# LARGE ROBUST GAMES

EHUD KALAI

ABSTRACT. With many semi-anonymous players, the equilibria of simultaneous-move games are *extensively robust*. This means that the equilibria survive even if the simultaneous-play assumption is relaxed to allow for a large variety of extensive modifications. Such modification include sequential play with partial and differential revelation of information, commitments and multiple revisions of choices, cheap talk announcements and more.

## 1. INTRODUCTION AND EXAMPLES

1.1. **INTRODUCTION.** Games with many players is an old topic in economics and in cooperative game theory. Prominent theorists, including Arrow, Aumann, Debreu, Hildenbrand, Scarf, Shapley, and Shubik, have shown that in large cooperative games major modeling difficulties disappear. For example, solution concepts such as the core, competitive equilibrium, and the Shapley value, which in general predict different outcomes for the same game, predict the same outcome when the number of players is large. As a special case, this coincidence offers a cooperative-coalitional foundation for competitive equilibria. For a general survey, see Aumann and Shapley (1974).

Less is known about general non-cooperative strategic games with many players. An early study of this subject was Schmeidler (1973), who shows the existence of pure-strategy Nash equilibria in normal-form games with a continuum of anonymous players.

More recently there have been many studies of specific large economic games; see, for example, Mailath and Postlewaite (1990) on bargaining, Rustichini, Satterthwaite, and Williams (1994) and Pesendorfer and Swinkels (1997) on auctions,

---

*Date*: February 5, 2004.

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

The author is the James J. O'Connor Distinguished Professor of Decision and Game Sciences at the Kellogg School of Management. Early versions of this paper benefited from presentations at many universities and conferences, including the 2001 Pazner Lecture, Tel Aviv University, and the 2000 Pareto Lecture of the Association of Southern European Economic Theorists. In particular, the author thanks Eddie Dekel, Olivier Gossner, Sergiu Hart and Adam Kalai for helpful conversations. Partial support for this research was made possible by the National Science Foundation, grant No. SES-9904931.

This version replaces the following earlier papers of the author published by the Center for Mathematical Studies in Economics and Management Science: "Private Information in Large Games" (2000), "Purification and Ex-Post Stability in Large Games" (2003), and "Large Robust Games" (2003).

and Feddersen and Pesendorfer (1997) on voting. For the most part, these papers concentrate on issues of economic efficiency.

Other aspects of large games are explored in the literature on repeated large games; see, for example, Green (1980) and Sabourian (1990)[1], where the focus is on short-term behavior exhibited by patient players, and Fudenberg and Levine (1988), who study the relationship between open-loop and closed-loop equilibria.

The objective of the current paper is to uncover properties of *general* strategic games with many players, beyond the one identified by Schmeidler. The results obtained parallel the ones obtained in cooperative game theory, in that they help overcome modeling difficulties and offer foundations for other economic concepts.

A particular modeling difficulty of non-cooperative game theory is the sensitivity of Nash equilibrium to the rules of the game, e.g., the order of players' moves and the information structure. Since such details are often not available to the modeler or even to the players of the game, equilibrium prediction may be unreliable. This paper demonstrates that this difficulty is less severe in general classes of games that involve many semi-anonymous players. In normal-form games and in simultaneous one-move Bayesian games with independent types and continuous and anonymous payoff functions, all the equilibria become *extensively robust* as the number of players increases. This is a new notion of robustness, different from other robustness notions used in economics and in game theory[2].

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 simultaneous-move game. Such versions allow for wide flexibility in the order of players' moves, as well as for 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 simultaneous one-move games.

Extensive robustness means in particular that an equilibrium must be *ex-post Nash*. Even with perfect hindsight knowledge of the types and selected actions of all his opponents, no player regrets, or has an incentive to revise, his own selected action. Similar notions have been extensively discussed in the implementation literature; see for example Cremer and McLean (1985), Green and Laffont (1987), and Wilson (1987).

The ex-post Nash and extensive robustness properties relate to some important issues in economic applications. Among them are the facts that being ex-post Nash offers a stronger (but asymptotic) version of the purification result of Schmeidler mentioned above, and that the Nash equilibria of certain large market games have a strong (extensive) rational-expectations property when the players' types are drawn independently. We delay discussion of these and other issues to later in the paper.

The proof of extensive robustness is in two steps. First we show that the equilibria of the games under consideration become ex-post Nash at an exponential rate, and then we show that this implies full extensive robustness. In order to prove the ex-post Nash part, we first develop a notion of strategic interdependence in a set of possible outcomes of the game. Strategic interdependence measures how a player's (ex-post) regret for choosing one action over another may vary as

[1]See Al-Najjar and Smorodinsky (2000) for more recent results and references.

[2]See for example Hansen and Sargent (2001), and Kovalenkov and Wooders (2001).

a function of opponents' choices. The continuity condition imposed on our games guarantees that strategic interdependence is low in small sets of outcomes near the expected play, where, by a law of large numbers, the outcomes must reside.

1.2. **EXAMPLES.** We illustrate some of the ideas of this paper with two examples. The first is a normal-form game where a mixed-strategy equilibrium becomes extensively robust, but only as the number of players becomes large. (In normal-form games, regardless of the number of players, all pure-strategy equilibria are extensively robust.) The second example is a Bayesian game with a *pure-strategy* (even strict) Nash equilibrium that becomes extensively robust, but only as the number of players becomes large.

**Example 1.** ***Village Versus Beach (2n players match pennies):***

*Each of the* $2n$ *players,* $n$ *females and* $n$ *males, have to choose between going to the beach (B) or staying in the village (V). A female's payoff is the proportion of males her choice* ***mis****matches and a male's payoff is the proportion of females his choice matches. Full randomization, with everybody choosing B or V with equal probabilities, is an equilibrium of the simultaneous-move game. For* $n = 1$*, the match-pennies game, it is the only equilibrium.*

For a small number of players, say, $n = 1$, the equilibrium is highly non-robust. One illustration of this is the lack of the ex-post Nash property. There is a high probability (certainty in this example) that after performing their randomizations and seeing both realized choices, one of the players would want to revise his/her choice. Similarly, sequential robustness fails, i.e., different orders of play will yield different outcomes. Actually each player would wish to wait and see the choice of the opponent before making his/her own choice. It is difficult to offer a reasonable equilibrium prediction for this situation.

What about large values of $n$? Now, laws of large numbers imply that the proportion of players of each gender at the beach will be close to a half. Thus, ex-post, no player could significantly improve his/her payoff by revising his/her realized pure choice. Moreover, as the results of this paper indicate, for every given acceptable range near the .50-.50 distribution, say, .499 to .501, the probability of either gender's empirical distribution falling outside this range goes down to zero at an exponential rate as the number of players increases.

Beyond ex-post stability, as we shall see, the randomizing equilibrium is fully extensively robust. Making the choices sequentially, observing players in bathing suits, counting busses on the way to the beach, controlling others' ability to make choices, etc., will give no player significant incentive to unilaterally deviate from his/her randomizing strategy.

**Example 2.** ***Computer-Choice Game:***

*Simultaneously, each of* $n$ *players has to choose computer I or computer M, and, independently of the opponents, each is equally likely to 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 she likes. Specifically, each player's payoff is 0.1 if she chooses her favorite computer, zero otherwise, plus 0.9 times the proportion of opponents her choice matches. Assuming that each player knows only her 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 where every player chooses her favorite computer.*

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 her predecessors, then everybody choosing I regardless of the observed history is a Nash (not subgame-perfect[3]) equilibrium of the extensive game. And they are robust to other modifications. For example, if a round of revision were allowed after the players observe the opponents' first-round choices, everybody choosing I with no revision remains an equilibrium of the two-round game. In other words, they are ex-post robust.

This is not the case when the players use the choose-your-favorite-computer strategies. For example, if the population consists of only two players and the game is played sequentially, there are positive-probability histories after which the follower is better off matching her predecessor than choosing her favorite computer. And in the revision game with two rounds, there is a significant probability of players revising in the second round.

As users of game theory know, this sensitivity to the order of moves creates modeling difficulties, since we do not know in what order players think about, rent, or buy computers. Also the real life situation may allow for other possibilities. For example, the players may make announcements prior to making their choices, repeatedly revise earlier choices after observing opponents' choices, make binding commitments, reveal or hide information, etc., and every such possibility may destabilize the equilibria of the simultaneous-choice game.

But the modeling difficulties become less severe if the number of players is large. In this case, even choosing-one's-favorite-computer is a highly robust equilibrium, i.e., it remains an approximate Nash equilibrium in the sequential and in all other extensive versions of the game. Such versions accommodate all the variations mentioned above.

Moreover, the above robustness property is not restricted to the equilibrium of choosing-one's-favorite-computer. Every Nash equilibrium of the one-shot game is extensively robust, and this is true even if the original computer-choice game with which we started is more complex and highly non-symmetric: There may be any finite number of computer choices and player types. Different players may have different (arbitrary) payoff functions and different (arbitrary) prior probability distributions by which their types are drawn. Players' payoff functions may take into consideration not just opponents' actions, but also opponents' types (e.g., a player may have a positive payoff when others envy his choice). Regardless of such specifications, all the equilibria of the one-shot game are highly extensively-robust when $n$ is large.

