarantee that player i does so at every strategy profile that is ex-post Nash. □

7.2. **Robustness and Stability.** While ex-post Nash means stability with perfect hindsight knowledge, it implies stability with partial hindsight knowledge. Formally, let a signalling function be any $g : C \to M$, where C is the set of composition profiles and M is an abstract set of messages. The Bayesian equilibrium σ is (i, g) ex-post Nash if for every positive probability composition profile c, conditioning on the information $g(c)$ with the knowledge of his own type t_i and selected action a_i, player i cannot improve his expected utility by switching to a different action a'_i.

Proposition 4. ***Invariance to Ex-Post Partial Information:*** *An ex-post Nash equilibria is (i, g) ex-post Nash for every player i and every signalling function g.*

Proof. Given any outcome c and an action a'_i different from a_i, consider all the positive probability outcomes $\bar{c}$ that are compatible with player i's information t_i, a_i, and $g(c)$, i.e., the outcomes that i cannot differentiate among with the knowledge that he has at c. Being ex-post Nash implies that switching at any such $\bar{c}$ from a_i to a'_i can only decrease player i's payoff. Thus his conditional expected gain from the switch could only be negative. □

Ex-post Nash equilibria are stable under changes in the prior distributions over types, and in the probabilities of mixed strategies of the opponents. Formally, given a Bayesian game $G = (N, T, \tau, A, u)$, a game $G' = (N, T, \tau', A, u)$ has *modified priors that admit no new types* if for every type profile t, $\tau'(t) > 0$ implies that $\tau(t) > 0$.

Proposition 5. ***Invariance to Prior Type Probabilities:*** *If σ is an ex-post Nash equilibrium in a game G then it is also an ex-post Nash equilibrium in every game G' with modified priors that admit no new types.*

Proof. This follows immediately from the fact that any positive probability composition profile generated by σ in G' has positive probability of being generated in G, and thus it is a Nash profile. $\square$

For exactly the same reasons the following is also true. Given any equilibrium σ of G, a *modified strategy profile that admits no new actions* is a strategy profile σ', with the property that for every player i and type t_i, $\sigma'_i(a_i|t_i) > 0$ implies that $\sigma_i(a_i|t_i) > 0$.

Proposition 6. ***Invariance to Mixing Probabilities:*** *If an equilibrium is ex-post Nash then every modified strategy profile that admits no new actions is also ex-post Nash.*

The above two propositions illustrate that in playing or in analyzing a simultaneous move Bayesian game, at an ex-post Nash equilibrium one should only be concerned with the type profiles that may be generated by nature, and not with their probabilities, and with the actions that may be selected by the players, but not with the probabilities used to select them.

Turning our attention to approximate ex-post Nash, some of the above properties should generalize to their approximate versions in a straight forward fashion. These include interchangeability, restricted dominance, ex-post implies ex-ante, and invariance with partial information. The invariance to prior probability and to mixing probabilities does not hold for approximate ex-post Nash. But it does have implications for continuity properties of such equilibria.

Kellogg School of Management, Northwestern University, Evanston IL 60208
E-mail address: `kalai@kellogg.northwestern.edu`