LARGE ROBUST GAMES

EHUD KALAI

Abstract. The equilibria of certain simultaneous-move normal form and Bayesian games are extensively robust when the number of players is large. Even if played sequentially, with information partially and differentialy revealed, with revision and commitment possibilities, with cheap talk announcements and more, the equilibria of the simultaneous-move one-shot game survives.

This robustness property is important for certain applications of game theory, and may also serve as strategic-informational foundation for rational expectations equilibrium under conditions of informational independence.

1. Introduction

Games with many players is a old topic in economics and in cooperative game theory. Prominent theorists, including Aumann, Debreu, Hildenbrandt, Scarf, Shapley and Shubik, have shown that in large cooperative games major modeling difficulties disappear. Specifically, solution concepts such as the core, competitive equilibrium and the Shapley value, that in general predict different outcomes for the same game, coincide to give the same prediction when the number of players is large. As a special case, this coincidence offers cooperative-coalitional foundation for competitive equilibrium. For a general survey see the book of Aumann and Shapley (1974).

Less is known for general non-cooperative strategic games with many players. The first study of this subject is 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-Postlewaite (1990) on bargaining, Rustichini-Satterthwaite-Williams (1994) and Pesendorfer-Swinkels (1997) on auctions, and Feddersen-Pesendorfer (1997) on voting. For the most part, these papers concentrate on issues of economic efficiency. Another general direction, is the literature on repeated large games, see for example Green (1980) and Sabourian (1990), where the focus is on short term behavior exhibited by patient players.

The objective of the current paper is to uncover properties of general strategic games with many players, beyond the one exhibited 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.

Date: March 12, 2002,
this version May 2003.

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

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

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 the players of the game, the prediction of the equilibrium may be unreliable. This paper illustrates that this difficulty is less severe in general classes of games that involve many semi-anonymous players. In normal form games and in 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.

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, 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 with the properties mentioned above.

Extensive robustness has direct positive implications to areas where game theory is applied. A mechanism designer, who succeeds in implementing a socially efficient outcome through a Nash equilibrium of a one-shot simultaneous-move game, does not have to be concerned that the players play a different extensive version of his game, see for example Green and Laffont (1987). 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. 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. And below, we discuss examples that show how extensively robust equilibrium may be a useful concept, exhibiting strong (extensive) rational expectations properties, for more robust modelling of behavior in market games.

When addressing robustness issues, previous studies imposed a weaker condition, known as ex-post Nash[1]. Applications of this idea to specific economic problems include Cremer-McLean (1985), Green-Laffont (1987) and Minehart-Scotchmer (1999). A general result illustrating that the ex-post Nash property is obtained for large Bayesian games is described in Kalai (2000,2003). Indeed, the proof of the main result of the current paper makes use of this weaker result and, very importantly, of the fact that it is obtained at an exponential rate as the number of players increases.

For a survey of rational expectations equilibrium we refer the reader to Jordan-Radner (1982). Forge-Minelli (1997,1998) and Minelli-Polemarchakis (1997) are earlier studies relating rational expectations equilibrium to game theory.

1.1. **Example: a Game of Evolving Standards.** Simultaneously, each of n players has to choose computer I or computer M, and, independently of the opponents, each is equally likely be a type who likes I or a type who likes M. Most of a player's payoff comes from matching the choice of the opponents but there is also a small payoff in choosing the computer he likes. Specifically, each player's payoff function is 0.1 if he chooses his favorite computer, zero otherwise, plus 0.9 times

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

the proportion of opponents his choice matches. Assuming that each player knows only his own realized type before making the choice, the following three strategy profiles are Nash equilibria of the simultaneous move game: the constant strategies, with all the players choosing I or with all the players choosing M, and the one where every player chooses his 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 his predecessors, then everybody choosing I regardless of the observed history is a Nash (not subgame perfect) equilibrium of the extensive game. And they are robust to other modifications. For example, if a round of revision was allowed, after the players observe the opponents' first round choices, everybody choosing I with no revision remains an equilibrium of the two round game.

This is not the case for 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 his predecessor rather than choosing his 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 modelling difficulties, since we do not know in what order players contemplate the choice of, 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 change the equilibria of the game.