## 2. THE MODEL

### 2.1. THE MODEL AND GENERAL DEFINITIONS.

Two *finite* non-empty abstract universal sets, $\mathcal{T}$ and $\mathcal{A}$, respectively describe all possible player *types* and all possible player *actions,* that may be involved in the games considered in this

---

[3]We elaborate on the issue of subgame perfection in the concluding section of the paper.

paper. For notational efficiency, a universal set $\mathcal{K} \equiv \mathcal{T} \times \mathcal{A}$ denotes all possible *type-action characters* of players.[4]

Throughout the rest of the paper we consider a general *family* $\Gamma = \Gamma(\mathcal{T}, \mathcal{A})$ that consists of Bayesian games, each described by a five-tuple $G = (N, T, \tau, A, u)$ as follows.

$N = \{1, 2, ..., n\}$, the set of *players*, is defined for some positive integer $n$.

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

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

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

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

In addition to the above, standard game-theoretic conventions are used throughout the paper. For example, for a vector $x = (x_1, x_2, ..., x_n)$ and an element $x'_i$, $x_{-i} = (x_1, ..., x_{i-1}, x_{i+1}, ..., x_n)$ and $(x_{i-1} : x'_i) = (x_1, ..., x_{i-1}, x'_i, x_{i+1}, ..., x_n)$. Also, while a profile of type-action characters $c = ((t_1, a_1), ..., (t_n, a_n))$, we sometimes describe it as a pair of profiles $c = (t, a)$ in the obvious way.

A Bayesian game $G$ is played as follows. In an initial stage, *independently* of one another, 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 (with the possible aid of a randomization device), thus defining his type-action character. Following this, the players are paid (according to their individual utility functions) the payoffs computed at the realized profile of type-action characters.

Accordingly, a *strategy* of player $i$ is defined by a vector $\sigma_i$ with $\sigma_i(a_i \mid t_i)$ describing 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 feasible type-action characters, $\gamma_i(c_i) = \tau_i(t_i) \times \sigma_i(a_i \mid t_i)$. Under the independence assumption, the profile of these distributions, $\gamma = (\gamma_1, \gamma_2, ..., \gamma_n)$, determines the overall probability distribution over *outcomes of the game* (i.e., type-action character profiles), by $\Pr(c) = \prod_i \gamma_i(c_i)$.

Through the use of mathematical 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 $\sigma$ is a (Bayesian) *Nash equilibrium* if for every player $i$ and every one of his strategies $\sigma'_i$, $u_i(\sigma) \geq u_i(\sigma_{-i} : \sigma'_i)$.

The above definitions are standard, and, except for the assumption that players' types are independent, a game is general in that a player's payoff may depend on other players' actions *and types*. To accommodate this generality it is useful to

[4]For example, famous characters in game theory are *weak types who eat quiche, weak types who drink beer,* etc. In this paper we have *characters who like I and choose I*, *characters who like I and choose M*, *characters who are informed about I and choose M*, etc.

introduce the notion of player's type-action character, as defined above, so that payoff functions and probability distributions are defined on a notationally simple space.

The family of games described above may be quite large. First, it may contain games with varying number of players $n = 1, 2, \ldots$ . But, in addition, for any fixed number of players $n$, it may contain infinitely many different $n$-player games. There are no restrictions (other than normalization) on the payoff functions and the prior probability distributions.

The restrictions that the $T_i$'s of the various games must all be subsets of the same finite set $\mathcal{T}$ and that the $A_i$'s of the various games must all be subsets of the same finite set $\mathcal{A}$ are substantial, and are needed for the method of proof we use. Thus, one may think of the current results as being relevant to games in which the number of players is large relative to the number of types and actions.

We also need to restrict ourselves to families of games that satisfy semi-anonymity and continuity conditions defined as follows.

**Definition 1.** ***Empirical distribution:*** *For every vector of type-action characters $c$, define the* empirical distribution *induced by $c$ on the universal set of type-action characters $\mathcal{K}$ by $emp_c(\kappa) =$*

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

**Definition 2.** ***Semi-anonymity****: The games of $\Gamma$ are semi-anonymous if for every game $G$, for every player $i$ in $G$, and for any two profiles of type-action characters $c$ and $\overline{c}$, $u_i(c) = u_i(\overline{c})$ whenever $c_i = \overline{c}_i$ and $emp_{c_{-i}} = emp_{\overline{c}_{-i}}$.*

*When this is the case, we may abuse notations and write $u_i(c_i, emp_{c_{-i}})$ instead of $u_i(c_1, c_2, ..., c_n)$.*

**Definition 3.** ***Continuity****: The payoff functions in the family of games $\Gamma$ are uniformly equicontinuous if for every positive $\varepsilon$, there is a positive $\delta$ with the following property: For every game in the family, for every player $i$ in the game, and for any two profiles of type-action characters $c$ and $\overline{c}$, $\mid u_i(c) - u_i(\overline{c}) \mid < \varepsilon$ whenever $c_i = \overline{c}_i$ and $\max_{k \in \mathcal{K}} \mid emp_{c_{-i}}(\kappa) - emp_{\overline{c}_{-i}}(\kappa) \mid < \delta$.*

Note that, technically, the continuity condition, as just stated, already implies that all the games in $\Gamma$ are semi-anonymous.

2.2. **CLARIFYING REMARKS ON THE MODEL[5].** The condition of semi-anonymity is less restrictive than may appear, since it imposes anonymity only on the payoff functions but with no further restrictions of symmetry or anonymity on the players. This means that information about named players can be incorporated into their types.

**Example 3.** ***Sellers and Buyers****:*[6]

*There are $n$ sellers, labeled as players $1, 2, ..., n$, and $n$ buyers, labeled as players $n+1, n+2, ..., 2n$. The payoff function of a seller depends entirely on his own strategy and on the empirical distribution of the strategies of the buyers, while the payoff function of a buyer depends on his own strategy and on the empirical distribution of the strategies of the $n$ sellers. In violation of the assumption of our model, the payoff functions are not anonymous. For example, a seller's payoff*

[5]A reader eager to get to the main results may choose to skip this section at first reading.

[6]For recent results and additional references on large market games, see Ellison and Fudenberg (2002).

*function treats players n+1,...,2n (the buyers) differently than the players labeled 1,2,...,n (the sellers). But if within each group the players are symmetric, we can describe the situation by a new Bayesian game in which the payoff functions are anonymous, as follows:*

*Allow each player to be of two possible types, a* seller *or a* buyer. *Assign probability one of being a seller type to players 1,...,n, and probability one of being a buyer type to players n+1,...,2n. Now we can write the payoff function of a seller, in the obvious, way to depend on the empirical distribution of types (i.e., to depend only on the players that are of the buyer type) and actions, without having to specify player labels. Since the same can be done to the payoff functions of the buyers, the game is semi-anonymous. Clearly, this description is possible because the model imposes no symmetry or anonymity on the prior distributions by which types are drawn.*

As in the example above, the model can accommodate many non-symmetric games. Players may play a variety of other roles, in addition to sellers and buyers as above. Also, they may be identified as belonging to different genders (as in our Village versus Beach example), geographical locations, or social or professional groups. The "Match The Expert" example below illustrates how even named individuals can be incorporated into a semi-anonymous model. However, the assumption of finitely many types does restrict the generality of such descriptions. Without it, semi-anonymity would represent no restriction at all.

**Example 4.** ***Random Number of Players:***

*In the Sellers and Buyers game above, extend the set of possible types to include a non-participating type. Assign each player 1,...,n positive probabilities of being a seller or a non-participating type, and assign each player n+1,...,2n positive probabilities of being a buyer or a non-participating type. Moreover, restrict the payoffs of seller types to depend only on their own strategy and the strategies of the buyer types (leaving out the non-participating types), and make the symmetric restriction on the buyer types. The result is a semi-anonymous Bayesian game with a* random number *of sellers and buyers.*

The above modeling method may be applied to a large variety of games with random number of players and with differential information about the composition of players in the game.

The next example shows how the uniform equicontinuity condition may fail.

