

Kalai, Ehud

Working Paper

Private Information in Large Games

Discussion Paper, No. 1312

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Kalai, Ehud (2000) : Private Information in Large Games, Discussion Paper, No. 1312, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221666>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Private Information in Large Games*

Ehud Kalai[†]

J.L. Kellogg Graduate School of Management

Northwestern University

Evanston IL 60208

kalai@kellogg.northwestern.edu

December 21, 2000

Abstract

This paper identifies a type of Bayesian equilibrium, called information-proof, defined for n-player one-shot Bayesian games. This equilibrium is immune to changes in the prior probability of types, in the probabilities of mixing pure actions, and in the order of play, as well as to information leakage and the possibility of revisions.

In large semi-anonymous games that satisfy continuity and independence conditions all the equilibria are shown to be approximately information proof.

Information-proof equilibrium, like rational-expectations equilibrium, may be used to model ongoing rational interaction, as opposed to just instantaneous one-shot play.

1 Introduction

A major concern with Nash and Bayesian equilibria is their sensitivity to details that are often not available in the real life situation they model. Of particular concern is the dependency of such equilibria on subjective prior probabilities, the order of player moves and the precise information that players possess when they are called to make a move.

Concentrating on simultaneous-move one-shot normal form and Bayesian games, this paper identifies Nash and Bayesian equilibria, called *information proof*, that are stable under changes in informational details, changes in the order of moves and revision possibilities. While such equilibria do not exist in

*This paper is dedicated to the memory of John C. Harsanyi. His concept of Bayesian equilibrium is the central topic of investigation of the current paper.

[†]The author thanks Eddie Dekel, Jeff Ely, Ted Groves, Matt Jackson, Adam Kalai, Andy Postlewaite, Alvaro Sandroni, 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.

many games, in classes of large semi-anonymous games that satisfy continuity and independence conditions, all the equilibria are approximately information proof..

To understand the concept of information proofness, recall that at a Bayesian equilibrium every player, with the possible aid of a randomization device, selects an action after being informed only of his own realized type. But his payoff depends on the realized types and selected actions of all the players. Not knowing his opponents realized types and selected actions, he chooses his action to be optimal relative to the distribution of their types and their action-selection rules.

A Bayesian equilibrium is information proof if, in addition to the above, every vector of types and actions that may be realized with positive probability has the following ex-post stability property. No player has an incentive to unilaterally change his own selected action, even after being fully informed of the realized types and selected actions of all his opponents. In other words, ex-post, knowing all the information regarding the outcome of the game, no player regrets his chosen action. Or, alternatively, the selected actions are always a pure-strategy Nash equilibrium of the complete information game induced by the realized types.

The interest in stability of this type is not new to this paper. Motivated by special problems in economics, Green and Laffont (1987) introduced the notion of ex-post Nash implementation. The idea is that in solving an implementation problem, after the actions of the players become publicly known, the incentives of the players to play the selected actions should not change. If the strategies are only ex-ante incentive compatible, and the selected actions get modified once they are observed, what is actually implemented may be drastically different from what was meant to be implemented.

If used by a designer of an implementation procedure, information-proof equilibrium guarantees among other things ex-post Nash implementation in the Green and Laffont sense. Therefore, the asymptotic results described in the current paper suggest that ex-post Nash implementation should be easier in a large group of participants. One may expect any Nash implementation to be automatically ex-post Nash, or at least approximately so, provided that certain anonymity, independence, and continuity conditions are satisfied by the selected mechanism.

It turns out, however, that beyond ex-post stability, information-proof equilibria have other stability properties. They are immune to certain changes in the prior probabilities by which player types are drawn, to changes in the probabilities that opponents use to mix their pure actions, and to sequential play with general information structure and revision possibilities. Possessing such robustness properties, like the economic notion of rational-expectations equilibrium, information-proof equilibrium may be used to model the steady state of an ongoing rational interactive process, as opposed to just the equilibrium of an instantaneous one-shot play. In addition, the incentive to learn opponents' types and to predict their selected actions, that is typical at a general Bayesian equilibrium of a one-shot or a repeated game, disappears when the

equilibrium is information proof. And information-proof equilibrium can serve as a purification device, generating pure strategy equilibria for the game under consideration.

In addition to the stability properties discussed above, information proofness in large games offers a different outlook on private information in such games. Being information proof means that even if all the information about all the relevant choices in the game were publicly available, for the purpose of assessing his own selected action, each player would be interested only in his own private information. This means that private information is a private concern, or alternatively, full decentralization in the use of information is compatible with individual incentives.

The next section illustrates some of the ideas through simple examples. The following presents a formal definition of information-proof equilibrium, its structure and relationship to Nash's notion of interchangeable equilibria, and its stability properties. The next section develops bounds on the level of information proofness of any given Bayesian equilibrium. These bounds are used in the following section to derive the asymptotic convergence to information-proof equilibria as the number of players becomes large.

Two additional features should be noted. First, the bounds developed on information proofness are sufficiently strong to imply that the asymptotic convergence is at an exponential rate in the number of players. Moreover, this convergence is uniform. The convergence to information proofness occurs simultaneously by all the equilibria of all the large games within the classes under consideration, and at each such equilibrium most of the outcomes are *ex post* Nash for all the players simultaneously. Second, the bounds are computed relative to a measure of strategic interdependence developed in this paper. In addition to being useful for this purpose, this measure, which is closely related to Nash's notion of interchangeability, is likely to have other applications. The final section contains additional discussion and questions for follow-up research.

The literature related to this paper is too large to be fully discussed here. The following are some suggestive examples. The first paper on general strategic anonymous games with many players is Schmeidler (1973). Unlike here, Schmeidler's paper concentrates on games with a continuum of players and complete information and studies existence of pure strategy equilibria. Despite these differences, it seems reasonable to expect connections between his paper and the information-proof equilibria discussed here. If one views the continuum as the limit of finite games with many players, and as the laws of large numbers become perfect in the limit, the incomplete information may fully disappear and the equilibria coincide.

Uses of laws of large numbers, similar to the one here, were made in many economic and game theory papers. This is due to the basic phenomenon that when a number of interactive participants is large, a player does not have to consider the actions of individual opponents. He can act, instead, against the aggregate behavior as predicted by the prior probabilities. Examples of such can be found in Hart, Hildenbrand and Kohlberg's (1974) work on large markets,

Groves and Hart's (1982) generalization of the work of Arrow and Radner (1979) on teams, Rustichini, Satterthwaite and Williams' (1994) and Satterthwaite's (1999) work on double auctions, Pesendorfer and Swinkels' (1997) work on auctions and Feddersen and Pesendorfer's (1997) work on voting.