But the modeling difficulties become less severe if the number of players is large. Now, even choosing-their-favorite-computers is a highly robust equilibrium, it remains an approximate Nash equilibrium in the sequential 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-their-favorite-computers. Every Nash equilibrium of the one shot game is extensively robust, and this is true even if the original computer selection game that we start with is more complex and highly non symmetric. There may be any finite number of computer choices and player types, different players may have different (arbitrary) payoff functions and different prior probability distributions by which their types are drawn, and players' payoffs may take into consideration opponent types, and not just opponent actions (e.g., a player may have a positive payoff when others envy his choice, i.e., they like the computer he chose even though they themselves did not choose it). Regardless of such specifications, all the equilibria of the one shot game are highly extensively-robust when n is large.

The robustness of all one shot equilibria described above is not restricted to computer choice or other market games. Any game with payoff functions that depend on aggregate data of the opponents in a continuous fashion has this property. So various social aggregation games, games of joining clubs - political parties or other groups, location games, transportation and congestion games, may all fit the description.

The model described in the paper deals with a family of games that contains, for each n=1,2,3,... , many different n-person games. Under assumptions of continuity, anonymity and independence, uniformly at an exponential rate all the equilibria of all the games in the family become extensively robust as the number of players increases.

We should emphasize, however, that the equilibria obtained in the extensive versions above are approximately Nash, and not approximately subgame perfect Nash. This is unavoidable. However, as we note by an example at the end of the paper, this may be less of an issue as the number of players becomes large.

We should also note that the equilibrium in the extensive versions is approximately Nash in a strong sense. With almost certainty, along the realized play path, *none of the many players*, no matter how long the play path is, ever has a significant incentive to deviate from the equilibrium strategy.

1.2. **Strategic Foundations for Rational Expectations Equilibrium 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 any inference from the prices gives no 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, at a well-defined Bayesian version of a market game of the Shapley and Shubik (1977) variety, agents endowments, information and preferences would be ex-ante input, and prices are resulting ex-post output, determined by the trade choices of the agents.

One can see modeling appeals to these two contradictory approaches. A game theorist may 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, may 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. (The conceptual difficulty is overcome by viewing rational expectations equilibrium as a fixed point of a bigger system, one that deals simultaneously with the ex-ante and ex-post information. This, however, often leads to non-existence).

As it turns out, extensive robust equilibria largely resolves the above issues, and moreover, offer a stronger version, an extensive one, of rational expectations equilibrium.

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 we model, which may not be played simultaneously, takes place in a relatively short period of time, so that players do no time discounting for their forthcoming expenses and payoffs.

The extensive robustness property implies that ex-post partial information, such as prices, gives no player an incentive to change his trade choices. Thus, the equilibrium possesses the rational expectations property. But this is also 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 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 equilibrium of the simultaneous move game is sustained, and all such price and behavior information will not affect its stability. So the equilibrium is not just a fixed point of private information and prices, but it remains a fixed point of much richer extensive systems.

The above properties may be viewed as non-cooperative strategic-informational foundation for extensive rational-expectations equilibrium[2], in a parallel way to cooperative equivalence theorems, where the core and Shapley value are used as cooperative game-theoretic foundations for competitive equilibrium.

Looking more directly at the phenomenon, due to laws of large numbers, at an equilibrium prices are anticipated correctly by the players, who know the prior probability over types. And more explicitly than at rational expectations equilibria, in the Nash equilibria above players have no incentives to deviate from their strategy even during the formation of the equilibria, as computer and software purchases and other information is being observed.

As the last example shows, the assumption of known priors is strong, as it leads to correct anticipation of complete and partly formed prices in equilibria of games with many players. Interesting questions about learning in many-players games with unknown priors are left for future research. It is also important to remember that our model leaves out correlated information, a central topic to much of the rational-expectations literature. These issues are partially discussed later in the paper.

2. General Definitions and Notations

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

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

$T = \times_i T_i$ is the set of *type profiles* (or vectors), with each T_i being a finite set that describes the *feasible types of player* i .

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

$A = \times_i A_i$ is the set of *action profiles*, with each A_i being a finite set that describes the *feasible actions of player* i.