**Example 5.** ***Match the Expert:***

*The game has n players, each having to choose action A or B, and there are two equally likely states of nature, a and b. The payoff of every player is one if he chooses the appropriate action - A in a and B in b - but zero otherwise. Player 1, who is the only informed player, is told what the realized state of nature is. Thus, we may think of her as having two types: a type who* is told *a and a type who* is told *b. All the other players are of one type that knows nothing.*

*Consider the equilibrium where Player 1 chooses the action corresponding to her type and everybody else chooses the two actions with probabilities* $.50 - .50$. *Obviously this equilibrium is not ex-post Nash, and hence, not extensively robust. Actually, it becomes less ex-post stable as the number of players n increases; with probability close to 1, approximately 1/2 of the players would want to revise their choice after observing the choice of Player 1.*

*Even though the ex-post Nash condition fails, we can formulate the above as a semi-anonymous game by identifying three possible types:* an expert informed of state a, an expert informed of state b, *and* a non-expert. *Assign Player 1 equal probability of being one of the first two types and every other player probability 1 of being of the third type. The payoff of every player may then be defined anonymously as follows. Player 1's payoff is one if she chooses according to her type and zero otherwise; every other player's payoff is one if his choice matches the information given to a strictly positive fraction of the realized expert types,*[7] *zero otherwise.*

*The failing of the ex-post Nash condition is due to the failing of uniform equicontinuity. To see this, consider two sequences of type-action character profiles* $c^n$ *and* $\overline{c}^n$ *defined as follows. In both profiles all players choose the action a. However, in* $c^n$ *Player 1 is given the information a and in* $\overline{c}^n$ *she is given the information b. Consider, for example, the sequence of payoff functions of players 2,* $u_2^n$*, in the n person games as* $n \to \infty$*. For all n,* $u_2^n(c^n) = 1$ *and* $u_2^n(\overline{c}^n) = 0$*, despite the facts that* $c_2^n = \overline{c}_2^n$ *and that the empirical distributions of player 2's opponents in the two type-action character profiles become arbitrarily close as* $n \to \infty$*.*

The uniform equicontinuity is the important condition that ties together the different games in the family. We do not restrict ourselves to replicating games, so that for each number of players n, the family $\Gamma$ may contain many different types of games (e.g., a market game with 100 players, a political game with 100 players, an auction with 101 players, etc.). Yet uniform equicontinuity is sufficiently strong to guarantee the asymptotic result as we increase the number of players in the variety of games under consideration.

The next example shows that without statistical independence of types, the ex-post Nash property may fail.

**Example 6.** ***Two States of the World with Dependent Types:***

*Consider, as above, a simultaneous-move n-player game where each player has to choose one of two actions, A or B, and there are two possible states of the world, a and b. Also as before, assume that every player's payoff is one if he chooses the action that corresponds to the (unknown) realized state of the world, zero otherwise. But assume now that every player is given a less than perfect signal about the realized state as follows. For every realized state, independently of one another, every player is told the correct (realized) state with probability .90 and the incorrect state with probability .10.*

*It is easy to see that every player choosing the action that corresponds to the state he was told is an equilibrium (even in dominant strategies). It is also clear that ex-post stability fails. When n is large, approximately 10% of the players will observe ex-post that theirs is the minority choice, and would want to revise.*

*Here, ex-post stability fails because of the dependency in the prior type probabilities (before conditioning on the state).*

## 3. EX-POST STABILITY AND PURIFICATION[8]

3.1. **EX-POST NASH.** A profile of type-action characters is ex-post Nash if, with full knowledge of the profile (i.e., types and selected actions of all players), no

[7]This is an indirect way of saying that the player's selected action matches the information of player 1, without having to mention player 1 by name so as to keep the payoff function anonymous.

[8]A reader eager to get to the extensive-robustness result may skip this section at first reading.

player has the incentive to unilaterally change his selected action. Alternatively, the vector of actions described by the profile is a Nash equilibrium of the complete information game determined by the corresponding profile of types. A strategy is ex-post Nash if it leads, with probability one, to profiles of type-action characters that have the above ex-post Nash property. This is a strong notion of the ex-post Nash property, since it requires that the realized pure actions, not the mixed strategies, constitute a Nash equilibrium of the complete information game with all the realized types being known.

To gain some understanding, we first consider a two-person normal-form game (only one type for each player) and a mixed-strategy equilibrium, as illustrated in the following payoff table:

| | .30 | .70 | 0 | 0 |
|---|---|---|---|---|
| .20 | **6**, **7** | **5**, **7** | 5,4 | 1, 2 |
| .30 | **6**, **5** | **5**, **5** | 8,1 | 9, 2 |
| .50 | **6**, **8** | **5**, **8** | 9,7 | 5, 4 |
| 0 | 2, 2 | 1, 4 | 0, 1 | 9, 8 |

This is an ex-post Nash equilibrium, since with probability one it leads to one of the six bold-faced entries, and each one of these is a pure-strategy Nash equilibrium of the game. Note that this implies that the pure-strategy profiles, in the support of the mixed strategies profile, must have a very special structure. For example, Player 1 must have identical payoffs in his first three rows, when the opponent is restricted to her first two columns. This implies that each pair of pure strategies in the support of the mixed strategies is interchangeable in the sense of Nash (see Luce and Raiffa [1957]). Alternatively, when Player 1 best responds to either column 1 or column 2, he automatically best responds to the other. Also, rows 1,2, and 3 of Player 1 constitute dominant strategies for Player 1, provided that Player 2 is restricted to one of her first two columns. Naturally, the symmetric properties hold for Player 2.

It follows that at an ex-post Nash equilibrium, the use of private information is somewhat decentralized and highly simplified. A player can simply compute his optimal action with respect to any of the positive-probability type-action character profiles of his opponents, and this action is guaranteed to be optimal with respect to any other positive-probability type-action character profile of his opponents. In other words, in deciding on one's own optimal action, the importance of information about opponents' realized types and selected actions is highly reduced. Additional discussion of the above properties and their generalizations to Bayesian games is postponed to later sections of the paper.

Since we prove the ex-post Nash property asymptotically as the number of players increases, we need to define a notion of *approximately ex-post Nash.*

**Definition 4.** ***Approximately ex-post Nash****: Let $\varepsilon$ and $\rho$ be positive numbers.*

*A profile of type-action characters $c = (c_1, ..., c_n) = ((t_1, a_1), ..., (t_n, a_n))$ is $\varepsilon$ best-response 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 type-action characters is $\varepsilon$ Nash if it is $\varepsilon$ best-response for* every *player.*

*A strategy profile $\sigma$ is $(\varepsilon, \rho)$* ex-post Nash *if the probability that it yields an $\varepsilon$ Nash profile of type-action characters is at least $1 - \rho$ .*

Stated differently, "$\sigma$ is $(\varepsilon, \rho)$ ex-post Nash" means that with probability $1-\rho$ it must produce a profile of actions that are $\varepsilon$ Nash equilibrium of the normal-form game determined by the realized profile of types.

**Theorem 1.** ***Ex-post Nash:*** *Consider a family of Bayesian games $\Gamma(\mathcal{T}, \mathcal{A})$ with continuous and anonymous payoff functions as above and a positive number $\varepsilon$, there are positive constants $\alpha$ $(= \alpha(\Gamma, \varepsilon))$ and $\beta$ $(= \beta(\Gamma, \varepsilon))$, $\boldsymbol{\beta < 1}$, such that, for every $m$, all the equilibria of games in $\Gamma$ with $m$ or more players are $(\varepsilon, \alpha\beta^m)$ ex-post Nash.*

**Remark 1.** *An equilibrium may be $(\varepsilon, \rho)$ ex-post Nash in the "strong" sense discussed above, yet fail to be so in a "weaker" sense. It may be that if only partial ex-post information about the opponents' actions and types is revealed, then revision possibilities may become more attractive with higher probability. The full extensive robustness result discussed later assures us that the ex-post stability theorem above continues to holds in all such partial senses as well.*

3.2. **PURIFICATION.** An immediate consequence of the ex-post Nash property obtained above is a *purification property* in large games. First, for normal-form games the ex-post Nash property provides stronger conclusions than Schmeidler's (1973) on the role of pure-strategy equilibria in large anonymous games.

Working in the limit with a continuum of players, Schmeidler shows that every "mixed-strategy"[9] equilibrium may be "purified." This means that for any mixed-strategy equilibrium one can construct a pure-strategy equilibrium with the same individual payoffs. A large follow-up literature on Schmeidler's result is surveyed in Ali Khan and Sun (2002); more recent results and references may be found in Cartwright and Wooders (2003).

The ex-post Nash theorem stated above shows (asymptotically) that in large semi-anonymous games there is no need to purify, since it is done for us automatically by the laws of large numbers. In the limit, as the number of players increases, any mixed-strategy equilibrium must yield, with probability one, pure-strategy profiles that are Nash equilibria of the game.[10] So every mixed strategy may be thought of as a "self-purifying device."

Going beyond Schmeidler's model, and with the stronger property of self-purification, the theorem above shows that the phenomenon holds not just for normal-form games, but also for Bayesian games. In such games, asymptotically, every (pure or) mixed-strategy equilibrium produces, with probability one, pure-action profiles that are Nash equilibria of the complete-information game determined by the profile of realized types. In other words, the equilibrium self-purifies to equilibria of the randomly produced games. Further elaboration on this phenomenon is postponed to future papers.

---

[9]As Schmeidler points out in his paper, it is difficult to define a "real mixed strategy" equilibrium due to failings of laws of large numbers in the case of continuously many random variables.

[10]More precisely, for arbitrarily small $\varepsilon > 0$ there is a critical number of players $n$, such that for every mixed-strategy equilibrium of a game with $n$ or more players, there is probability of at least $1-\varepsilon$ that the realized vector of pure strategies is an $\varepsilon$ Nash equilibrium of the normal form game.