While this paper discusses the relationship of information-proof equilibrium to rational-expectations equilibrium, a more thorough exposition to rational-expectations equilibrium can be found in Radner (1981 and 1991) and in Jordan and Radner (1982). Moreover, related studies connecting rational-expectation equilibrium to games of incomplete information can be found in Forges and Minelli (1997 and 1998), Minelli and Polemarchakis (1997) and Minehart and Scotchmer (1999).

Other connections, ones that suggest potential future research, are offered by the literature on player smallness, see for example Green (1980), Sabourian (1990), Fudenberg, Levine and Pesendorfer (1998), and Al-Najjar and Smorodinsky (2000), and by the literature on information smallness, see for example Mailath and Postlewaite (1990), Gul and Postlewaite (1992), and McLean and Postlewaite (1999). As discussed in the concluding section of this paper, phenomena described in this paper may be found to have a greater generality when facts from this literature are brought into the analysis.

Finally, another topic that is discussed in the concluding section is the connection of information proofness to the Bayesian learning literature, see Jordan (1991) and Kalai and Lehrer (1993). The somewhat surprising conclusion may be that learning and playing becomes easier as the number of opponents increases beyond some critical number.

1.1 Examples and Elaboration

The first two examples, a coordination game and a match pennies game, offer a quick illustration of the ideas involved. To connect them better to economic applications, they are presented in a language of economic choices.

Example 1 : *Coordinated computer choices.* *Simultaneously, each of two players have to choose computer B or M. If their choices match they will each be paid 1, and if they mismatch, they will each be paid 0.*

Example 2 : *Mis/matching computer choices.* *Each of two players, a game theorist and a poet, have to choose computer B or M. If their choices match the game theorist is paid 1 and the poet 0, and if they mismatch the game theorist is paid 0 and the poet 1.*

Compare a pure strategy Nash equilibrium of the coordination game, say when they each choose computer B, with the mixed strategy Nash equilibrium of the mis/matching game, where each player chooses one of the computers randomly with equal probability.

Starting with the question of information proofness, it is easy to see that the pure strategy equilibrium of the coordination game is information proof, while the mixed strategy equilibrium of the mis/matching game is not. In the former,

after being informed of the selected choices, (B, B), no player has incentive to unilaterally change his own choice. This is clearly not the case for the mixed strategy equilibrium of the mis/matching game. There, there are four possible outcomes, (B,B), (B,M), (M,B), and (M,M), and at each one of them, when the outcome becomes publicly known one player has incentive to revise his choice. This strongly violates the condition of information proofness that at *no outcome* should any player have an incentive to revise.

Indeed in n -person normal form games, as obvious from its definition, any pure strategy Nash equilibrium is information proof. This is not the case for mixed strategy equilibrium. There is only a small group of mixed strategy equilibria that are information proof. To be so, as discussed later in the paper, the support of the mixed strategy profile must consist of pure strategy equilibria that are interchangeable in Nash's sense. But the situation is less clear cut when we move from normal form to Bayesian games where there are many pure strategy equilibria that fail to be information proof.

The difference in stability of the two equilibria above are also apparent. If the coordinated matching game was played sequentially, say with player one moving first and his choice being observed by player two before making her own move, the (B,B) equilibrium would still be a Nash equilibrium of the extensive form game (player one's strategy is to play B and player two's strategy is to play B independently of what she observes, see the concluding section for a discussion of subgame perfection). Again, this is not the case for the extensive form of the mis/matching game. Here, if the matcher moves first and uses the mixed strategy it is no longer optimal for the mismatcher to mix her choices. She would rather observe the realized choice of the matcher and choose her own to mismatch his with certainty. As discussed later in the paper, invariance to sequential play is a general property of all information-proof equilibria, not just pure strategy equilibria in normal form games.

A second issue of stability relates to sustainability over time. If the computers above were to be used for some continuous period of time, rather than an instantaneous one-shot use, then the first equilibrium remains stable while the second one does not. Whatever the outcome of the randomization is in the second game, as soon as the players see the computers selected by the opponents, one of them will switch. Following this, the other player would switch, and so on. It would be hard to model the outcome of the game in this situation. But for the pure strategy equilibrium, as for all information-proof equilibria, there is no problem. The players can continuously use the computers (B,B) without anyone having the incentive to change.

To see how large numbers make equilibria approximately information proof, even when mixed strategies are used, consider the following.

Example 3 : Mis/matching computer choices with many players. *Simultaneously, each of n game theorists and n poets have to choose computer B or M. The payoff to every game theorist equals the proportion of poets his computer matches. The payoff to every poet equals one minus the proportion of game theorists his computer matches.*

With $n = 2$ this example is the earlier mis/matching game. Thus, it is impossible to have information-proof equilibria here for all values of n . But for large values of n , the equilibria are nearly information proof. If every player, game theorist or poet, chooses one of the two computers randomly with equal probability, then within each group the proportions of the two selected computers are likely to be close to one half, and no player would be able to gain much by switching. This is a strong version of an approximate stability since there is a high probability for the events of no possible improvement greater than some given epsilon holding *simultaneously* for all players.

It is impossible, however, to have approximate information-proof equilibria in every large game, as illustrated by the following.

Example 4 : Majority mis/matching game. *Again, simultaneously, n game theorists and n poets must each choose computer B or M. The payoff to a game theorist is one, if his computer matches at least one half of the poets and zero otherwise. The payoff to a poet is zero, if his computer matches at least one half of the game theorists, and one otherwise.*

If n is odd, for example, no matter what strategies are played, at every publicly known outcome at least one half of the players will have a strong incentive to unilaterally revise their choices. This failure of information proofness is due to a discontinuity in the payoff functions. In particular, small variations in the empirical distribution of computer choices in the large population may drastically alter the payoffs of some players.

Discontinuity in the payoff may become more severe when Bayesian games, as opposed to normal form games, are considered. Now, even strict pure-strategy equilibria in large games may fail to be approximately information-proof.

Example 5 : Matching the expert. *Again, simultaneously, each of n players have to choose computer B or M. But every player may be of two different types. Player one is equally likely to be an expert who likes B or an expert who likes M. Independently, every other player is equally likely to be a nonexpert who likes B or a nonexpert who likes M. The payoff to player one is one, if he chooses the computer he likes, and zero otherwise. To every other player the payoff is the sum of two components. The first component is $1/3$, if he likes the computer he chooses, and zero otherwise. The second component in his payoff is $2/3$ if the computer he chooses is the one the expert likes, and zero otherwise.*

The strategies consisting of every player choosing the computer they like constitute a strict pure-strategy Nash equilibrium. However, it is not approximately information proof, no matter how large n is. With high probability, once the choice of the expert is revealed, close to one half of the players would have strong incentive to revise their own choices.

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. To save notations, 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.

$N = \{1, 2, \dots, n\}$ 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* .

