LARGE ROBUST GAMES
(DRAFT, COMMENTS WELCOME)

EHUD KALAI

Abstract. A major modeling difficulty of non-cooperative game theory is the sensitivity of Nash equilibria to details that are not defined by the real life situation they describe. The difficulty is less severe in games with many semi-anonymous players. All the equilibria of such games are extensively robust, they are immune to changes in the order of play, information transmission, cheap talk, commitments, revision possibilities and more. This is illustrated for normal form games and for one-shot Bayesian games with statistically independent types, subject to a suitable continuity condition on the payoff functions.

1. Introduction and Summary

A difficulty with the use of Nash equilibrium as a modeling tool is its sensitivity to the rules of the game, e.g. the order of players moves and the information they have when they move. Since these details are often not defined by the real life situation being modeled, the prediction of the equilibrium is often unreliable. This paper illustrates that this difficulty is less severe in general classes of games that involve many semi-anonymous players. In normal form games and in one-shot Bayesian games with independent types all the equilibria become *extensively robust* as the number of players increases, provided that a certain continuity condition holds.

For this purpose, we define an equilibrium of a game to be extensively robust, if it remains an equilibrium in *all extensive versions* of the game. Such versions allow for wide flexibility in the order of players moves, information leakage, commitment and revision possibilities, cheap talk, and more. The robustness property is obtained, uniformly at an exponential rate in the number of players, for all the equilibria in general classes of one-shot games with the properties mentioned above.

In addition to interest in the phenomenon for its own sake, the robustness property may have direct positive implications to areas where game theory is applied. A mechanism designer, who succeeds in implementing a socially efficient outcome through a Nash equilibrium of a one-shot simultaneous-move game, does not have to be concerned that the players may play a different extensive version of his game, see for example Green and Laffont (1987). So even if players decide to act sequentially and to share information prior to making their moves, or to go back and

---

Date: March 12, 2002.

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

The author thanks Eddie Dekel, Avinash Dixit, Jeff Ely, Olivier Gossner and Sergiu Hart for helpful conversations. Partial support for this research was made possible by the National Science Foundation, grant No. SES-9904931.

revise choices after seeing the outcome of the implementation, the equilibrium that he constructed remains viable. In various social aggregation methods, extensive robustness means that the outcome of a vote is immune to institutional changes, and public poles should not alter the outcome of the equilibrium. And below, we discuss examples that show how extensively robust equilibrium may be a useful concept, similar to rational expectations equilibrium, for more robust modelling of behavior in market games.

Studies of large games is not a new topic in game theory and economics. The book of Aumann and Shapley (1974) surveys years of research on large cooperative games. The first study of large non-cooperative anonymous games, Schmeidler (1973), deals with existence of pure strategy Nash equilibria in such normal form games. More recently there have been many studies of specific economic games, see for example Mailath-Postlewaite (1990) on bargaining, Rustichini-Satterthwaite-Williams (1994) and Pesendorfer-Swinkels (1997) on auctions, and Feddersen-Pesendorfer (1997) on voting. Many of these papers concentrate more on issues of economic efficiency, and less on robustness.

When addressing robustness issues, previous studies imposed a weaker condition, known as ex-post Nash[1]. Applications of this idea to specific economic problems include Cremer-McLean (1985), Green-Laffont (1987) and Minehart-Scotchmer (1999). A general result illustrating that the ex-post Nash property is obtained for large Bayesian games is described in Kalai (2000,2002). Indeed, the proof of the main result of the current paper makes use of this weaker result and, very importantly, of the fact that it is obtained at an exponential rate as the number of players increases.

A connection of extensively robust equilibrium to rational expectations equilibrium is discussed below. We refer the reader to the survey of Jordan-Radner (1982) for a general discussion, and to Forge-Minelli (1997,1998) and Minelli-Polemarchakis (1997) for additional earlier studies relating rational expectations to game theory.