## 4. EXTENSIVE ROBUSTNESS

### 4.1. EXAMPLES AND MOTIVATION[11].

In the sections that follow, we define an equilibrium of a Bayesian game $G$ to be robust if it remains an equilibrium in every extensive version of the simultaneous-move game. As explained below, this means that playing *constant-play versions* of the (simultaneous-move) equilibrium strategies is an equilibrium in every extensive version of the game. A simple and important special type of an extensive version of a game $G$ is the following:

**Definition 5.** ***The game with one round of revisions:*** *In the first round of this two-round game, the original game $G$ is played. In a second round, the information about the realized types and selected pure actions of all the players becomes common knowledge; then, simultaneously, every player has the one-time opportunity to revise his first-round choice to any of his feasible $G$ actions. The players' payoffs are determined by the profile of final choices according to the payoff functions of $G$.*

In the above game, playing the constant-play versions of the equilibrium strategies means that in the first round the players choose their actions with the same probabilities as they do in the equilibrium of $G$. In the second round, every player chooses his action to be the same (hence, constant play) as the realized action of his first round, i.e., no revisions. Clearly, any strategy of the simultaneous-move game has a constant-play version in the game with revisions.

Notice that an equilibrium is ex-post Nash as discussed in the previous section, if and only if its constant-play version is an equilibrium in the game with the one round of revisions. This means that being ex-post Nash is a weaker condition than extensive robustness.

In various applications that use ex-post Nash conditions (for example, Green and Laffont [1987]), the information revealed between rounds may be only partial; for example, the players may learn their opponents' choices but not the opponents' types.

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. Here again, every player in his turn can play just as he would in the simultaneous-move game. Constant play is meaningless here, since every player moves only once.

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 which other players, various commitment possibilities may be available, cheap talk announcements may be made, and players may have control over the continuation game.

Basically, an extensive version of a simultaneous-move game $G$ is any extensive game in which (1) every complete play path generates a type-action character profile of the simultaneous-move game, and (2) all constant-play versions of the simultaneous-move strategies are possible.

[11] Again, the reader interested in the formal presentation alone, may proceed to the next section after reading the following definition of a game with one round of revisions.

Below are some simple examples of extensive versions of a game with two players who have to choose one of two computers. In all three extensive versions of Figure 1, nature moves first, choosing the (only) central arc with probability one, resulting in Player 1 being a type who likes I (i) and Player 2 being a type who likes M (m). Extensive version 1 is simply the simultaneous-move game, while Extensive Version 2 describes the sequential game where Player 2 makes her choice after being fully informed of the choice of Player 1.

In the third extensive version, as in the second, Player 2 follows with full knowledge of Player 1's choice. She can choose the same computer as Player 1, or she can choose the other computer. But when she chooses the other computer, she can do so in two ways: (1) she can do so and end the game, or (2) she can choose the central arc, in which case Player 1 will be informed of her contradictory choice and be offered the opportunity to revise. Notice that in the circumstances where Player 1 plays twice, the final outcome is determined by his last choice, since repeated choices represent revisions.

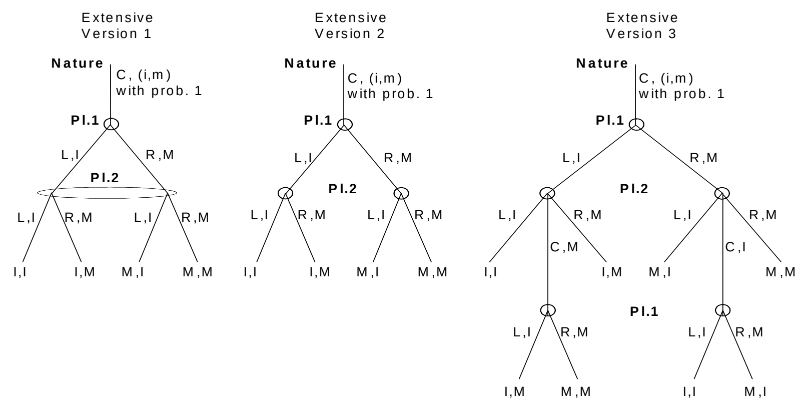


FIGURE 1

Suppose, for example, that in the one simultaneous-move game each player chooses one of the two computers with equal probabilities; then Extensive Version 3 permits a multiplicity of constant-play versions. Such versions require that at his first information set Player 1 choose L,I and R,M with equal probabilities. Player 2 must assign probability of .5 to L,I at her left information set and .5 to R,M at her right information set, with the remaining two arcs in each information set being assigned any probabilities sum to .5. Finally, Player 1 should not revise his computer choice in his two second information sets, i.e., he should choose L,I in the left one, and R,M in the right one.

Extensive versions of a game allow for incomplete information, beyond the types allowed for in the simultaneous-move Bayesian game. For example, in Figure 2, when Player 1 makes his initial selection, he does not know which of the three games in Figure 1 he is playing, but Player 2 does.

Possibilities of commitment are illustrated in Figure 3. By going right, Player 1 commits to choosing M; by going left, he reserves the right to revise his choice after observing Player 2's choice; and by playing center, he commits to matching the choice made by Player 2.

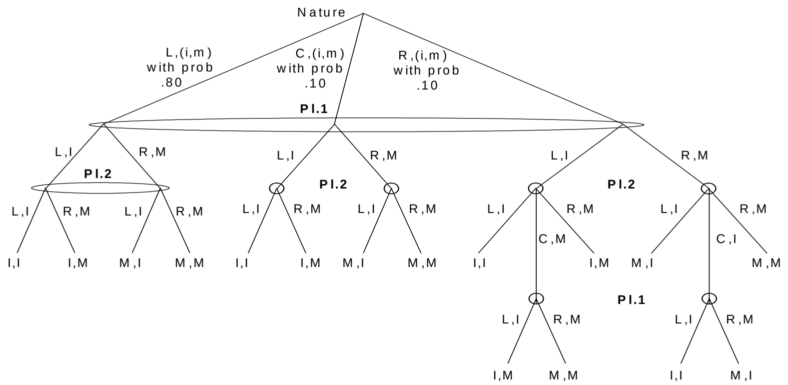


FIGURE 2

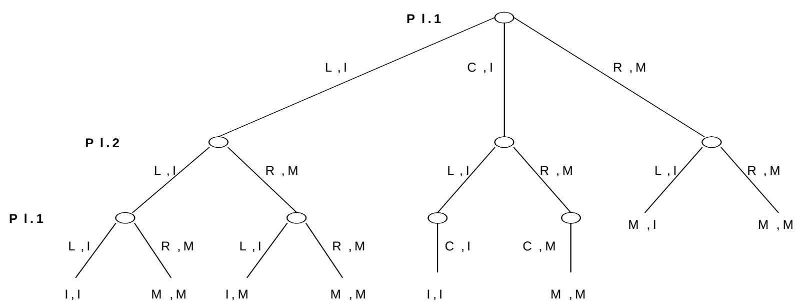


FIGURE 3

However, the game of Figure 3 will be ruled out in our model, since the middle part, where Player 1 commits to matching the opponent's choice, may *force* Player 1 to revise his earlier choice, interfering with his ability to perform constant play.

A reader familiar with the rich modeling possibilities of extensive games should be able to imagine the vast number of possible extensive versions that one can describe for any given simultaneous-move game.

**Remark 2.** ***On the complexity of extensive versions.***

*Real-life interaction that involves any reasonable level of uncertainty is most often too complex to describe by an extensive game, even when the number of players is moderate. Consider, for example, the game of choosing computers. For any other pair of players, a player must know who among them moves first. If he does not, he must have two trees, allowing for both possibilities, with information sets that link the two games to indicate his ignorance on this question. Since most players will have no such information about other pairs, we already have a very large number of trees linked by complex information sets.*

*But the situation is even worse since the extensive game is supposed to answer* all *informational questions. For example, Player A most likely does not know whether Player B knows some things (e.g., what computer Player C chose, or who plays first, D or E). Player A must then have different trees that describe every such possibility regarding the information of B, and again must link these trees with his information* *sets. This is already bad enough, but it goes on to information about information about information, etc., and the product of possibilities is enormous. The only way to stop this exponentially growing complexity is to make the answers to almost all questions common knowledge at a low level in the hierarchy of knowledge, which is not likely in real-life situations that involve even a moderate number of players.*

*Because of this modeling difficulty, it is essential to have strategies and equilibria that are highly not sensitive to the details of the tree, as is done in the sequel.*

4.2. **EXTENSIVE VERSIONS OF A GAME AND STRATEGIES.** The following abstract definition of an extensive version accommodates the modifications discussed earlier and more. Starting with the given simultaneous-move Bayesian game $G = (N, T, \tau, A, u)$, define an *extensive version* of G to be any finite perfect-recall Kuhn-type extensive-form game $\overline{G}$ constructed and defined as follows.

First, **augment G** with any finite non-empty set of abstract moves, $M$. $M$ may include words such as L, C, and R, as in Figures 1-3 above, and various (cheap) announcements, signals, etc.

**The initial node** in the game tree belongs to nature, with the outgoing arcs being labeled by the elements of $M \times T$. Thus, at the initial stage (as in the original game), nature chooses a profile of types. But in addition, it chooses an abstract move that may lead to different continuation games. 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 $\tau$ of the underlying game $G$.