$\tau = (\tau_1, \tau_2, \dots, \tau_n)$ 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 $u = (u_1, u_2, \dots, u_n)$, assuming a suitable normalization, are of the form $u_i : C \rightarrow [0, 1]$.

In addition, standard game theoretic conventions are used throughout the paper. For example, for a vector $x = (x_1, x_2, \dots, x_n)$ and an element x'_i , $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $(x_{-i} : x'_i) = (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$.

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 device. 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 $\sigma_i = (\sigma_i(a_i | t_i))$ where $\sigma_i(a_i | t_i)$ 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, $\gamma_i(c_i) = \tau_i(t_i) \times \sigma_i(a_i | t_i)$. The profile of these distributions, $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$, under the independence assumption, determines the overall probability distribution over *outcomes of the game*, namely composition profiles, by $\Pr(c) = \prod_i \gamma_i(c_i)$. Unless otherwise specified, all probabilistic statements in this paper are made with respect to this distribution.

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

3 Information-Proof Equilibrium

Definition 6 : Information Proofness.

A profile of compositions c is information proof for player i if for every action a'_i , $u_i(c_{-i} : (t_i, a'_i)) \leq u_i(c)$.

A profile of compositions c is information proof if it is so for every player.

A strategy profile σ is information proof if it yields information-proof composition profiles with probability one.

In words, at an information-proof composition profile, even under full knowledge of the complete profile, no player has incentive to unilaterally change his own selected action. Alternatively, the vector of actions described by the profile is a Nash equilibrium of the complete information game determined by the vector of profile types. This means that a strategy profile is information proof if with probability one it would lead to outcomes that are ex-post Nash. Ex-post here is in a strong sense, since it requires that the realized actions, not the mixed strategies, constitute a Nash equilibrium of the complete information game.

3.1 Structure

To gain further insight into this concept, consider first the simple case of normal form games. Since such games are a subclass of Bayesian games (with every player being restricted to one possible type) the definition of information proofness is applicable. As discussed in the introduction, it is easy to see from its definition that every pure strategy Nash equilibrium is information proof. But in general, for a mixed strategy equilibrium to be information proof it has to have a highly restricted special structure. To understand this, it is useful to generalize Nash's notion of interchangeability, see for example Luce and Raiffa (1957), from two to n player games and from a pair of equilibria to any number of them.

Definition 7 : Interchangeability. Consider a set M of pure strategy Nash equilibria of a given n -person normal form game. The equilibria in M are interchangeable, if every action profile a which is a coordinatewise combination of elements in M , i.e., every coordinate $a_i = a'_i$ for some $a' \in M$, is a Nash equilibrium of the game.

When a set of equilibria is interchangeable, maximizing against one of the profiles from the set means maximizing against all of them, or even against any coordinatewise combination of them. So every player may be playing a different equilibrium, yet everybody is maximizing to the actual play which is at equilibrium. Or in other words, for the purpose of choosing their own individual actions, the players do not have to be concerned with which equilibrium is played.

Proposition 8 A (mixed strategy) equilibrium of an n -person normal form game is information proof if and only if its support consists of a set of interchangeable pure strategy Nash equilibria.

Proof. Clearly if the support of a strategy profile consists of Nash-equilibrium action-profiles, it must be information-proof. And, as was already discussed, every action profile in the support of an information-proof equilibrium must be a Nash equilibrium of the game. Combining this with the fact that the support of a strategy profile must contain all its coordinatewise combinations, the conclusion follows. ■

The structure described above carries over easily to general Bayesian equilibrium under the assumption of type independence assumed throughout this paper. Consider a set of composition profiles S , and recall that every one of its elements $c = (c_1, \dots, c_n)$ can be viewed as a pair of profiles (t, a) describing the corresponding types and actions of the n players. The elements of S are information-proof interchangeable if every coordinatewise combination of such elements, i.e., c with every $c_i = c'_i$ for some $c' \in S$, has the information-proof property, i.e. a is a Nash equilibrium of the complete information game induced by t .

Proposition 9 *A strategy profile of a Bayesian game is an information-proof equilibrium if and only if its support consists of information-proof interchangeable composition profiles.*

Proof. Under the assumption of type independence the support of a strategy profile includes all the coordinatewise combinations of its elements. Thus, by its definition, every information-proof equilibrium must have an interchangeable support. Conversely, if the support of the strategy profile consists of composition profiles that has the information-proof property then it is by definition an information-proof equilibrium. ■

Note that the same interpretation of Nash interchangeability carries over to information-proof interchangeability. Within a set of information-proof interchangeable profiles, a player who learns that he is of type t_i and plans to take the action a_i does not have to be concerned with the multiplicity issue regarding the possible types and actions of his opponents. If, given his type, his action is optimal against one of the realizations of his opponents' parameters, it is also optimal against any one of the other realizations in the interchangeable set.

3.2 Stability

While the condition of information proofness requires ex-post stability after complete revelation of information about the final outcome of the game, it is also stable to revelation of partial information. Formally, let a signalling function be any $g : C \rightarrow M$, where C is the set of composition profiles and M is an abstract set of messages. The Bayesian equilibrium σ is (i, g) -information proof 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 10 *An information-proof equilibria is (i, g) -information proof for all players i and all signalling functions 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's i information t_i, a_i , and $g(c)$. The assumed information proofness implies that switching at any $\bar{c}$ 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. ■*

Information-proof equilibria are stable under certain changes in the prior distribution, and to 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 with no new types* if for every type profile t , $\tau'(t) > 0$ implies that $\tau(t) > 0$.

Proposition 11 *If an equilibrium σ is information proof in a game G then it is also an information-proof equilibrium in every game G' with a modified priors with 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 has the desired ex-post robustness property. ■*

For exactly the same reasons the following is also true. Given any equilibrium σ of G , a *modified strategy profile with 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 12 *If an equilibrium is information proof then every modified strategy profiles with no new actions is an information-proof equilibrium.*

The above two propositions illustrate that in playing or in analyzing a simultaneous move Bayesian game, at an information-proof 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. The following observations discuss stability with regard to sequential play, information leakage and revision possibilities.

Given a Bayesian game with a set of players N , for any ordered partition of the set of players $O = (N_1, N_2, \dots, N_r)$, consider the sequential extensive form game described as follows. In an initial stage all the players are privately informed of their own realized types, drawn according to the prior distributions of the given game. Then, simultaneously in a first stage, the players in N_1 choose actions from their feasible sets; each player may use a private randomization device. The realized chosen actions are made public knowledge. Following this, simultaneously in a second stage, the players in N_2 choose actions which are made publicly known. The game ends after the r 'th stage with the simultaneous final selection of actions by the players in N_r . Payoffs of this extensive game

are made at the realized composition profile according to the utility functions of the given game.