Let $C_i \equiv T_i \times A_i$ describe the *feasible characters*[3] *of player* i, and $C = \times_i C_i$ denote the set of *feasible character profiles.*

The *players' utility functions* described by the vector $u = (u_1, u_2, ..., u_n)$, assuming a suitable normalization, are of the form $u_i : C \to [0, 1]$.

In addition, standard game theoretic conventions will be used. For example, for a vector $x = (x_1, x_2, ..., x_n)$ and an element x_i', $x_{-i} = (x_1, ..., x_{i-1}, x_{i+1}, ..., x_n)$ and $(x_{i-1} : x_i') = (x_1, ..., x_{i-1}, x_i', x_{i+1}, ..., x_n)$.

---

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

[3] Clarifying the terminology, we use actions and types as is normally done in normal form and Baysian games. So actions may be (go to the) movie, ballet, eat quiche, drink beer, choose I, choose M, etc. Types may be strong, weak, likes I, likes M, has information about I, etc. Characters represent possible compositions of types with actions. Famous game theory characters are a weak type who eats quiche, a weak type who drinks beer etc. And in the current paper we have characters who like I and buy I, like I and buy M, are informed about I and buy M, etc.

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

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

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

The above definitions are standard, and, except for the assumption that player types are independent, the 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 introduce the notion of player's character, as above, so that payoff functions and probability distributions may be defined on a notationally simple space.

2.2. **Extensive Versions of the Game and Strategies.** In the sequel, we define an equilibrium of a Bayesian game G to be robust, if playing it, as constant-play strategies, is an equilibrium in all extensive versions of the game. A simple example of an extensive version of a game G is the following.

Definition 1. ***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, the information about the realized types and selected pure actions of all players becomes common knowledge, and then, simultaneously, every player has the one time opportunity to revise his first-round choice to any of his feasible G actions..*

In the above game, the constant-play strategies of the equilibrium have the players choose their G (possibly mixed) strategies in the first round, and stay with their realized choices as their choices for the second round, i.e., no revisions. This special case, where the equilibrium with no revisions remains an equilibrium in the two rounds game, is an alternative way to define the ex-post Nash condition, as discussed later.

In various applications with the ex-post Nash condition, for example Green-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.

Notice that in the one-revision-round game all constant-play strategies are possible, since in the first round a player is allowed to choose (deterministically or randomly) all of his G strategies, and in the second round he is allowed to keep his original choice with no revision.

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

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

To make the above more concrete, we offers some simple examples of extensive versions of a game with two players that have to choose one of the two computers. In all three extensive versions of Figure 1 nature moves first, chooses the (only) central arc with probability one, resulting in player 1 being a type who likes I (i) and player two being a type who likes M (m). The first extensive version 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 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), 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 that player 1 plays twice, the final outcome is determined by his last choice, since repeated choices represent revisions.

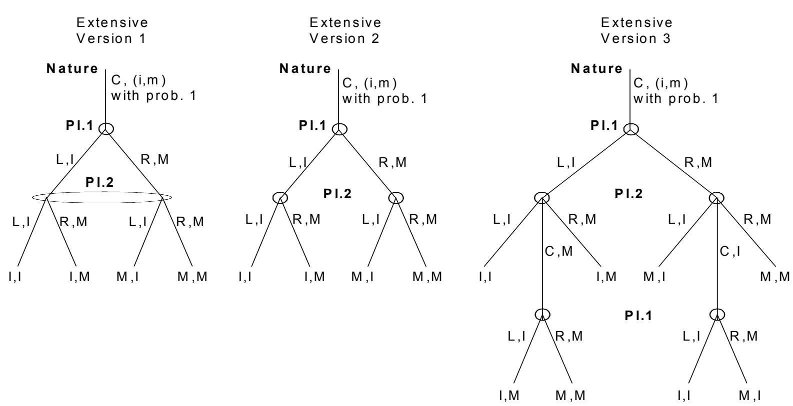

FIGURE 1

Extensive versions of a game allow for incomplete information, beyond the types allowed for in the simultaneous move Bayesian game. For example, when player 1 makes his initial selection he may not know which of the three games in Figure 1 he is playing. Such a game is described by the extensive version in Figure 2, where player 1 believes that he is most likely playing the simultaneous move game, but there is a small probability that player 2 may know his choice when she moves, and there is also a small probability that she may be able to offer him the possibility to revise.