**Every other node** in the game tree belongs to one of the players $i$, with the outgoing arcs labeled by elements of $M \times A_i$.

At **every information set** the active player $i$ has, at a minimum, complete knowledge of his own type $t_i$ (i.e., 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 one of the information sets of every player at least once. This guarantees that a player in the game chooses an action $a_i \in A_i$ at least once.

**The resulting type-action character profile** associated with a complete path in the game tree is $c = (t, a)$, with $t$ being the profile of types selected by nature at the initial arc of the path, and with each $a_i$ being the *last* action taken by player $i$ in the play path. The last action is the one that counts 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 $G$, computed at the resulting type-action character profile of the path.

In addition to the above, in this paper we restrict ourselves to extensive versions that satisfy the following condition:

**All constant-play versions of strategies of $G$ can be played**. For every player $i$:

1. All his G actions are feasible at any one of his *initial* information sets. More precisely, for any initial information set of Player $i$ and every action $a_i \in A_i$, at least one of the arcs leading out of the initial information set is labeled by $a_i$.

2. It is possible not to revise. More precisely, if at an information set $X$ Player $i$ selects an outgoing arc labeled with an action $a_i \in A_i$, then at any of his next

information sets (that follow that arc), $Y$, it is possible for him to select $a_i$ again, i.e., one of the outgoing arcs at $Y$ is labeled with the action $a_i$ (unlike in Figure 3).

**Remark 3.** *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.*[12] *But it also means that the version that is being played may be random, reflecting possible uncertainties about the real-life version in the mind of either the modeler or the players, as in Figure 2.*

*Similarly, a greater generality is obtained by including abstract moves, in addition to selected actions, at the nodes of the players. For example, using such moves, we can model a player's choice to reveal information, to seek information, to make cheap talk announcements, and to affect the continuation game in other ways. Figure 4 illustrates an extensive version where Player 1 can make cheap talk announcements that do not correspond to his real choices.*

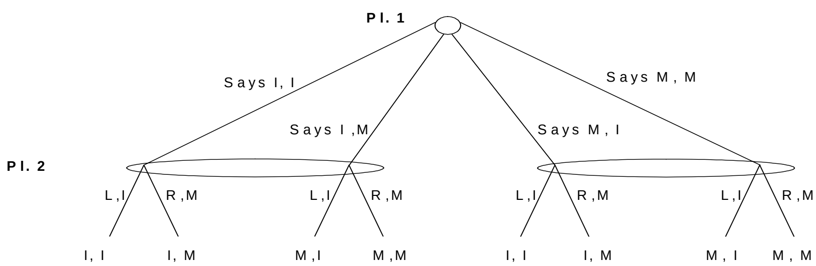


FIGURE 4

**Definition 6.** ***Constant-play strategies:*** *Given an individual strategy $\sigma_i$ in a Bayesian game $G$ and an extensive version $\overline{G}$, a* constant-play version *of $\sigma_i$ is any strategy $\overline{\sigma}_i$ in $\overline{G}$ that initially chooses actions with the same probabilities as $\sigma_i$ 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 $\sigma_i(a_i \mid t_i)$,*[13] *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 him in his previous information set. (This is well defined under the perfect-recall assumption.)*

*$\overline{\sigma} = (\overline{\sigma}_1, ..., \overline{\sigma}_n)$ is a profile of constant-play strategies if each of its individual strategies $\overline{\sigma}_i$ is constant-play, as described above.*

**Remark 4.** ***On the simplicity of constant-play strategies.***

*Constant-play strategies are attractive from perspectives of bounded-rationality. Notice that even if the extensive version of the game is highly complex, playing a constant-play strategy is relatively simple. For example, communicating the instructions for playing such a strategy is as simple as the definition above. In*

---

[12] Recall, as argued by Kreps and Wilson (1982), that one can move all the random choices in a game tree to an initial nature node; the relevant parts of the outcomes realized at this initial node will only be partially and differentially revealed at their corresponding place in the tree.

[13] The sum of the probabilities of all the out-going arcs labeled with the action $a_i$ equals $\sigma_i(a_i \mid t_i)$.

*playing it, only minimal knowledge of the extensive version being played is needed. A player only needs to keep track of whether it is his initial move or not, and when it is not, he must only remember what his previous choice was.*

4.3. **EXTENSIVELY ROBUST EQUILIBRIA.** A Nash equilibrium $\sigma$ in a Bayesian game $G$ is *extensively robust* if in every extensive version of $G$, every profile of its constant-play versions $\overline{\sigma} = (\overline{\sigma}_1, ..., \overline{\sigma}_n)$ is a Nash equilibrium. In other words, at such an equilibrium, when player $i$ plays a constant-play version of $\sigma_i$, he is assured to be best-responding to his opponents without even having to know what extensive version is being played. This greatly simplifies the computation he needs to perform.

But we need to define a notion of being *approximately robust*, so that in every $\overline{G}$, every $\overline{\sigma}$ is required to be only an approximate Nash equilibrium and only with high probability. This notion assures that the incentives of any player to deviate unilaterally at any positive-probability information set are insignificant.

For a given extensive version $\overline{G}$ and a vector of behavioral strategies $\overline{\eta} = (\overline{\eta}_1, ..., \overline{\eta}_n)$, we use the natural probability distribution induced by $\overline{\eta}$ over the outcomes of the game: the complete play paths. The payoff of player $i$ is defined to be the usual expected value, $E_{\overline{\eta}}(u_i)$.

Given an information set $A$ of player $i$, a *modification of player i's strategy at A* is any strategy $\overline{\eta}'_i$ with the following property: at every information set $B$ of player $i$ which is not a follower of the information set $A$,[14] $\overline{\eta}'_i$ coincides with $\overline{\eta}_i$. Player $i$ can *unilaterally improve his payoff by more than* $\varepsilon$ *at the information set* $A$ if $E_{\overline{\eta}'}(u_i|A) - E_{\overline{\eta}}(u_i|A) > \varepsilon$, where $\overline{\eta}' = (\overline{\eta}_1, \overline{\eta}_2, ..., \overline{\eta}'_i, ..., \overline{\eta}_n)$ for some $\overline{\eta}'_i$ that is a modification of $\overline{\eta}_i$ at $A$.

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

Note also that not having an $\varepsilon$ improvement, as defined above, is restrictive in two important ways. First, it rules out significant improvements, even if attained through coordinated changes in the player's later information sets. Moreover, the improvements must be small, even when viewed conditionally on being at the information set (not just from the ex-ante perspective).

**Definition 7.** ***Approximate Nash equilibrium:*** *A strategy profile* $\overline{\eta}$ *of* $\overline{G}$ *is an* $(\varepsilon, \rho)$ *Nash equilibrium, if the probability that some player has a better than* $\varepsilon$ *improvement at some information set is not greater than* $\rho$.

**Definition 8.** ***Approximate robustness:*** *An equilibrium of* $G$, $\sigma$, *is* $(\varepsilon, \rho)$ extensively robust, *if in every extensive version* $\overline{G}$ *every profile of constant-play versions of* $\sigma$, $\overline{\sigma}$, *is an* $(\varepsilon, \rho)$ *Nash equilibrium.*

*An equilibrium* $\sigma$ *is* $(\varepsilon, \rho)$ ex-post Nash *if its constant-play version is an* $(\varepsilon, \rho)$ *Nash equilibrium in the extensive version with one round of revisions (defined earlier).*

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

Clearly, being $(\varepsilon, \rho)$ ex-post Nash is a consequence of being $(\varepsilon, \rho)$ extensively robust. The converse is not true, as can be seen by the following example.

**Example 7.** ***Buying Insurance:***

*This two-person game is between Fate and a risk-taker, who has to choose between buying or not buying insurance for his car. Fate has flat preferences of zero, while the payoffs of the risk-taker are given by the following table:*

| | *Fate* | |
|---|---|---|
| *Risk Taker* | *accident* | *no accident* |
| *insurance* | *0.99* | *0.99* |
| *no insurance* | *0* | *1.00* |

*Consider an equilibrium with Fate choosing an accident with probability* .001*, no accident with probability* .999*, and with the risk-taker buying no insurance. This equilibrium is* $(0, .001)$ *ex-post Nash, since the probability of the risk-taker being able to gain any positive amount by revising, after seeing Fate's choice, is not greater than* .001*.*

*To check the level of extensive robustness, consider the extensive version in Figure 5, which may be described verbally as follows. In addition to the accident/no-accident choice, Fate also decides on an intermediary event,* the driver and the car are late*. The decision maker observes the late/ no-late choice, before he decides whether or not to buy insurance. Assume that Fate chooses* $\Pr(late|accident) = 1$ *and* $\Pr(late|\text{ no accident}) = 0.05$*. (Notice that the Pr(late and no acc.)+Pr(not late and no acc.)=* $.04995 + .94905 = .95$*, as required by the definition of an extensive version).*