Every strategy profile σ of the given Bayesian game induces a *corresponding* profile of constant strategies in every O ordered sequential play as follows. Player i , on his turn to play, ignores all the public information received about earlier players, and randomizes over his choices just as he does in $\sigma_i(\cdot|t_i)$. A Nash equilibrium of the Bayesian game is *order invariant*, if the corresponding constant strategies it induces are Nash equilibria in every ordered play of the game as described above.

Proposition 13 *Information-proof equilibria are order invariant.*

Proof. *The order invariance follows from the fact that information-proof equilibria are ex-post stable also with respect to partial information, which is what the players have in an ordered play. ■*

But in general, information-proof equilibria satisfy a stronger form of invariance. This allows for the order of play to be determined randomly, as the game evolves, with players possessing only partial knowledge about the order, who has already played, who will get to revise, etc., and with opportunity for revisions evolving as the game progresses.

Formally, starting with the given simultaneous move Bayesian game $G = (N, T, \tau, A, u)$ and a given finite set of messages M , define a *sequential revisional form* to be any finite extensive form game of the following type. The initial node in the game tree belongs to nature with the outgoing arcs being labeled by the elements of $T \times M$. Any probabilities may be assigned to these arcs as long as the marginal distribution over T coincide with the prior probability distribution of the underlying Bayesian game. 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 . At every information set the moving player i has 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 coordinate t_i). Every play path in the tree visits at least once one of the information sets of every player. Finally, associate to every complete path in the tree a composition profile $c = (t, a)$ with t being the initial choice of types selected by nature, and with each a_i being the last action, in the play path, taken by player i . The payoffs associated to a complete path equal the payoffs of the players from the resulting composition profile in the given simultaneous move game.

Given a strategy profile σ of a Bayesian game and a sequential revisional form, define the *corresponding* constant strategies to be the ones where every player at every one of his information sets randomizes over his actions a_i with probabilities $\sigma_i(a_i|t_i)$, where t_i denotes his type at the information set. In what follows there is an abuse of language where σ is used to denote a vector of strategies in the underlying Bayesian game, but also the corresponding constant strategies in its sequential revisional forms.

Before stating the general result, however, the following example shows that for a given information-proof equilibrium σ of a given Bayesian game, there may be sequential revisional forms in which σ is not an equilibrium.

Example 14 Simultaneously, each of two players has to choose among two actions, Y or N . Player 1's payoff depends entirely on player 2, he is paid 1 if she chooses Y and 0 if she chooses N . Player 2 is completely indifferent, her payoff is zero no matter what was chosen. Every player randomizing, with equal probability between the two actions, is an information proof equilibrium. Whatever pair of actions are realized nobody can gain by unilaterally changing his own action, because nobody really controls his own payoff in this game.

But in the following order of play and revision opportunities the corresponding constant strategies are no longer in equilibrium. Player 2 moves first and with complete information of her choice player 1 follows. If player 1 follows with a Y the game ends and his payoff is determined by her first choice. If he follows with an N then player 2 makes a fresh choice, and player 1's payoff is determined by her new choice.

Notice that if player 2 randomizes with equal probabilities at every one of her three nodes, then randomizing by player 1 at his two nodes is a suboptimal. The unique optimal response for him is to play Y and terminate the game after Player 2 chooses the good action for him, Y , and play N to have player 2 randomize again if her initial choice is the bad one for him, N .

It turns out, however, that an information-proof equilibrium remains an equilibrium in all sequential revisional forms that satisfy a suitable restriction on the revision possibilities. A node of player i in a sequential revisional form is an *ambiguous final revision* if two play paths that go through this node have the following properties. On one play path player i plays again after the given node (has a later decision node) and on the other play path he never moves again. In other words it is not clear whether player i will play again without knowing what he and the other players will do.

Proposition 15 *An equilibrium of the Bayesian game is information proof if and only if it is an equilibrium in every sequential revisional form that has no ambiguous final revisions.*

Proof. First, if σ is invariant in the sense above, for every player i consider a sequential revisional form with all the players other than i moving sequentially first, and player i moving last with complete information of all the realized types and all the realized actions of the earlier players. The fact that $\sigma_i(\cdot|t_i)$ is optimal for him under this full information means that every positive probability composition profile is information proof for him.

To show the converse assume that σ is an information-proof equilibrium of the given Bayesian game and consider a given sequential revisional form. It suffices to show that σ is an equilibrium of the perfect information version (every information set is a singleton) of the sequential revisional form. Moreover, if σ is an equilibrium in every subgame that follows nature choice, then it is an equilibrium of the entire game. This means that without loss of generality, the proof of the converse can be restricted to normal form games (each player has only one type) and to perfect information sequential revisional forms.

The lemma that follows shows that in such an extensive game, if the players choose strategies that correspond to σ , the induced probability of an action profile

a is $\prod_j \sigma_j(a_j)$ (conditioning on t is omitted due to the assumption that the game is a normal form one), whereas if player i changes his strategy, even to a non-constant one, this probability is of the form $\eta_i(a_i) \prod_{j \neq i} \sigma_j(a_j)$, where η_i is some probability distribution over A_i . Thus, the proof follows from the fact that σ_i is i 's best response in the simultaneous move-one shot game. ■

Lemma 16 Consider a finite n person normal form game, a perfect information sequential revisional form with no ambiguous final revisions and an information-proof equilibrium σ . For every node x let $A(x)$ denote the players that are still active in the subgame that starts at x (i.e., the players that have a move on some, equivalently all, the paths that follow the node x).

1. The distribution generated over the play paths in the subgame, i.e., over $\times_{i \in A(x)} A_i$, is given by $\Pr(a) = \prod_{i \in A(x)} \sigma_i(a_i)$.

2. For a strategy profile of the form (σ_{-i}, σ'_i) , the distribution generated over play paths in the subgame is of the form $\Pr(a) = \eta_i(a_i) \prod_{j \in A(x)} \sigma_j(a_j)$, with η_i being some distribution over the action set of player i , A_i .

Proof. The proof is by induction on the depth of the subgame (i.e., the maximal length of a path in it). If player i is the first mover in the subgame and this is his final revision, then both parts 1 and 2 follow directly from the induction hypothesis. If i is the first mover in the subgame and this is not his final revision, part 1 still follows directly from the induction. Part 2 also follows, but after noticing that in this case the overall distribution is a convex combination of a finite number of distributions each having the form in part 2, which yields a distribution of the same form. For part 1 one needs to also analyze the case where the first mover in the subgame is a player $j \neq i$. But the same analysis of the convex combinations that are created by the initial randomization of j leads to the same conclusion. ■

3.3 Ongoing Interaction and Rational-Expectation Equilibrium

Consider for example a computer rental market. Economic agents include owners and potential users of the existing computers. Each agent has some private knowledge, a private signal regarding the quality and other characteristics of the various computers; and his overall preferences for final computer choices depend not only on his own signal, but also on the unknown signals of the other agents and their own computer choices. Assuming that the agents are rational and know the prior probabilities by which signals are generated, what is an appropriate equilibrium concept to model the rental activity?