1.1. **Example: a Game of Evolving Standards.** Simultaneously, each of n players has to choose computer I or computer M, and, independently of the opponents, each is equally likely be a type who likes I or a type who likes M. Most of a player's payoff comes from matching the choice of the opponents but there is also a small payoff in choosing the computer he likes. Specifically, each player's payoff function is: 0.1 if he chooses his favorite computer (zero otherwise) plus 0.9 times the proportion of opponents his choice matches. Assuming that each player knows only his own realized type before making the choice, the following three strategy profiles are Nash equilibria of the simultaneous move game: the constant strategies, with all the players choosing I or with all the players choosing M, and the one with all the players choosing their favorite computers

The constant strategies are robust, no matter what the size of the population is. For example, if the choices are made sequentially in a known order, with every player knowing the choices of his predecessors, then everybody choosing I regardless of the observed history is a Nash (even if not subgame perfect) equilibrium of the extensive game. This is not the case for the choosing-your-favorite-computer strategies. For example if the population consists of only two players, there are positive probability

[1]We ignore other economic and (cooperative) game theoretic notions of robustness, see for example Hansen and Sargent (2001) and Kovalenkov and Wooders (2001).

histories after which the follower is better off matching his predecessor rather than choosing his favorite computer.

As users of game theory know, this sensitivity to the order of moves creates modelling difficulties, since we do not know in what order players buy (or rent) computers. Also the real life situation may allow for other possibilities, for example the players may repeatedly revise earlier choices after observing opponents choices, players may make binding commitments, may make announcements etc., and every such version may drastically change the equilibria of the game.

But the modeling difficulty becomes less severe if the number of players is large. Now, even choosing-their-favorite-computers is a highly robust equilibrium, it remains an approximate Nash equilibrium in the sequential game and in all other extensive versions of the game. These versions accommodate modifications of the following types.

The order of moves may be decided, deterministically or stochastically, based on the history of play, and information about earlier choices may be partially and differentially revealed as the game progresses. Multiple opportunities to revise choices may be available to players. Player may commit to choices and reveal such commitments to selected opponents. Players may commit to not observe earlier choices. Cheap talk may take place during the play of the game. Regardless of all such modifications, if the number of players is large, choosing-their-favorite-computers remains an approximate Nash equilibrium in the extensive version of the game.

Moreover, the above robustness property is not restricted to the equilibrium of choosing-their-favorite-computers. Every Nash equilibrium of the one shot game is extensively robust, and this is true even if the original computer selection game that we start with is more complex and highly non symmetric. There may be any finite number of computer choices and player types, and different players may have different (arbitrary) payoff functions and different prior probability distributions by which their types are drawn. Players' payoffs may take into consideration opponent types, and not just opponent actions (e.g., a player may have some utility in impressing some opponent types with his chosen computer). Regardless of such specifications, all the equilibria of the one shot game are approximately robust when n is large.

The robustness of all one shot equilibria described above is not restricted to computer choice or other market games. Any game with payoff functions that depend on aggregate data of the opponents in a continuous fashion has this property. So various social aggregation games, games of joining clubs, political parties or other groups, location games, transportation and congestion games, may all fit the description. And for a given collection that consists of many such games, all the equilibria of all the games becomes extensively robust, uniformly at an exponential rate, as the number of players increases.

We should emphasize, however, that the equilibria obtained in the extensive versions above are approximately Nash, and not approximately subgame perfect Nash. This is unavoidable. It is interesting to note though, as done by an example later in the paper, that the level of violating subgame perfection decreases as the number of players increases. We should also note that the equilibrium in the extensive versions is approximately Nash in a strong sense. There is a high probability of following a complete play path along which *none of the many players*, no matter how long the play path is, ever has a significant incentive to deviate from the equilibrium strategy.

1.2. **Connection to Rational Expectations Equilibria.** Their strong robustness property implies that the equilibria of the large games above may play a role similar to rational expectations equilibria. Consider the example of the equilibrium above, where players buy their favorite computers. Assume that the real game of buying computers, which may not be played simultaneously, will be played in a relatively short period of time, so that players do no time discounting for their forthcoming expenses and payoffs.