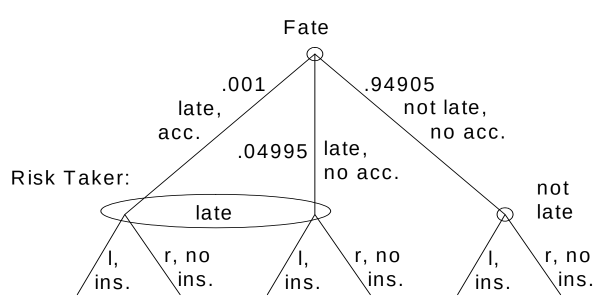


FIGURE 5

*Doing the Baysian computations reveals that* $\Pr(late) = .051$*,* $E(payoff \mid insurance\ and\ late) = .99$ *and* $E(payoff \mid no\ insurance\ and\ late) = .98$*. So at the information set* late *the risk taker can gain at least* .01 *by switching from* no insurance *to* insurance*, and the probability of visiting this information set is* .051*. In other words, the equilibrium, which is* $(0, .001)$ *ex-post Nash, can be no better than* $(.01, .05)$ *extensively robust.*

**Remark 5.** *One can check that in a complete-information normal-form game, every pure-strategy Nash equilibrium is ex-post Nash and even extensively robust, regardless of the number of players. This is no longer the case for incomplete-information games, as was illustrated by the example in the introduction. In the* *two-player game of choosing computers, the pure strategy of choosing-one's-favorite-computer is clearly not robust. And this is so despite the fact that it is even a strict Nash equilibrium.*

**Theorem 2.** ***Extensive robustness****: Consider a family of games $\Gamma(\mathcal{T},\mathcal{A})$ with continuous and anonymous payoff functions as above, and a positive number $\varepsilon$. There are positive constants $\alpha=\alpha(\Gamma,\varepsilon)$ and $\beta=\beta(\Gamma,\varepsilon)$, $\beta<1$, such that all the equilibria of games in $\Gamma$ with $m$ or more players are $(\varepsilon,\alpha\beta^m)$ extensively robust.*

**Remark 6.** ***Failure of converse and focul point of simplicity****. As illustrated at the concluding section of the paper, a converse theorem fails: there are extensive versions of large games with equilibria that are not constant-play versions of the simultaneous-move equilibria. Still, the constant-play versions of the simultaneous-move equilibria stand out in a couple of ways.*

*First, as already noted, they consist of simple strategies and thus may have a strong focul-point effect on a group of players having to choose an equilibrium. Players may be drawn to reason "lets not bother with the complex rules of the extensive version (play the game as if it was one simultaneous move)," in light of the fact that it is a best response to do so. Second, the fact that the simultaneous move game is an extensive versions of itself implies that the constant-play versions of the simultaneous-move equilibria are the only elements in the intersection of the equilibria of all the extensive versions.*

## 5. PROOFS AND FURTHER RESULTS

While one may view the ex-post Nash theorem as an immediate corollary of extensive robustness, the method of proof we use first establishes the ex-post Nash property. The fact that the ex-post Nash property is obtained at an exponential rate is then shown to imply the full extensive robustness property.

We also show that a weaker property than continuity, namely, low local strategic interdependence, is the essence behind ex-post Nash. We first introduce this concept and use it to develop bounds on the level of the ex-post Nash property obtained at an arbitrary given equilibrium. These bounds are then used to prove the asymptotic ex-post Nash result.

### 5.1. LOW STRATEGIC INTERDEPENDENCE IMPLIES EX-POST NASH.

Every equilibrium is $(\varepsilon,\rho)$ ex-post Nash for sufficiently large $\varepsilon$ or $\rho$. This section concentrates on a fixed Nash equilibrium of a fixed Bayesian game and develops bounds on the levels of the ex-post Nash property it must have. These bounds are used later to prove the main result. As it turns out, for an equilibrium to be highly ex-post Nash, we do not need uniform continuity of the players' payoff functions; rather, continuity in a region near the expected play of the equilibrium is sufficient. And in fact, the only property that is really needed is low strategic interdependence in such a restricted region.

**Definition 9.** ***Strategic interdependence:*** *The strategic dependence of player $i$ in a set of type-action character profiles $M$, $sd_i(M)$, is defined to be*

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

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

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

This means that if $sd_i(M)$ is small and the type-action character profiles are likely to be in M, then the gain to player $i$ in a switch from action $a_i'$ to $a_i''$ is almost independent of his type and of the type-action characters of the opponents. If this is the case, any uncertainty about opponents' types and opponents' selected actions may make only a minor difference in his decision about what action to choose.

**Lemma 1.** *A Bayesian equilibrium is $(\varepsilon, \rho)$ ex-post Nash if for some set of type-action character profiles $M$,*

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

*where the maximum is taken over all players $i$ and all $c_i$'s that are part of a type-action character profile $c\epsilon M$.*

*Proof.* It suffices to show that at any $c\epsilon M$, no player can improve his payoff by more than $\varepsilon$ by switching from his $a_i$ to another action. From the definition of strategic interdependence, if switching from $a_i$ to $a_i'$ at $c$ improves player $i$'s payoff by $r$, then the same switch must improve his payoff by at least $r - si(M)$ at any other $\bar{c}\epsilon M$ with $\bar{t}_i = t_i$. (The improvement referred to is the following: fix the opponents' type-action characters 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 $a_i$, relying on the 0-1 normalization of his utility function we see that 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 by $i$ to be an optimal response, the last expression must be non-positive, which yields the desired bound. □

The above lemma illustrates that if $\sigma$ generates a low strategic-interdependence set $M$, which has high-probability in the conditional sense just described, then $\sigma$ is highly ex-post Nash. The following discussion illustrates that under natural restrictions on the game, and assuming a large number of players, such sets $M$ are natural.

Recall first that for every vector of type-action characters $c$, as a function of all possible individual type-action characters $\kappa$ in the universal set of type-action characters $\mathcal{K}$, the empirical distribution was defined to be $emp_c(\kappa) =$

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

Since the profile of type-action characters is randomly generated at an equilibrium, the empirical distribution is a $|\mathcal{K}|$-dimensional random variable. We next identify a $|\mathcal{K}|$-dimensional vector of numbers, called the expected distribution, that constitute, coordinate by coordinate, the expected values of the empirical distribution.

Starting with a strategy profile $\sigma$, the induced vector of measures $\gamma = (\gamma_1, \gamma_2, ..., \gamma_n)$ may be viewed as a vector of extended distributions, each being defined over the universal set of all possible type-action characters $\mathcal{K}$ (as opposed to each $\gamma_i$ being defined on $C_i \subset \mathcal{K}$). Specifically, for any possible type-action character $\kappa \in \mathcal{K}$, $\gamma_i(\kappa) = \gamma_i(c_i)$ if $\kappa$ equals some $c_i \in C_i$, and $\gamma_i(\kappa) = 0$ otherwise.

**Definition 10.** ***Expected distribution**: For a strategy profile $\sigma$ and the induced distribution $\gamma = (\gamma_1, \gamma_2, ..., \gamma_n)$, define the* expected distribution *on the universal set of type-action characters $\mathcal{K}$ by $\exp_\sigma(\kappa) = \sum_i \gamma_i(\kappa) \ / \ n$.*

We now proceed to argue that when n is large (relative to $|\mathcal{K}|$), for any profile of strategies $\sigma$ there is a high probability of realizing a profile of type-action characters $c$ with $emp_c(\kappa)$ being close to $\exp_\sigma(\kappa)$ for every $\kappa \in \mathcal{K}$.

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

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

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

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

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

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

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

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

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

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

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

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

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

**Theorem 3.** ***Bounds on the level of ex-post Nash****: For any $\delta > 0$, a Bayesian equilibrium $\sigma$ is $(\varepsilon, \rho)$ ex-post Nash if*

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

5.2. **PROOF OF THE EX-POST NASH PROPERTY.** We fix the family of semi-anonymous games $\Gamma$, with the collection of uniformly equicontinuous payoff functions as in the statement of the Ex-Post Nash Theorem. It is sufficient to prove the theorem only for values of $n$ from some $m$ and above. That is, given the family of games and a positive $\varepsilon$, there is an $m$ and constants $\alpha$ and $\beta$, $\beta < 1$, such that for all $n \geq m$, all the equilibria of $n$-player games in the family are $(\varepsilon, \alpha\beta^n)$ ex-post Nash. (Once one proves it for such an $m$, one can simply increase $\alpha$ to an $\alpha'$, so that the conclusion, with $\alpha'$ and $\beta$, is trivially satisfied for all values of $n < m$.)

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

*Proof.* Recall that $si[nbd(\exp_\sigma, \delta)]$ is defined to be the maximum of the expression below, when one considers all players $i$, all pairs of actions $a_i'$ and $a_i''$, all types $t_i$, and all pairs of type-action character profiles $c^1, c^2 \in nbd(\exp_\sigma, \delta)$ having $t^1 = t^2 = t_i$:

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

But by rearranging terms, this expression becomes:

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

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

To return to the proof of the theorem, recall from the previous section that an equilibrium $\sigma$ is $(\varepsilon, \rho)$ ex-post Nash if for some $\delta > 0$, the following two inequalities are satisfied:

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

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

The proof is now completed by setting $\alpha = 2\,|\mathcal{K}|$ and $\beta = e^{-2\delta^2}$.