A naive game-theoretic approach may be to have every agent decide on a rental strategy as a function that depends only on his own signal. If every profile of strategies of this nature leads to a well-defined expected payoff, this establishes a well-defined game whose Bayesian equilibrium may be proposed as the proper solution. This approach, however, is too simplistic since it leaves out the market rental prices of the various computers. These are important for two reasons. First, even in a world of complete information (say everybody

shares the same signal) agents would like to know the prices before they make their choices. But, in addition, the prices are important since they may reveal information about the private signals of other agents. So if the prices are not revealed to the agents until after the play took place, then this model fails to capture an important aspect of trade, namely, the fact that prices are taken directly into consideration at the time decisions are made.

Recall that rational-expectations equilibria resolve the above difficulty by including prices, in addition to trades, in the domains of the strategies. A rational-expectations equilibrium identifies a fixed point of a system in which trades and prices are determined simultaneously. Moreover, it requires agents do Bayesian updating regarding opponents signals for the prices they observe, and maximize relative to these updated beliefs. In this sense, rational-expectations equilibrium may be viewed as a steady state of an ongoing process of trade rather than an instantaneous one-shot trade. But rational-expectations equilibrium does not allow for mixed strategies, and it may fail to exist.

Returning to the Bayesian equilibrium and its criticism above, notice that the criticism is highly diminished if the equilibrium is information proof. If a vector of prices is a function of the trade strategies and the realized private signals, then it may be viewed as a signalling function in the sense described earlier. The stability of information-proof equilibrium to ex-post signals then implies that no player would have an incentive to change his choice once the realized equilibrium prices are revealed. This means that the agents would be willing to carry on with their equilibrium trades even after observing the prices. Or put in other words, at a Bayesian equilibrium of such a market, price distribution was anticipated by the agents correctly and was optimally responded to even in an ex-post sense. Being information proof implies that once a realized price is revealed, still no agent wants to revise his choice.

3.4 Approximate Information Proofness

As stated in the introduction, approximate information-proofness is automatic in certain classes of large games. The sense of approximation is the following.

Definition 17 : (ε, ρ) *information proofness*. Let $\varepsilon \geq 0$ and $\rho \geq 0$ be given.

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

A profile of compositions c is ε information proof if it is so for every player.

A strategy profile is (ε, ρ) information proof if it yields ε -information-proof composition profiles with probability of at least $1 - \rho$.

Note that when ε and ρ are small the above approximation is strong in the following sense. It assigns a high probability, $1 - \rho$, to the event that, simultaneously, none of the players may improve their individual payoffs by more than the threshold ε .

It is also important to note that, due to the assumption that payoffs were normalized to be in the interval $[0, 1]$, the threshold ε describes the payoff improvement as a percentage of the total possible payoff variability in the game.

What happens to the robustness properties, discussed above for full information-proof equilibrium, as we move to approximate-information proof equilibrium? This question, which may become fairly involved and is left for future research, is important because the convergence properties that follow assure only approximate information proofness in finite large populations.

4 Bounds on Information Proofness of a Given Equilibrium

This section concentrates on a fixed Nash equilibrium of a fixed Bayesian game as defined above.

4.1 A Measure of Strategic Interdependence

The following measure is useful for constructing minimal bounds on the level of information proofness of the given equilibrium. It deals with situations where the play of the game is restricted to a set of composition profiles M , and conditional on that, it measures the extent to which a player's choice of an action is sensitive to changes in the composition profiles of the opponents. In this paper M describes all the possible compositions that may result with positive probability at the play of the given equilibrium. But this measure may be of interest in other contexts, for example, when one wishes to study the strategic interdependence within a subgame of an extensive form game. In this case, M will include all the outcomes that end up in the subgame.

Definition 18 : Strategic Interdependence.

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

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

with the maximum taken over all players i , all actions a'_i and a''_i , all types t_i , and all composition profiles $c_i^1, c_i^2 \in M$ with $t_i^1 = t_i^2 = t_i$.

The strategic interdependence in M is defined by

$$si(M) = \max_i sd_i(M).$$

Useful interpretations of strategic dependence can be obtained by rearranging terms within the absolute value in the definition above. One interpretation regards the change in any player i 's payoff when the opponents change from c^1 to c^2 and he holds action a'_i , versus the same change when he holds action a''_i . The difference in these changes can be at most $si(M)$. Thus, if the play is likely to take place in M and $si(M)$ is small, even though the opponents' types and actions may be highly relevant for his payoff, in deciding which action to select the player should not be highly concerned with the opponents' types

and actions. A similar interpretation, one which is used below, is obtained by rearranging the absolute value above to read as follows.

$$\max | [u_i(c_{-i}^1 : (t_i, a_i')) - u_i(c_{-i}^1 : (t_i, a_i''))] - [u_i(c_{-i}^2 : (t_i, a_i')) - u_i(c_{-i}^2 : (t_i, a_i''))] |$$

This interpretation shows that if $sd_i(M)$ is small and the composition profile is 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.

Low strategic dependence may be obtained if, when being of a certain type, the player's payoff depends mostly on his own actions, and not on the type or actions of his opponents, or if his payoff depends mostly on the types and actions of his opponents, and not on his own actions. In the first case, the player's choice of an action can be done with high disregard to the information on opponents. In the second case, information about opponents matters, but not in determining his own action since his own action is highly irrelevant. In the extreme case of zero dependence, the following characterization is obtained.

Proposition 19 *The strategic dependence of player i in a set of outcomes M is zero, $sd_i(M) = 0$, if and only if player i 's utility for composition profiles in M can be written as the sum of two functions*

$$u_i(c) = g_i(c) + h_i(c)$$

with g_i being dependent only on player i 's own compositions and h_i being independent of player's i own actions. More precisely $h_i(c^1) = h_i(c^2)$ whenever $(t_i^1, c_{-i}^1) = (t_i^2, c_{-i}^2)$, and for $c^1, c^2 \in M$ $g_i(c^1) = g_i(c^2)$ whenever $c_i^1 = c_i^2$.

Proof. It is easy to see that functions of the form of g_i and h_i have zero strategic dependence. Moreover the sum of functions with zero strategic dependence must also have zero strategic dependence. Thus if u_i is of the form above it has zero strategic dependence.