At a rational expectations equilibrium, the players strategies, in addition to depending on their private information, depend on the market prices of the computers, which in turn depend on the players computer choices. From a game theoretic perspective, it seems that such a rational expectations equilibrium mixes together ex-ante information (players' types) with ex-post results (prices, which depend on players choices). Or put differently, it is a fixed point of a "bigger system" in which ex-ante and ex-post information is considered simultaneously

Game theory, on the other hand, differentiates between ex-ante information, expressed as variables in the domain of individual strategies, and ex-post information, which is available at too late a stage for any player to respond to. For example, if in the above one-shot game of buying computers we introduced prices, as a function of total demand for the two computers, the prices will be available only as ex-post information, and players could not react them. This would make the game theoretic Bayesian equilibrium less attractive as a tool for modeling the buying of computers.

But in the case of the large games described in this paper such discrepencies disappear. The extensive robustness property implies that partial information, such as evolving computer prices revealed at any stage of the game (even if not played simultaneously), gives no player an incentive to change his computer choice. So the equilibrium possesses the properties of rational expectations equilibrium, even in a strong extensive sense.

This is an attractive property, because the game theoretic model, unlike the rational expectations equilibrium that is only looking for a fixed point, allows us to make explicite what is ex-post and ex-ante at every stage of the game and allows for all the extensive modifications described earlier. The fact that we obtain an extensively robust generation of rational expectations equilibrium can therefore be viewed as non-cooperative game theoretic foundation for rational expectations equilibrium[2] (in a parallel way to cooperative equivalence theorems, where the core and Shapley value are used as cooperative game-theoretic foundations for competitive equilibrium).

Looking more directly at the phenomenon, due to laws of large numbers, at the above equilibrium prices are anticipated correctly by the players, who know the prior probability over types. (this would also be true for the sellers, if we add them to the model). And more explicitly than at rational expectations equilibria, in the Nash equilibria above players have no incentives to deviate from their strategy even during the formation of the equilibria, as computer purchases and other information is being observed.

[2]The author thanks Avinash Dixit for making this observation.

As the last example shows, the assumption of known priors is strong, as it leads to correct anticipation of prices in equilibria of games with many players. Interesting questions about learning with many players in games with unknown priors are left for future research.

2. General Definitions and Notations

2.1. **The Bayesian Game.** 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 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 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.*

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 will be used. 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 $(x_{i-1} : x'_i) = (x_1, ..., x_{i-1}, x'_i, x_{i+1}, ..., x_n)$.

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* for 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, $u_i(\sigma) \geq u_i(\sigma_{-i} : \sigma'_i)$.

2.2. **Extensive Versions of the Game and Strategies.** In the sequel, we define an equilibrium of a Bayesian game G to be robust, if it remains an equilibrium in all extensive versions of the game. A simple example of an extensive version of a game G is one with *one round of revisions.* In the first round of this two round game, the original game G is played. In a second, the information about the realized types and selected pure actions of all players becomes common knowledge, and then, simultaneously, every player has the one time opportunity to revise his first-round choice. Naturally, one can imagine situations where the pre-revision information is only partial, and differentially so across players.

A second example of an extensive version is when the game is played sequentially, rather than simultaneously, with later movers receiving full or partial information about the history of the game. The order may depend on the history of the game.

Combining changes in the order of play and multiple rounds of revisions already permits the construction of many interesting extensive versions of a game. But many more modifications are possible. For example, players may determine whether their choices become known and to what other players, various commitment possibilities may be available, cheap talk announcements may be made and players may have control over the rules of the game that follows. Basically, we would like an extensive version of a simultaneous-move game G to be any extensive game in which every play path ends up generating a composition profile for the simultaneous move game.

The following abstract definition of an extensive version of a game accommodates the modifications discussed above and many more. Starting with the given simultaneous move Bayesian game $G = (N, T, \tau, A, u)$ and any auxiliary finite set of moves M, define a *version* of the game G to be a finite perfect-recall Kuhn type extensive form game $\overline{G}$ with the following structure and payoffs.