5.3. **PROOF OF EXTENSIVE ROBUSTNESS.** The proof of extensive robustness follows two steps: (1) to show that all the equilibria in the family become ex-post Nash at an exponential rate as the number of players increases, and then (2) to show that this implies that they become extensively robust at an exponential rate.

The first step is the Ex-Post Nash theorem just proven, which we restate in a slightly more convenient version as the following lemma.

**Lemma 7.** *For every positive $\varepsilon$ there are positive constants $\alpha$ and $\beta$, $\beta < 1$, such that all the equilibria of games in $\Gamma(\mathcal{T}, \mathcal{A})$ with $m$ or more players are $(\varepsilon, \rho_m)$ ex-post Nash with $\rho_m \leq \alpha\beta^m$.*

The above, together with the next proposition, directly yields the proof of the theorem but with $\rho_m \leq m\alpha\beta^m$. This, however, is sufficient for completing the proof, since we can replace $\alpha$ and $\beta$ by bigger positive $\alpha'$ and $\beta', \beta' < 1$, for which $m\alpha\beta^m < \alpha'\beta'^m$ for $m = 1, 2, \ldots$ .

**Proposition 1.** *If $\sigma$ is an $(\varepsilon, \rho)$ ex-post Nash equilibrium of an $n$-player game $G$, then for any $\zeta > 0$, $\sigma$ is $(\varepsilon + \zeta, n\rho/\zeta)$ extensively robust.*

*Proof.* It 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: player $i$ has a better than $\varepsilon + \zeta$ improvement at some information set visited by a play path in $V$. Let $W$ be the set of type-action character profiles $c$ with the property that player $i$ has a better than $\varepsilon$ improvement at $c$, i.e., by a unilateral change of his action at $c$, he can improve his payoff by more than $\varepsilon$.

$\rho \geq \Pr(W)$ by assumption and the probability-coincidence lemma below,

$\Pr(W) \geq \sum_{A \subset V} \Pr(W|A) \Pr(A)$ by the decomposition lemma below, and

$\sum_{A \subset V} \Pr(W|A) \Pr(A) \geq \zeta \Pr(V)$ by the bounds lemma below. The weak inequality is needed to accommodate the case in which $\Pr(V)$ and $\rho$ are both zero.

Combining the above inequalities completes the proof of the proposition. □

**Lemma 8.** ***Probability coincidence****: The probabilities of any type-action character profile c computed either (1) directly from a strategy profile $\sigma$ in G, or (2) by any constant-play profile of strategies $\overline{\sigma}$ in $\overline{G}$, are the same.*

*Proof.* Notice that we can simplify $\overline{G}$ by removing every non-initial information set of every player from the tree, without affecting the distribution over the resulting type-action character profiles. This is due to the constant-play (no revision) property of $\overline{\sigma}$. So we assume without loss of generality that $\overline{G}$ has every player moving only once, and that he randomizes with the same probabilities as in $\sigma$ when it is his turn to play.

Conditional on every realized profile of types and order of players' moves, the probability of c computed from the tree coincides with the probability computed directly from $G$. □

**Lemma 9.** ***Decomposition****: $V$ can be represented as a disjoint union of information sets $A$, when we view each information set $A$ as the event containing the play paths that visit the information set $A$.*

*Proof.* From the definition of $V$, every information set of player $i$ must be either fully included in $V$ or fully included in the complement of $V$. Moreover, due to the perfect recall assumption, the information sets are well ordered in the tree, as a partial order in the timing of play of the tree. This means that they are well ordered by containment as events, with the events corresponding to later information sets being subsets of the events corresponding to earlier ones. Thus, if for every play path we take the first information set that exhibits the violation of V, we have a collection of disjoint events whose union equals $V$. □

**Lemma 10.** ***Bounds****: For any positive-probability information set $A$ of player $i$, if player $i$ has better than $\varepsilon + \zeta$ improvement at $A$, then $\Pr(W|A) > \zeta$.*

*Proof.* We first assert that at any positive-probability information set $A$, any modification of player $i$'s strategy at $A$ does not affect the probability distribution over the profile of opponents' type-action characters. This assertion can be checked node by node. At every node, the opponents who played earlier in the game tree will not revise (by the definition of $\overline{\sigma}$) and thus, their type-action characters stay fixed regardless of the modification. The opponents who did not play prior to reaching the node will randomize according to $\sigma$ 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 $A$ (that he never revises later on).

Now we can put a bound on the possible levels of such improvement as follows. For type-action character 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 type-action character profiles in $W^c$, the largest possible improvement is $\varepsilon$. 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 $\Pr(W|A) + \varepsilon \geq 1\Pr(W|A) + \varepsilon\Pr(W^c|A) > \varepsilon + \zeta$, which validates the claim made above. □

## 6. ADDITIONAL COMMENTS

6.1. **IMPLICATIONS FOR MECHANISM DESIGN.** The ex-post Nash property overcomes certain modeling difficulties in the implementation literature; see, for example, Cremer and McLean (1985), Green and Laffont (1987), and a significant follow-up literature.[15] Also, as argued by Milgrom and Wilson,[16] in real-life design of auction and resale markets, one is concerned that the process stops. Otherwise ex-post interaction may upset the equilibrium computed by the designer for the one-shot game.

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 as concerned that the players play an extensive version of his game. Even if the players can engage in cheap talk, act sequentially, share information, revise choices after the implementation game is complete, and more, the equilibrium constructed for the one-shot simultaneous-move game remains viable.[17] In various social aggregation methods, extensive robustness means that the outcome of a vote is immune to institutional changes, and public polls should not alter the outcome of the equilibrium.

6.2. **EXTENSIVE RATIONAL EXPECTATIONS UNDER INFORMATION INDEPENDENCE.** At a rational-expectations equilibrium, agents base their trade choices on their own preferences and information, and on the observed market prices. The system is at equilibrium because no inference from the prices gives any agent an incentive to alter his trade choices.

From a game-theoretic perspective, such equilibrium seems to "mix together" ex-ante information with ex-post results. For example, in a well-defined Bayesian version of a market game of the Shapley and Shubik (1977) variety,[18] agents' endowments, information, and preferences are ex-ante input, and prices are resulting ex-post output that is determined by the trade choices of the agents.

---

[15]See Chung and Ely (2000) for additional discussion and more recent references.

[16]Private communication, May 2003.

[17]In social choice terminology, the implementation in an extensive version is not "exact," since there may be other equilibria of the extensive version that do not meet the goals of the implementor. But as discussed earlier, the constant-play equilibria do have a focal point of simplicity.

[18]For an example of such a model and additional references, see Peck (2003).

One can see modeling advantages to these two contradictory approaches. A game theorist might argue that in a well-defined game you cannot have it both ways; if prices are consequences of trading decisions, you cannot have prices as input at the time players contemplate their trade choices. An economist, on the other hand, might argue that players do know the prices when they make trade choices, so not including prices in the domain of the decision rules is improper. To model this, economic theorists view rational-expectations equilibrium as a fixed point of a bigger system, one that deals simultaneously with the ex-ante and ex-post information (see for example Jordan and Radner [1982]). This, however, often leads to nonexistence (see Grossman and Stiglitz [1980]).

As it turns out, extensive robust equilibria largely resolve the above contradictions. Moreover, they offer a stronger version, actually an extensive one, of rational-expectations equilibrium[19].

To illustrate this, consider an equilibrium of a Bayesian Shapley-Shubik game of trading computers and related products, with continuous payoff functions and many players with independent types. Assume, for simplicity of our discussion here, that the real process of trade being modeled, which may not be played simultaneously, takes place in a relatively short period of time, so that no time discounting of future expenses and payoffs is necessary.

The extensive robustness property implies that partial ex-post information, including prices, gives no player an incentive to change his trade choices. Thus, the equilibrium possesses the rational-expectations property. But this is even true for partial intermediary information, such as partially formed prices, revealed at any stage of the game (if not played simultaneously), and partially observed choices and behavior of opponents. For example, some players may observe current prices of related software, before deciding on what computers to buy; others may have statistics on what was bought before them; and some may buy with no information, or just information on what some of their neighbors buy. Still, the equilibria of the simultaneous-move game is sustained, and all such price and behavior information will not affect their stability. Thus, such an equilibrium is not just a fixed point of private information and prices, but remains a fixed point of much richer extensive systems.

The above properties may be viewed as a non-cooperative strategic/informational foundation for extensive rational-expectations equilibria, in a way that is parallel to cooperative equivalence theorems, where the core and Shapley value are used as cooperative game-theoretic foundations for competitive equilibrium.[20] The limitation to independent types, a severe restriction in the case of rational expectations equilibrium, offers a strong motivation to study extensive robustness with correlated types.

6.3. **LEARNING.** There is an interesting connection between the ex-post Nash property and learnability in repeated games in the sense of Kalai and Lehrer (1993). Learning within an equilibrium of a Bayesian repeated game, in the Kalai-Lehrer sense, means that from some time on, the strategies used by the players are a Nash

[19]Existence is not an issue for us, since we deal with large but finite games.