For the converse, fix an arbitrary action for player i , a_i^0 , and define for every c , $h_i(c) = u_i(c_{-i} : (t_i, a_i^0))$ and $g_i(c) = u_i(c) - h_i(c)$. Clearly h_i is independent of player i 's actions and u_i is the sum of the two functions. To see that g_i depends only on player i 's compositions let $c^1, c^2 \in M$ with $(t_i, a_i) = c_i^1 = c_i^2$.

$$\begin{aligned} g_i(c^1) - g_i(c^2) &= [u_i(c^1) - h_i(c^1)] - [u_i(c^2) - h_i(c^2)] \\ &= [u_i(c^1) - u_i(c_{-i}^1 : (t_i, a_i^0))] - [u_i(c^2) - u_i(c_{-i}^2 : (t_i, a_i^0))] \\ &= [u_i(c^1) + u_i(c_{-i}^2 : (t_i, a_i^0))] - [u_i(c^2) + u_i(c_{-i}^1 : (t_i, a_i^0))] \\ &= 0 \text{ because the zero strategic dependence of } i \text{ in } M. \blacksquare \end{aligned}$$

Given the decomposition above, it follows immediately that in normal form games, full strategic independence implies full interchangeability of all Nash equilibria.

Proposition 20 *If the set of all action profiles in a normal form game has zero strategic interdependence, then the set of all Nash equilibria is interchangeable.*

Proof. Under the decomposition above, the Nash equilibria of the game are exactly the action profiles a with the property that each a_i is a maximizer of the individual function g_i . $\blacksquare$

4.2 Bounds on Information Proofness

Lemma 21 *A Bayesian equilibrium is (ε, ρ) information proof if for a set of composition profiles M*

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

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

Proof. Suffice to show that at any $c \in M$, 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 $\bar{c} \in M$ with $\bar{t}_i = t_i$ (the improvement referred to is the following: fix the opponents' compositions and i 's type to be as in $\bar{c}$, 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 nonpositive, 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 information proof. The following discussion illustrates that under natural restrictions on the game, and assuming a large number of players, such sets are natural.

Starting with a strategy profile σ the induced vector of measures $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$ may be viewed as a vector of extended distributions, all 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 22 : Expected play. For a vector of strategy profile σ and the induced distribution $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_n)$ define the expected play on the universal set of compositions $\mathcal{C}$ to be the distribution $\exp_\sigma(\kappa) = \sum_i \gamma_i(\kappa) / n$

Definition 23 : Empirical play. For every composition profile c define the empirical distribution induced by c on the universal set of compositions $\mathcal{C}$ by $\text{emp}_c(\kappa) =$

(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 play 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 play 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 play.

Lemma 24 Chernoff-Hoeffding additive bounds. *Let $X_1, X_2, \dots, X_n$ be a sequence of independent 0-1 random variables with $\Pr(X_i = 1) = \mu_i$. Let $\bar{X} = \sum X_i/n$ and $\bar{\mu} = \sum \mu_i/n$. Then for every $\delta > 0$,*

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

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

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

Proof. Let $\bar{X}_{-i} = \sum_{j \neq i} X_j/(n-1)$ and $\bar{\mu}_{-i} = \sum_{j \neq i} \mu_j/(n-1)$. $|\bar{X} - \bar{\mu}| > \delta$ implies that $|\bar{X}_{-i} - \bar{\mu}_{-i}| > \delta n/(n-1) - 1/(n-1) = (n\delta - 1)/(n-1)$. The lemma follows by applying the previous lemma to $\bar{X}_{-i}$. ■

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 26 *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 possible configurations, 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 27 *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 information proofness to $M = nbd(\exp_\sigma, \delta)$, one obtains the following bounds.

Theorem 28 *The Bayesian equilibrium σ is (ε, ρ) information proof if for some $\delta > 0$,*

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

While somewhat long and complex, the above expressions offer a good insight about rates of convergence. For any fixed $\mathcal{C}$ and δ , the bound on ρ becomes small at a rate which is exponential in the number of players. A similar conclusion is possible for ε , provided however that the strategic interdependence is a function that goes to zero at a fast rate in small neighborhoods of the expected play. The next section exploits this observation to obtain convergence to information-proof equilibrium which is uniform for all equilibria in some interesting classes of games.

5 Information Proofness in Large Games

This section deals with large (many-players) Bayesian games in which the payoff influence of each player on his opponents is small. The players are not symmetric in two senses. First, every player may have a different prior probability distribution over his possible types. Second, different players may have different payoff functions. However, the game is *payoff anonymous* in the following sense. At any composition profile the payoff of every player depends entirely on the player's own composition and the empirical distribution over compositions induced by the opponents. Since it is the *distribution* of opponents compositions, this means also that a player's payoff cannot depend on the number of opponents.

Due to laws of large numbers, at equilibrium a player's anticipation of opponents' aggregate characteristics should be approximately correct, which implies that the equilibrium should be approximately information proof. This intuition is confirmed by an asymptotic result which shows a strong uniform rate of convergence.

As before, there is a universal finite set of possible player types, $\mathcal{T}$, a universal set of possible player actions, $\mathcal{A}$, and a resulting universal finite set of possible player compositions, $\mathcal{C} = \mathcal{T} \times \mathcal{A}$.

But in addition, there is a universal *set of possible payoff functions* U , which contains functions of the form $g : \mathcal{C} \times \Delta(\mathcal{C}) \rightarrow [0, 1]$. 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.

Formally, let the family of payoff-anonymous Bayesian games $\Gamma = \Gamma(\mathcal{A}, \mathcal{T}, U)$ 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 is imbedded in some function $g \in U$ so that $u_i(c) = g(c_i, emp_{c_{-i}})$.

Theorem 29 *Assume that the set of possible payoff functions U consists of a finite number of continuous functions. For every positive ε and ρ there is an integer m such that all the equilibria of the Bayesian games in Γ with more than m players are (ε, ρ) information proof.*

While all the examples given earlier may be modeled with the set U containing a single possible payoff function, the fact that U contains more functions

allows for the possibility of including several families of games at once. Actually, substantially richer families of games can be included, as shown by the lemma below that includes the theorem above as a special case.

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 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 than δ . Recall that every finite collection of continuous functions defined on the same compact domain is uniformly equicontinuous, and thus the above theorem is an immediate consequence of the following.

Lemma 30 *The conclusion of the theorem holds under the assumption that U consists of any collection of uniformly equicontinuous functions.*

Proof. Recall that an equilibrium σ is (ε, ρ) information proof if for some $\delta > 0$ the following two inequalities are satisfied.

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

Relying on the assumption that $|\mathcal{C}|$ is fixed for all the games in the collection, for any $\delta > 0$ all the equilibria of the games in Γ can be made to satisfy simultaneously the second inequality, provided that only games with more than some large number n of players are considered. So, it suffices to show that by making δ sufficiently small, and n sufficiently large, the first inequality can be simultaneously satisfied for all large games. The second term in the right hand side of this inequality can be made uniformly arbitrarily small, by choosing n to be sufficiently large, again for every given positive δ . Thus, it remains to be shown that $si[nbd(\exp_\sigma, \delta)]$ can be made uniformly smaller than any given positive ε by choosing a sufficiently small positive δ , and restricting attention to games with more than some large number of players n .