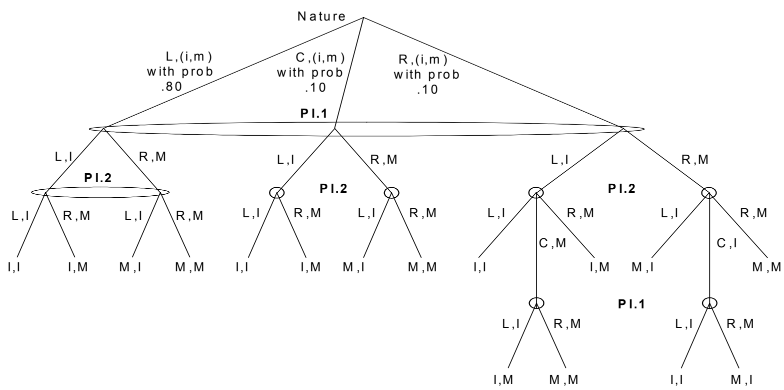

FIGURE 2

A reader familiar with the rich modelling 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 1. *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 suppose to answer all informational questions. For example, most likely player A does not know if 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 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 moderate number of players.

Because of this modeling difficultly, 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.

The following abstract definition of an extensive version of a game 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 includes words such as L,C, and R, as in the examples above, various annoucements, 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 type for each player. 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 τ 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 $M \times A_i$. The first coordinate represents the continuation game that player i chooses, the second coordinate represents the choice of an action from G.

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

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

The resulting 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 , at the resulting character profile of the path.

All constant-play versions of strategies of G can be played. More precisely:

1. for any initial information set of a player i, the labels of the outgoing arcs include all the actions of his simultaneousness game, A_i, and

2. if in information set X player i selects an outgoing arc labeled with an action a_i then the next information set of player i, Y, must have an outgoing arc labeled with the action a_i.

Remark 2. *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, since one can move all the random choices to the beginning, to be included as a part of a large initial nature's move, that will be partially and differentially revealed to the players at the appropriate places in the tree, see for example Kreps and Wilson (1982). 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 the modeler, or the players, as in the example above.*

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

Definition 2. ***Constant-play strategies***. *Given an individual strategy σ_i in a Bayesian game G and an extensive version $\overline{G}$, a* constant-play version *of σ_i is any strategy $\overline{\sigma}_i$ in $\overline{G}$ that initially chooses actions with the same probabilities as σ_i and does not modify earlier choices in all subsequent information sets. Formally, at any initial information set of player i the marginal probability that $\overline{\sigma}_i$ selects the action a_i is $\sigma_i(a_i \mid t_i)$, where t_i is the type of player i at the information set. In every non initial information set of player i $\overline{\sigma}_i$ selects, with certainty, the same action that was selected by player i in his previous information set (this is well defined under the perfect recall assumption).*

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

Remark 3. *Constant-play strategies are appealing from bounded-rationality considerations. The complexity of playing such a strategy is low, and moreover, the information (about the extensive game being played), needed in order to play it, is minimal.*

2.3. **Extensively Robust Equilibrium.** A Nash equilibrium σ 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. But we need to define a notion of being approximately robust, where in every $\overline{G}$, every $\overline{\sigma}$ is only required to be an approximate Nash equilibrium . This notion assures that the incentives of any player to unilaterally deviate at any positive probability information set are small, even when such a deviation is coordinated with further deviations in his later information sets.

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 of player i, A, a *modification of player i strategy at A* is any strategy $\overline{\eta}'_i$ with the property that at every information set of player i, B, which is not a follower[4] of the information set A, $\overline{\eta}'_i$ coincides with $\overline{\eta}_i$. Player i can *unilaterally improve his payoff by more than ε at the information set A* if $E_{\overline{\eta}'|A}(u_i) - E_{\overline{\eta}|A}(u_i) > \varepsilon$, where $\overline{\eta}' = (\overline{\eta}_1, \overline{\eta}_2, ..., \overline{\eta}'_i, ..., \overline{\eta}_n)$ for some $\overline{\eta}'_i$ that is a modification of $\overline{\eta}_i$ at A.