The initial node in the game tree belongs to nature with the outgoing arcs being labeled by the elements of $T \times M$. Thus, at the initial stage as in the original game, nature chooses a type for each player. But in addition, it chooses an abstract move that may have different consequences to different players during the play of the game. Any probabilities may be assigned to these arcs as long as the marginal distribution over the set of type profiles T coincides with the prior probability distribution τ of the underlying game G.

Every other node in the game tree belongs to one of the players i, and the arcs coming out of it are labeled by the elements of $A_i \times M'$, for some non-empty subset of the set of moves M'. The interpretation is that when a player is called upon to act he has to make an initial choice, or a revision earlier choices, from his set of feasible actions in the underlying game, as well as a move that may affect the continuation of the game.

At **every information set** the active player i has, at a minimum, complete knowledge of his own type t_i (all the paths that visit this information set start with nature selecting t's with the same type t_i).

Every play path in the tree visits at least once one of the information sets of every player. This guarantees that a player in the game chooses an action, possibly without any revisions, at least once.

The resulting composition profile associated with a complete path in the game tree is $c = (t, a)$ with t being the profile of initial types selected by nature, and with each a_i being the last action, in the play path, taken by player i. The last action is the one that matters because multiple choices by the same player represent revisions of earlier choices.

The players' payoffs at a complete play path are defined to be their payoffs, from the underlying game G , at the resulting composition profile of the path.

Remark 1. *The inclusion of abstract moves by nature at the beginning of the tree significantly extends the set of possible versions. For one thing, it means that excluding nature from having additional nodes later in the tree involves no loss of generality (one can move all the random choices to the beginning, to be included* *as a part of a large initial nature's move that will be partially and differentially revealed to the players at the appropriate places in the tree). But it also means that the version that is being played may be random, reflecting, for example, possible uncertainty about the real life version in the mind of the modeler.*

Similarly, a greater generality is obtained by including abstract moves, in addition to selected action, at the nodes of the players. For example, using these we can model a player's choice to reveal information, to seek information, to make announcements, affect the rules regarding the future progression of the game, etc.

Definition 1. ***Extensive Versions of a Strategy***. *Given an individual strategy σ_i in a Bayesian game G and a version $\overline{G}$, an* extensive version *of σ_i is any strategy $\overline{\sigma}_i$ in $\overline{G}$ that initially chooses actions with the same probabilities as σ_i and does not modify earlier choices in all subsequent information sets. Formally, at any initial information set of player i the marginal probability that $\overline{\sigma}_i$ selects the action a_i is $\sigma_i(a_i \mid t_i)$, where t_i is the type of player i at the information set. In every non initial information set of player i $\overline{\sigma}_i$ selects, with certainty, the same action that was selected by player i in his previous information set (this is well defined under the perfect recall assumption).*

An extensive version of a profile of strategies $\sigma = (\sigma_1, ..., \sigma_n)$ is a profile $\overline{\sigma} = (\overline{\sigma}_1, ..., \overline{\sigma}_n)$ with all the individual strategies $\overline{\sigma}_i$ being defined in the same extensive version $\overline{G}$ of the game G.

2.3. **Extensively Robust Equilibrium.**

A Nash equilibrium σ in a Bayesian game G is extensively *robust* if every profile of its extensive versions $\overline{\sigma} = (\overline{\sigma}_1, ..., \overline{\sigma}_2)$ is a Nash equilibrium in the extensive game $\overline{G}$ in which they are defined. But we need to define a notion of being approximately robust, where $\overline{\sigma}$ is only required to be an ε Nash equilibrium in $\overline{G}$ with high probability. This notion assures that the incentives of any player to unilaterally deviate at any positive probability information set are small, even when such a deviation is coordinated with further deviations in his later information sets.

Fix a version $\overline{G}$, a vector of behavioral strategies $\overline{\eta} = (\overline{\eta}_1, ..., \overline{\eta}_2)$ and a player i. The payoff of player i is defined to be the usual expected value, $E_{\overline{\eta}}(u_i)$, computed according to the distribution on play paths (outcomes) induced by $\overline{\eta}$.