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 $c_1, c_2 \in nbd(\exp_\sigma, \delta)$ having $t_1 = t_2 = t_i$:

$$|[u_i(c_{-i}^1 : (t_i, a'_i)) + u_i(c_{-i}^2 : (t_i, a''_i))] - [u_i(c_{-i}^1 : (t_i, a''_i)) + u_i(c_{-i}^2 : (t_i, a'_i))]|.$$

But rearranging terms this expression equals

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

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 $u_i(c_{-i}^1 : (t_i, a'_i)) - u_i(c_{-i}^2 : (t_i, a'_i))$ can be made arbitrarily small by restricting attention to c 's in $nbd(\exp_\sigma, \delta)$. However, the uniform 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 . ■

6 Concluding Remarks

6.1 Subgame Perfection and Other Refinements

As discussed earlier, an information-proof equilibrium is order invariant in the sense that if the game is played sequentially, the equilibrium remains a Nash equilibrium also in the sequential game. However, the induced equilibrium in the sequential game is not necessarily subgame perfect. Consider for example the 2 person coordinated computer choice game discussed earlier with both players preferring to match, but with player 1 preferring to match on computer B and with the other player preferring to match on computer M, i.e., a "battle of the sexes" instead of a coordination game. As discussed earlier (B,B) is a Nash equilibrium outcome of the sequential game where player 2 moves first and player 1 follows. But the only subgame perfect equilibrium outcome of this game is (M,M). This is unavoidable. It is worth observing though, that the severity of the problem, or the level of violation of subgame perfection, may be diminished as the number of players increases.

Example 31 : *n*-Person Battle of the sexes. *There are n male players and n female players. Simultaneously, every player has to choose computer B or computer M. The payoff to a male who chooses B equals the proportion of the total population that he matches, while if he chooses M his payoff equals 0.9 times the proportion of the total population that he matches. For a female the opposite is true, i.e., she is paid the proportion that she matches when she chooses M and 0.9 of the proportion she matches when she chooses B.*

Consider the above game played sequentially, one player at a time, with the females moving first and the males following. Every player is informed of the choices made by all earlier players.

One can verify that the only subgame perfect Nash equilibrium results in everybody choosing M. So the constant B strategy, where everybody chooses B independent of history, is not subgame perfect. Nevertheless it is highly perfect in the following sense, see Kalai and Neme (1992) for the general measure of perfection that is applied below.

At a pure strategy subgame perfect equilibrium of a perfect information game, every action in every node of the tree is optimal given the actions taken at all other nodes. Consider now a pure strategy equilibrium which is not subgame perfect. Since it is a Nash equilibrium, at every node on the play path the selected action is optimal. Thus it takes some deviations from the play path and possibly from other equilibrium prescribed actions to reach an "irrational action," i.e., a node where the equilibrium action is suboptimal. Let the minimal number of deviations, i.e., equilibrium actions that are changed, needed to reach a suboptimal action be the level of perfection of the equilibrium. Thus, every Nash equilibrium is 0-perfect equilibrium, since with zero deviations no irrational action is reached, and every subgame perfect equilibrium is ∞ -perfect, since any number of deviations will fail to reach an irrational action.

And if an equilibrium is m -perfect it means that simultaneously, m deviations from the equilibrium strategies must take place to lead to an "irrational action."

As m becomes large the stability of m perfect equilibrium increases. Recall that a logic of subgame perfection is that if the game was diverted to a node where an irrational action is taken, then, when called upon to play, the irrational actor will change to a rational action. And if this would benefit the players who can guide the play to such a node they would choose to do so. But when m is large it requires the coordinated effort (or the cumulation of mistakes, depending on your interpretation of subgame perfection) of many players participating in many deviations to guide the play to such a node. And the faith of the first deviator that the others will follow may be diminished as m gets larger.

This is the case for the constant B equilibrium in the Battle of the Computers game. It yields every female a payoff of 0.9 and a female that deviates, 0. Would the first female deviate? Her successor, the second female, if she stays with the equilibrium strategy will be paid close to 0.9 and if she deviates, without the rest of the females following, will be paid close to zero. To make the deviation of the first female worthwhile she needs to count on successful deviation by at least 0.9 of the players that follow, including almost all the early females who share similar concerns.

The level of subgame perfection of the constant strategy B can be shown to be approximately n .

Related issues of refining and of weakening the Nash equilibrium concept can be found in recent economic literature. For example Dekel and Wolinsky (2000) study large auctions with independent types and show that strategy profiles that are only rationalizable already have strong implications regarding the outcome of the auction. On the other hand, Chung and Ely (2000) stay with Bayesian equilibrium as a solution concept, but strengthen the requirement on the ex-post stability. Rather than just being a Nash equilibrium they require that it be refined through the sequential elimination of dominated strategies.

6.2 Voting games: more general sufficient conditions.

The voting equilibrium below shows that even when a game is discontinuous in the sense of the previous section, it may still have equilibria that are highly information proof. This is obtainable from the general bounds defined earlier in the paper, which did not assume full continuity of the payoff functions. The second voting equilibrium below shows that one may obtain information proofness even when these general bounds are not applicable. In other words, the sufficient conditions for information proofness offered in this paper could possibly be made more general.

Consider a political game, where simultaneously n voters have to each select between two candidates, A and B. Suppose n is odd and the candidate with the majority vote is elected. A proportion p of the voters prefers A to B, say they each receive a utility one when A is elected and zero when B is elected, and a proportion $1 - p$ has the symmetrically opposite preferences.

The payoff functions of these games cannot be imbedded in a continuous payoff function as required in the previous section. At the point where the votes for the two candidates are split evenly there is a payoff discontinuity for all the voters. But if the prior probability of support for A is not in the middle, $p \neq 0.5$, then the equilibrium in which every voter votes for his preferred candidate is approximately information proof. With high probability the outcome will have close to a proportion p of the voters vote for A, and the remaining ones vote for B. Near these expected proportions, strategic interdependence is very low (actually zero due to the fact that no the outcome cannot be changed by a player unless the proportion becomes very close to 0.5), so by the bounds given in the earlier section the equilibrium must be $(0, \rho)$ information proof with a low value for ρ .

But even when $p = 0.5$, there is still a high degree of information proofness due to a different use of the law of large numbers. Consider the same equilibrium, with every player voting for his favorite candidate, so that the expected support is 0.5 for each candidate. This is exactly where the strategic interdependence is high even for small neighborhoods of this point. Yet it is clear that with high probability there will be some positive gap between the realized empirical proportions and the expected proportions. Moreover, if n is large, with high probability no one player should be able to close the gap in order to swing the outcome of the election in the opposite direction. Thus, the outcome is likely to stay information proof.