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

Note also that not having an ε improvement, as defined above, is in the proper, more restrictive, sense in two ways. We rule out improvements even through coordinated changes at an information set and its follow-up information sets. Moreover,

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

the improvement is small when viewed conditionally on being at the information set, not just from the ex-ante perspective.

Definition 3. ***approximate Nash equilibrium.*** *A strategy profile $\overline{\eta}$ of $\overline{G}$ is an (ε, ρ) Nash equilibrium, if the probability that some player has a better than ε improvement at some information set is not greater than ρ.*

Definition 4. ***approximate robustness.*** *An equilibrium of G, σ, is (ε, ρ)* extensively robust, *if in every extensive version $\overline{G}$ every profile of constant-play version of σ, $\overline{\sigma}$, is an (ε, ρ) Nash equilibrium.*

An equilibrium σ is (ε, ρ) ex-post Nash *if its constant-play version is an (ε, ρ) Nash equilibrium in the extensive version with one round of revisions defined earlier.*

Clearly being ex-post Nash is a weak form of being robust, yet it is the "strongest"[5] form of being ex-post Nash, since every player knows everything, realized types and selected pure actions of all the players, before playing the revision round.

Remark 4. *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-your-favorite-computer is clearly not robust. And this is so despite the fact that it is even a strict Nash equilibrium.*

The notion of being (ε, ρ) ex-post Nash is not monotone in the information that is revealed prior to the possible revision. In other words, even when ε and ρ are small there may be intermediary less-informative (than knowing the precise outcome) events that have a high probability of being observed and after which the incentives to revise are significant. Put differently, the example shows why there is a need for Proposition 1, the main mathematical result established in this paper.

Example 1. ***Insure the car?*** *The following 2-person game is between nature and a risk taker, who has to choose between buying or not buying insurance for his car. Nature has flat preferences of zero, while the payoffs of the risk taker are given by the following table.*

	Nature	
Risk Taker	*accident*	*no accident*
insurance	*0.99*	*0.99*
no insurance	*0*	*1.00*

Consider an equilibrium with nature 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 the accident/no-accident outcome, is not greater than .001.

Suppose, however, that the intermediary event the driver and the car are late *(late) may be observed before it is known whether the car had an accident. Assume* $\Pr(late|accident) = 1$ *and* $\Pr(late|\text{ no accident}) = 0.05$*. 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 there is probability higher than 5% that*

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

the gain from revising is greater than 0.01. Or, in other words, the equilibrium, which is $(0, .001)$ *ex-post Nash, can be no better than* $(.01, .05)$ *extensively robust.*

3. Extensive Robustness in Large Games

3.1. The Main Result.

Two finite universal sets, $\mathcal{T}$ and $\mathcal{A}$, describe respectively *all possible player types* and *all possible player actions* that may appear in the games discussed in this sections. The set $\mathcal{K} \equiv \mathcal{T} \times \mathcal{A}$ denotes the universal set of *all possible player characters.*

We consider a *family* $\Gamma = \Gamma(\mathcal{A}, \mathcal{T})$ that consists of Bayesian games $G = (N, \times T_i, \tau, \times A_i, (u_i))$ with $T_i \subseteq \mathcal{T}$ and $A_i \subseteq \mathcal{A}$.

Definition 5. ***Empirical distribution:*** *For every vector of characters c define the* empirical distribution *induced by c on the universal set of 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 6. ***Semi-anonymity****: A game G is semi anonymous if all its payoff functions are anonymous, i.e., for every player i and for any two profile of 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 7. ***Continuity****: The payoff functions in the family of games* Γ *are uniformly equicontinuous if for every positive* ε *there is a positive* δ *with the following property. For every game in the family, for every player in the game and for any two profiles of 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, this continuity condition as stated, already implies that all the games in Γ are semi-anonymous.

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

3.2. Proof of the Main Result.

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

The first step is the following result from Kalai (2000 & 2003).

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

The above, together with the next proposition, yield directly 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 α and β by bigger positive α' and $\beta', \beta' < 1$, for which $m \alpha \beta^m < \alpha' \beta'^m$ for $m = 1, 2, ...$.