Given an information set of player i, A, a *modification of player i strategy at A* is any strategy $\overline{\eta}'_i$ with the property that at every information set of player i, B, which is not a follower[3] of the information set A, $\overline{\eta}'_i$ coincides with $\overline{\eta}_i$. Player i can *unilaterally improve his payoff by more than ε at the information set A* if $E_{\overline{\eta}'|A}(u_i) - E_{\overline{\eta}|A}(u_i) > \varepsilon$, where $\overline{\eta}' = (\overline{\eta}_1, \overline{\eta}_2, ..., \overline{\eta}'_i, ..., \overline{\eta}_n)$ for some $\overline{\eta}'_i$ that modifies $\overline{\eta}_i$ at A.

Note that such ε unilateral improvements are only defined at positive probability information sets, and that the event player i has a better than ε improvement at some information set is well defined, by simply considering the play paths that visit such information sets. Similarly, the event that *some* player has a better than ε improvement at some information set is well defined, since it is the union of all such individual events.

[3]Recall that by Kuhn's perfect recall condition, every node in B follows some node in A or no node in B follows some node in A.

Definition 2. ***approximate Nash equilibrium.*** *A strategy profile $\overline{\eta}$ of $\overline{G}$ is an (ε, ρ) Nash equilibrium, if the probability that some player has a better than ε improvement at some information set is not greater than ρ.*

Definition 3. ***approximate robustness.*** *An equilibrium of G, σ, is (ε, ρ) robust, if in every extensive version $\overline{G}$ every profile of extensive versions $\overline{\sigma}$ is an (ε, ρ) Nash equilibrium.*

An equilibrium is (ε, ρ) ex-post Nash *if it is an (ε, ρ) Nash equilibrium in the version with one round of revisions discussed above.*

Clearly being ex-post Nash is a weak form of being robust, yet it is the "strongest"[4] form of being ex-post Nash, since the players know everything, types and selected pure actions, before deciding on revisions.

Remark 2. *One can check that in any complete information normal form game every pure strategy Nash equilibrium is ex-post Nash and even robust. This is no longer the case for incomplete information games, as can be seen by the example in the introduction. In the two player game of choosing computers, the pure strategy of choosing-your-favorite-computer is clearly not robust.*

The notion of being (ε, ρ) ex-post Nash is not monotone in the information that is revealed prior to the possible revision. Even in a single person decision making, a player strategy may be, with high probability, ε optimal after receiving complete information about the state of the world, but may no longer be so under partial information. Consider a parent who decides not to insure the car of his child. His probability of having serious regret, due to learning that there was an accident, is small. But the probability of having (somewhat less but still) serious regret due to partial information, e.g., the child is late coming home at night, is high.

3. Extensive Robustness in Large Games

3.1. The Main Result.

Two finite universal sets, $\mathcal{T}$ and $\mathcal{A}$, describe respectively *all possible player types* and *all possible player actions* that may appear in the games discussed in this sections. The set $\mathcal{C} \equiv \mathcal{T} \times \mathcal{A}$ denotes *all possible player (type-action) compositions.* A set of *possible payoff functions $\mathcal{U}$* consists of functions of the form $g : \mathcal{C} \times \Delta(\mathcal{C}) \longrightarrow [0, 1]$, where the first argument describes a player's own composition and the second argument describes the empirical distribution of opponents' compositions.

Definition 4. ***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 5. *The family of* ***semi-anonymous Bayesian games*** $\Gamma = \Gamma(\mathcal{A}, \mathcal{T}, \mathcal{U})$ *consists of all the Bayesian games $(N, \times T_i, \tau, \times A_i, (u_i))$ satisfying $T_i \subseteq \mathcal{T}$ and $A_i \subseteq \mathcal{A}$, and where every u_i can be imbedded in some function $g \in \mathcal{U}$. More specifically, let $C_i = T_i \times A_i$ and $C = \times C_i$, then for every payoff function u_i there is a function $g \in \mathcal{U}$ such that for every $c \in C$, $u_i(c) = g(c_i, emp_{c_{-i}})$.*

Recall that a collection of common-domain functions is *uniformly equicontinuous* if for every positive ε there is a positive δ such that for every two points x, y in

[4] See remark below about the non monotonicity of the ex-post Nash property in the information received.

the common domain and for every function f in the collection, $|f(x) - f(y)| < \varepsilon$ whenever the distance between x and y is less then δ. For example, every finite collection of continuous functions defined on a common compact domain is uniformly equicontinuous.