[20]The author thanks Avinash Dixit for making this observation. For earlier studies of game theory and rational expectations, see Forges and Manelli (1997, 1998). Additional references and newer results (obtained after publication of the first draft of the current paper) may be found in Perry and Reny (2003).

equilibrium of the repeated game with the realized types being commonly known. In other words, even if the players initially play the repeated game without knowing the realized types of their opponents, with time they converge to play as if they do.

The ex-post Nash property discussed in this paper may be thought of as the players learning to play, in the Kalai-Lehrer sense above, already at time zero. Even in the first, and the only, period of play, the players play strategies that are optimal, as if they knew the realized types of their opponents. So in a large semi-anonymous game, learning, or more precisely playing as-if-you-know, is immediate.

The above observation suggests a promising line of investigation: trying to connect the speed of learning to the number of players in Bayesian repeated games with semi-anonymous players. It seems reasonable that if we require fast learning rather than immediate knowledge, as discussed above, we may be able to obtain results about convergence to extensive robustness, and to rational expectations equilibria in the case of market games.

6.4. **OPEN-LOOP VERSUS CLOSED-LOOP EQUILIBRIA.** Consider a normal-form game repeated twice, with the following two possible information structures: (1) no information (the second round is played with no information about opponents' first round choices), and (2) perfect monitoring (the second round is played with perfect information about opponents' first round choices). Referring to an equilibrium of the no-information game as open-loop and to an equilibrium of the perfect-monitoring game as closed-loop, Fudenberg and Levine (1988) study the relationships between the two types of equilibria. A main conclusion is that if the number of players is large, so that the influence of each player on his opponents is negligible, then every open-loop equilibrium is also closed-loop.

In the terminology of the current paper, the twice-repeated game above may be viewed as a game with one round of revisions.[21] Constant-play strategies in this game can be used in both the no-information version and the perfect-monitoring version. An equilibrium being ex-post Nash (in the current terminology) implies that it is a closed-loop equilibrium (in Fudenberg-Levine terminology). Thus, our results, and in particular the finding that every equilibrium is ex-post Nash, are consistent with the Fudenberg-Levine finding that open-loop equilibria must be closed-loop in games with many players.

However, since the full extensive robustness applies to all extensive versions and information structures, not just to two rounds of repetitions, it should be clear to the reader that extensive robustness is substantially stronger than "open-loop implies closed-loop."

Also, unlike the Fudenberg-Levine paper, the current paper deals with Bayesian and not just normal-form games. Indeed, as Fudenberg and Levine argue, their result may serve as a foundation to competitive equilibria, while what we have in this paper may serve as a foundation to rational-expectations equilibria.

Two additional issues, raised below, are related to issues discussed in Fudenberg and Levine (1998). One is the converse to our result, which mirrors the question of when a closed-loop equilibrium is open-loop. The other is the issue of subgame-perfection.

[21]One difference is that in the Fudenberg-Leving paper the players are paid twice, once in each round, while in the game with one round of revisions of the current paper, there is only one-time payoff at the end of the second round. But this difference does not affect any of the points we make below.

6.5. **SUBGAME PERFECTION WITH MANY PLAYERS.** As mentioned in the introduction, an extensively robust equilibrium is required to remain a Nash equilibrium, without being subgame-perfect, in every extensive version of the game. It turns out to be impossible to have extensive robustness together with subgame perfection.

However, it seems that when the number of players is large, a lack of subgame perfection may become a less severe deficiency. The following example illustrates these points.

**Example 8.** ***Big Battle of the Sexes:***
*Simultaneously, each of $n$ male players and $n$ female players 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 only* $0.9$ *of the proportion she matches when she chooses I.*

Consider the above game played sequentially in a predetermined commonly-known order in which the females choose first and the males follow, with every player being informed of the choices made by all earlier players. The reader may verify that the only subgame-perfect equilibrium has all players choose M.

Since the symmetrically opposite result must hold for the case of all males moving first, i.e., the only outcome is everybody choosing I, it is clear that we cannot have a subgame-perfect equilibrium that is immune to varying the order of play. In other words, we cannot have extensive robustness together with subgame perfection.

It is also clear that the extensively robust equilibrium all-choose-I is not subgame-perfect in the case in which the females move first. How severe is this lack of subgame perfection?

When we have a small number of players - for example two - all-choose-I is highly incredible due to the following standard argument: If the female, who moves first, deviates to M, the male must choose between obeying the equilibrium strategy and being paid .5 for sure, or deviating himself to M and being paid .9 for sure. Being rational, the male would have to choose M. Knowing this, the female deduces that if she sticks to the equilibrium strategy, she will be paid .9, and if she defects to M, she will be paid 1. Thus her conclusion is clear and the all-I equilibrium fails.

Now, let us imagine the same scenario of an equilibrium with all players choosing I, but with one million females moving first and one million males following, and let us view the situation once again from the point of view of the first female.

In order to make her deviation from the all-I equilibrium worthwhile, she must believe that substantially more than one million followers will deviate too. Otherwise, deviating on her part may be quite costly. Moreover, her immediate follower has similar concerns that may prevent her from deviating to M. Players can no longer rely on direct incentives of their immediate followers; instead, they must rely on a long chain of followers who rely on the rationality of their followers, which is also based on such long chains of cumulative reasoning. So, unlike the case of the two-player game above, where a deviation by the first deviator induces direct immediate incentives to deviate in 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 conviction by more than a million 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 one million-subgame perfect. In the example with two million players, at least one million deviations from the equilibrium path are required before we reach a node where a player's choice is clearly not optimal.

6.6. **A CONVERSE THEOREM FAILS.** Consider again the Village versus Beach example discussed earlier, but now with 1001 females and 1001 males. The payoff of a female equals the proportion of males she mismatches and the payoff of a male equals the proportion of females he matches. At an equilibrium of the simultaneous-move game, the (expected) payoff of every player, male or female, must be .5.

But consider the extensive version, with the game being played sequentially in a known order, with perfect monitoring, and all the females moving first. Consider an equilibrium with 500 females choosing the village, 500 choosing the beach, and the last one randomizing with equal probabilities. The males all choose the location realized by the randomizing female. While the payoffs of the males are .501 each, 501 females receive payoff 0, and 500 females receive payoff 1.

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

| | .40 | .60 | 0 | 0 |
|---|---|---|---|---|
| .20 | $\mathbf{6,7}$ | $\mathbf{5,7}$ | $1,2$ | $7,3$ |
| .30 | $\mathbf{6,5}$ | $\mathbf{5,5}$ | $9,2$ | $5,1$ |
| .50 | $\mathbf{6,8}$ | $\mathbf{5,8}$ | $5,4$ | $4,3$ |
| 0 | $2,2$ | $3,4$ | $9,8$ | $7,7$ |

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

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

**Proposition 2.** ***Nash Interchangeability of Outcomes***. *A strategy profile of a Bayesian game with independent types is ex-post Nash if and only if its support consists of Nash type-action character profiles that are interchangeable.*

The reader can construct the direct couterparts to the above statements for the case of approximate ex-post Nash equilibria.

## 7. REFERENCES

Ali Khan, M., and Y. Sun (2002): "Non-Cooperative Games with Many Players," Chapter 46 in the *Handbook of Game Theory with Economic Applications*, Vol 3, pp. 1761-1808, R.J. Aumann and S. Hart editors, North Holland, Amsterdam.

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 and Sons, New York.

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

Cartwright, E. and M. Wooders (2003): "On Equilibrium in Pure Strategies in Games with Many Players," University of Warwick.

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

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

Ellison, G. and D. Fudenberg (2002): "Knife-Edge or Plateau: When Do Market Models Tip?" MIT/Harvard discussion paper.

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, pp. 388-406.

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

Fudenberg, D. and D. Levine (1988): "Open-Loop and Closed-Loop Equilibria in Dynamic Games with Many Players," *Journal of Economic Theory* 44, pp. 1-18.

Green, E., (1980): "Noncooperative Price Taking in Large Dynamic Markets," *Journal of Economic Theory*, 22, pp. 155-82.

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.

Grossman, S.J. and J.E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70(3), pp. 393-407.

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.

Kreps, D., and R. Wilson (1982): "Sequential Equilibria," *Econometrica* 50, pp. 863-894.

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.

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.

Peck, J., (2003):"Large Market Games with Demand Uncertainty," *Journal of Economic Theory* 109, pp. 283-299.

Perry, M. and P. Reny (2003): "Toward a Strategic Foundation for Rational Expectations Equilibrium," Hebrew University.

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.

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

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

Shapley, L.S. and M. Shubik (1977): "Trade Using One Commodity as a Means of Payment," *Journal of Political Economy*, 85, pp. 937-68.

Wilson, R. (1987): "Game-Theoretic Analyses of Trading Processes," *Advances in Economic Theory: Fifth World Congress*, Truman Bewley (ed.); Chapter 2, pp. 33-70. Cambridge UK: Cambridge University Press.

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