6.3 Large auctions: a continuum of types and actions

This paper exhibits a strong form of convergence to information-proof equilibrium. For given ε and ρ the convergence is at an exponential rate and it is so simultaneously for all games, all equilibria, and where a high proportion of the outcomes are information proof for all players simultaneously. To a significant degree this is due to the restriction of the games to have finite sets of possible types and actions. However, as the following discussion suggests, there are games where a model with a continuum of types and actions may be better.

Consider an n bidder first price sealed bid auction of an item whose private valuation by each bidder is a random integer from 1 to 100, generated by independent prior probability distributions. Simultaneously, each bidder submits a bid that consists of an integer from 1 to 100. The highest bidder wins the object and pays for it the amount of his bid. If several bidders tie at the highest bid one of them is selected at random to be the winner. The payoffs may be represented as follows. If bidder i is a highest bidder his payoff is $(v_i - b_i)/m$ with v_i being his valuation, b_i being his bid, and m being the number of highest bidders. Otherwise his payoff is zero.

The above payoff functions cannot be imbedded in anonymous g 's as required in the formulation of the previous section. More specifically, the above tie breaking rule requires knowledge of the number of players with the highest bid, but the functions g are expressed in the language of proportions. This means that no finite number of such g 's can describe simultaneously all the payoff

functions for auctions with increasing number of bidders.

On the other hand, it seems that a model with a continuum of types and actions, between 0 and 100, could possibly work here. Consider the strategies where every player bids his true value. When the number of players is large and the prior probabilities are sufficiently diffused this is an approximate Bayesian equilibrium that is approximately information proof. The probability of a tie at this equilibrium is zero and, due to the crowding of the bids, no player can improve much by changing his bid ex-post.

6.4 Repeated play and learning

In the context of rational Bayesian learning, being information proof can be interpreted to mean that there is no need for any learning. Consider an infinitely repeated Bayesian game, where each of n players is drawn an independent type that is used in an infinitely repeated play of a given one-shot n person game. According to a result of Kalai and Lehrer (1993), at every Bayesian equilibrium, where each player chooses his repeated game strategy based on his private knowledge of his type alone, with time the play converges to a Nash equilibrium of the complete information repeated game, as if the realized types are known to all.

Consider now an information-proof equilibrium of the one-shot Bayesian game. Because of information proofness the selected actions are a Nash equilibrium of the complete information game with the realized types being known to all. In other words, learning in the Kalai and Lehrer sense is automatic, even in the one-shot game without the need for any repetition.

More formally, starting with the information proof Bayesian equilibrium of the one-stage game, construct the following equilibrium for the repeated game. After being informed of his type, every player will choose a pure action as he does at the equilibrium of the one shot game, and will play this action repeatedly in the repeated game. It is easy to see that such strategies are a Nash equilibrium of the repeated Bayesian game, no matter what the discount parameter is, and that learning, in the Kalai and Lehrer sense, took place already in the first stage.

The above observation suggests an interesting direction of research, investigating a possible connection between the speed of learning and the number of players in the repeated game. On one hand, one would expect learning to be easier in a game with a small number of players, since in such a situation a player has to predict the behavior of a smaller number of opponents. On the other hand, at an information-proof equilibrium, which is obtained for continuous games with many players, learning the relevant information is immediate (or unnecessary). Is there a monotonicity in the ease of learning as the number of players become large beyond some critical level?

6.5 Information and action smallness

By assuming anonymity, continuity and many players, this paper implicitly deals with players whose actions and information have only a small effect on

their opponents. But as mentioned in the introduction there is a growing recent literature that deals explicitly with measurements of "action smallness" and "information smallness" and their effects on the play of the game. Moreover it seems reasonable to expect that action and information smallness imply low strategic interdependence and thus a high level of information proofness. In other words, by combining assumptions and results from the above literature one may be able to obtain approximate information proofness under more general conditions than the ones assumed in this paper.

7 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.
- Arrow K., and R. Radner (1979): "Allocations of Resources in Large Teams," *Econometrica* 47, pp 361-385.
- Chung, K.S, and J.C. Ely (2000): "Efficient and Dominance Solvable Auctions with Interdependent Valuations," Northwestern University working paper.
- Dekel, E., and A. Wolinsky (2000): "Rationalizable Outcomes of Large Independent Private-Value First-Price Discrete Auctions," Northwestern University Discussion Paper.
- Feddersen, T., and W. Pesendorfer (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica* 65 (5), pp 1029-58.
- Fudenberg, D., D. Levine, and W. Pesendorfer (1998): "When are Non-Anonymous Players Negligible," *Journal of Economic Theory*, 79, pp 46-71.
- 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, E.J. (1980): "Non-Cooperative Price Taking in Large Dynamic Markets," *Journal of Economic Theory*, 22, pp 155-182.
- Green, J.R., and J.J. Laffont (1987): "Posterior Implementability in a Two-person Decision Problem," *Econometrica*, 55 (1), pp. 69-94.
- Green, J.R., and J.J. Laffont (1989): "Posterior Implementability, Sequential Games and International Agreements," in *Cooperative Models in International Relations Research*, M. Intriligator and F. Interbacher (eds.).
- Groves, T., and S. Hart (1982): "Efficiency of Resource Allocation by Uninformed Demand," *Econometrica*, 50 (6), pp1453-82.
- Gul, F., and A. Postlewaite (1992): "Asymptotic Efficiency in Large Economies with Asymmetric Information," *Econometrica* 60, pp 1273-92.
- Hart, S., W. Hildenbrand, and E. Kohlberg (1974): "On Equilibrium Allocations as Distributions on the Commodity Space," *Journal of Mathematical*

Economics 1, pp 159-166.

Jordan, J.S., (1991): "Rational Learning in Normal Form Games," *Games and Economic Behavior*, 3, pp 60-81.

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

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

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

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

McLean, R., and A. Postlewaite (1999): "Informational Size and Incentive Compatibility," CARESS Working Paper #99-14, University of Pennsylvania.

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.

Radner, R., (1981): "Equilibrium Under Uncertainty", in K. J. Arrow and M. Intriligator, eds., *Handbook of Mathematical Economics*, North-Holland, Amsterdam, Ch. 20, pp. 923-1006.

Radner, R., (1991): "Intertemporal General Equilibrium," Ch. 15 of L. W. McKenzie and S. Zamagni (eds.), *Value and Capital Fifty Years Later*, Macmillan, London.

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.

Sabourian, H., (1990): "Anonymous Repeated Games with a Large Number of Players and Random Outcomes," *Journal of Economic Theory*, 51, pp 92-110.

Satterthwaite, M.A., (1999): "Strategy Proofness and Markets," Northwestern University Discussion Paper.

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