Theorem 1. ***Robust Large Games**: Consider a family of semi-anonymous games $\Gamma(\mathcal{A}, \mathcal{T}, \mathcal{U})$ with uniformly equicontinuous $\mathcal{U}$ as above, and a positive number ε. There are positive constants A and B, $B < 1$, such that (simultaneously) all the equilibria of all the games in Γ with m or more players are (ε, ρ_m) robust with $\rho_m = AB^m$.*

3.2. **Proof of the Main Result.** The proof of the above theorem follows two steps: (1) all the equilibria above become ex-post Nash at an exponential rate as the number of players increases, and (2), that this implies that they become robust.

The first step is the following result from Kalai (2002).

Lemma 1. *For every positive ε there are positive constants A and B, $B < 1$, such that (simultaneously) in all the games in $\Gamma(\mathcal{A}, \mathcal{T}, \mathcal{U})$ with m or more players all the equilibria are (ε, ρ_m) ex-post Nash with $\rho_m = AB^m$.*

The above, together with the next lemma, yield directly the proof of the theorem but with $\rho_m = mAB^m$. This, however, is sufficient for completing the proof, since we can replace A and B by bigger positive A' and $B', B' < 1$, for which $mAB^m < A'B'^m$for $m = 1, 2, \ldots$.

Lemma 2. *If σ is an (ε, ρ) ex-post Nash equilibrium of an n player game G, then for any $\zeta > 0$, σ is an $(\varepsilon + \zeta, n\rho/\zeta)$ robust.*

Proof. Suffices to show that for any versions $\overline{G}$ and $\overline{\sigma}$ and for any player i, $\Pr(V) \leq \rho/\zeta$, where V is the "violation" event, that player i has a better than $\varepsilon + \zeta$ improvement at some information set. Let W be the set of composition profiles c with the property that player i has a better than ε improvement at c, i.e., by a unilateral change of his action he can improve his payoff by more than ε. By assumption, $\Pr(W) \leq \rho$.

We claim and argue below that for any positive probability information set of player i, A, if player i has better than $\varepsilon + \zeta$ improvement at A, then $\Pr(W|A) > \zeta$. Notice that we can write $V = \cup A_j$ as a disjoint union of such information sets A_j. So that using this claim in the following set of inequalities, $\rho \geq \Pr(W) \geq \sum \Pr(W \cap A_j) = \sum \Pr(W|A_j) \Pr(A_j) > \zeta pr(V)$, establishes the proof of the lemma.

Before arguing the validity of the claim above, we first note that at any positive probability information set A, any modification of player i strategy at A, does not affect the probability distribution over the profile of opponent compositions. This assertion can be checked node by node. At every node, the opponents that played earlier in the game tree will not revise (by the definition of $\overline{\sigma}$) and their compositions are fixed regardless of the modification. The opponents that did not play prior to reaching the node, will randomize according to σ when their turn to play comes, disregarding what other players, including player i, did before them.

The assertion just stated implies that without loss of generality we can check the validity of the claim at information sets A where player i can improve by more than $\varepsilon + \zeta$ through the use of a modification of $\overline{\sigma}_i$ that uses a pure strategy b at A and never revises it later on.

Now we can put a bound on possible level of such improvement as follows. For composition profiles in W, the largest possible improvement by player i using a different pure action is 1 (due to the normalization of the utility functions), and for composition profiles in W^c, the largest possible improvement is ε. This means that the highest possible improvement in the information set A is $1\Pr(W|A) + \varepsilon\Pr(W^c|A)$. So if the possible improvement at A is greater than $\varepsilon + \zeta$, we have $\Pr(W|A) + \varepsilon > 1\Pr(W|A) + \varepsilon\Pr(W^c|A) > \varepsilon + \zeta$, which validates the claim made above. □

4. Further Elaboration

4.1. **On the level of anonymity.** The condition of anonymity imposed above is less restrictive than may appear, since it only imposes anonymity within the payoff functions but without further restrictions of symmetry on the players. 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 opponents anonymously, since the buyers, i.e., the players called n+1,...,2n, play a role different from the other players. But if we assume that 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 every player to be one of two possible types, a seller or a buyers. Assign every player 1,...,n a prior probability one of being a seller type, and every player n+1,...,2n a prior probability one of being a buyer type. Now we can write the payoff function of the above seller in the obvious way to depend on the empirical distribution of types and actions, without dependence on 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, our 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.

4.2. **On the continuity assumption.** The continuity assumption, when combined with the assumption of semi anonymity, may be more imposing than appears. 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 that prefers I, or a type that prefers M. This player's payoff is 1 when he chooses the computer that he prefers and 0 otherwise. All 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 every other player probability one of being of the third type.) In this game, player 1 choosing the computer he prefers and every other player randomizing with equal probability between the two computers is an equilibrium of the one shot simultaneous move game.

It is easy to see, however, that the above equilibrium is not approximately robust, no matter how large the number of buyers is. For example, in the version where player 1 chooses first and the other players observe him before making their own choices, the above equilibrium no longer holds. The difficulty is that the players payoff functions cannot be imbedded in a collection of uniformly equicontinuous payoff functions g as required by the theorem. In the games described above, the percentage of expert types goes to zero as the number of players increases. So if the game is large, a player's payoff must be close to what one of these functions g pays at a composition profile with zero proportion of experts. But arbitrarily close to such a profile there are profiles with payoff 1 and profiles with payoff 0.

We should also note that one needs less than full continuity to obtain robustness of an equilibrium. The proof of the theorem used two parts: (1) that under the assumptions made all equilibria are shown to be (ϵ, ρ) ex-post Nash, with ρ going down to zero at an exponential rate, and (2) that such an equilibrium must be $(\epsilon, n\rho)$ robust. The assumptions of anonymity and continuity were used for the first part. So if we can weaken them and still obtain the conclusion of the first part, we would have a stronger overall result.

Indeed, in Kalai (2000, 2002), where the first part is proven, there are two such possibilities. First, it is shown that the continuity assumption does not have to hold globally, but only near the expected play of an equilibrium, in order to make it nearly ex-post Nash. Moreover, less than local continuity is needed. All that one needs is that the strategic dependence of a player on his opponents goes to zero as the possible variations in opponents' types and actions goes to zero.

4.3. **Subgame Perfection with many players.** As mentioned in the introduction, an extensively robust equilibrium is required to remain a Nash equilibrium, without subgame perfection, in every extensive version of the game. Consider for example the 2 person complete information computer choice game with both players preferring to match, but with Player 1 preferring to match on computer I and with Player 2 preferring to match on computer M, i.e., a "battle of the sexes" instead of a coordination game.

If Player 1 moves first the only subgame perfect equilibrium results in both choosing I, and if Player 2 moves first the only subgame perfect equilibrium results in both choosing M. So we cannot have an extensively robust equilibrium, an equilibrium that is sustained simultaneously in all extensive versions, that is also subgame perfect.

On the other hand, both choosing I and both choosing M are each a simultaneous Nash equilibrium of both extensive games above. The following example shows that while these are not fully subgame perfect, they are *highly* subgame perfect when the number of players is large.

Example 1. ***n-Person Battle of the sexes:*** *Simultaneously, each of n male players and n female players have to choose computer I or computer M. A male's payoff equals the proportion of the total population that he matches if he chooses I, but only* 0.9 *times the proportion he matches if he chooses M. For a female the opposite is true, she is paid the full proportion that she matches when she chooses M and* 0.9 *of the proportion she matches when she chooses I.*

Consider the above game played sequentially in a predetermined known order in which the females choose first and the males follow. Every player is informed of the choices made by all earlier players.

We first verify that for any n the only outcome of a subgame perfect equilibrium is for all players to choose M. Clearly, assuming that all his predecessors chose M, the last male's optimal choice is M. Given that, in the move just before this, the $n-1$st male's optimal choice is M. Following this standard backwards reasoning, we conclude that all players choosing M is the only subgame-perfect-equilibrium play path.

So what is wrong with the Nash equilibrium in which all players choose I? Lets start first with the case of only two players, $n = 1$, and follow the standard objection to non subgame perfect equilibrium. Despite the prevailing equilibrium of both players choosing I, if the female, who moves first deviates, the best response of the following male is to also defect from the equilibrium and choose M. The female knowing this, should therefore deviate, counting on his rationality, and improve her overall payoff.

But let us imagine the same senario, an equilibrium with all players choosing I, but with one million females moving first and one million males following. It is true that if all players except for the last males deviated to play M, his rationality would force him to also play M. But let us view the situation from the point of view of the first female.

In order to make her deviation worth while, she must believe that substantially more than n followers will deviate too. Otherwise deviating on her part may be quite costly and highly suboptimal. Moreover, her immediate follower has similar concerns that may prevent her from deviating to M. Players no longer count just on the rationality of their followers, they must count on followers counting on the counting on the counting... of the rationality of followers. So, unlike in the case of the 2 player game above, where a deviation by the first deviator induces direct immediate incentives to deviate by her follower, the incentives here are much weaker. The idea of deviating itself, is almost a move to another equilibrium that has to be taken on with simultaneous beliefs by more than n players.

We refer the reader to Kalai and Neme (1992) for a general measure of subgame perfection that formalizes this idea. The constant I equilibrium, while not fully (or infinitely in the Kalai and Neme measure) subgame perfect, is n-subgame perfect. In the example with two million players at least one million deviations from the equilibrium path are required before we can get to a node where a player's choice is not optimal, given the history and the equilibrium strategies of his followers.

5. References.

Aumann, R.J., and L.S. Shapley (1974): *Values of Nonatomic Games*, Princeton University Press, Princeton.

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

Feddersen, T,. and W. Pesendorfer (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica* 65 (5), pp 1029-58.

Forges, F., and E. Minelli (1997): "Self-fulfilling mechanisms and rational expectations", *Journal of Economic Theory* 75, 388-406.

Forges, F., and E. Minelli (1998): "Self-fulfilling mechanisms in Bayesian games", *Games and Economic Behavior* 25, pp 292-310.

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

Hansen, L.P. and T. Sargent (2001):"Robust Control and Model Uncertainty," *American Economic Review*; 91(2), pp 60-66...

Jordan, J.S., and R. Radner (1982): "Rational Expectations in Microeconomic Models: An Overview," *Journal of Economic Theory*, 26 (2), pp 201-223.

Kalai, E., and A. Neme (1992):, "The Strength of a Little Perfection," *International Journal of Game Theory* 20, pp 335-355.

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

Kalai, E. (2002), "Ex-Post Stability in Large Games," Northwestern University discussion paper.

Kovalenkov, A., and M. Wooders (2001): "Epsilon cores of games with limited side payments; Nonemptiness and
equal treatment" *Games and Economic Behavior* 36, no. 2, pp. 193-218.

Mailath, G.J., and A. Postlewaite (1990): "Asymmetric Information Bargaining Problems with Many Agents," *Review of Economic Studies*, 57, pp 351-67.

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

Minelli, E., and H. M. Polemarchakis (1997): "Knowledge at Equilibrium," forthcoming in *Rationality and Knowledge*, TARK VI, Morgan-Kaufmann.

Pesendorfer, W., and J. Swinkels (1997): "The Loser's Curse and Information Aggregation in Auctions," *Econometrica* 65 (6), pp 1247-82.

Rustichini, A., M.A. Satterthwaite, and S.R. Williams (1994): "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica* 62 (5), pp 1041-64.

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

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