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

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

$\rho \geq \Pr(W)$ 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
$\sum_{A\subset V} \Pr(W|A)\Pr(A) > \zeta \Pr(V)$ by the bounds lemma below.
Combining the above inequalities completes the proof of the proposition. □

Lemma 2. ***Probability coincidence****: The probabilities of any character profile c computed either (1) directly from a strategy profile σ 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 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 at his turn he randomizes just as he does in σ.

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

Lemma 3. ***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. Every information set of player i must be either fully included in A or fully included in the complement of A, since a violation is observed at the information set, without knowledge of the play path that is being played. 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 later information sets being subsets of earlier ones. Thus, if for every play path we take the first information set that exhibit the violation of V, we have a collection of disjoint events whose union equals V. □

Lemma 4. ***Bounds****: for any positive probability information set of player i, A, if player i has better than $\varepsilon+\zeta$ improvement at A, then* $\Pr(W|A) > \zeta$.

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

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

Now we can put a bound on possible level of such improvement as follows. For 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 character profiles in W^c, the largest possible improvement is ε. This means that the highest possible improvement in the information set A is $1\Pr(W|A) + \varepsilon \Pr(W^c|A)$. So if the possible improvement at A is greater than $\varepsilon + \zeta$, we have $\Pr(W|A) + \varepsilon \geq 1\Pr(W|A) + \varepsilon \Pr(W^c|A) > \varepsilon + \zeta$, which validates the claim made above. □

4. Further Examples and Elaboration

4.1. On Types and Anonymity.

The following example shows that semi anonymity is weaker than anonymity in some important ways. The fact that the prior probability over player types and payoff functions can be individualized, allows us to model games that are not anonymous.

Example 2. ***Village or beach*** *; 2n players match pennies:*

The players of the simultaneous move game are n females, labeled as players $1, 2, ..., n$, and n males, labeled as players $n+1, n+2, ..., 2n$, each having to choose between going to the beach (B) or staying in the village (V). A females payoff is the proportion of males her choice ***mis****matches and a males payoff is the proportion of females his choice matches. Thus, for $n = 2$ this is a match pennies game, and the only equilibrium, everybody chooses randomly between B and V, is highly non robust. Failing ex-post Nash stability for example, after every realization some player would want to revise his choice. What about large values of n?*

Unfortunately, this game has non anonymous payoff functions. For example, player $n+1$'s payoff depends on the actions of the players named $1, 2, ..., n$, and is independent of the actions of the other players.

But we can reformulate the above as a Baysian game with anonymous payoff-functions. Let the possible types be female *and* male, *assign players $1, 2, ..., n$ prior probability one of being a* female *type, and players $n+1, n+2, ..., 2n$ prior probability one of being a* male *type. With this, we can define the payoff functions of all players, in the obvious way, to depend on types and actions only, and not on player labels. For example, player $n+1$'s payoff is the proportion of* female types *that his choice matches.*

Under the above formulation the game fits into the model of this paper, and considering the family of all the Beach/Village games with any number $2n$ of players, we may conclude that all its equilibria are essentially robust for sufficiently large values on n.

Thus, for any given ε-deviation from the expected .50-.50 distribution of choices by the females or by the males, the probability of observing such deviation goes to zero at an exponential rate as the number of players increases. If the game is played sequentially, if players get to observe people dressed in bathing suits, buses going to the beach, etc., everybody randomizing with no revisions remains an equilibrium.

In the example above we did not have to have equal number of females and males. But we must have the size of both populations go to infinity, otherwise continuity may fail, as illustrated by an example below.

Example 3. ***Random number of players***: *In the Village/Beach example above, enrich the set of types to be: a female type, a male type, and a non playing type. Assign players* 1 *through* n *positive probabilities of being a female or a non playing type, and players* $n+1$ *through* $2n$ *positive probabilities of being male or a non playing type. A player of the non-playing type has flat payoff of zero. For a female type, the payoff is the proportion of male types her choice mismatches, and for a male type it is the proportion of female types his choice matches.*

Using the types in the manner above, we can fit into the model games with random number of players. One has to be careful though, the continuity of the payoff functions may depend crucially on how the proportions above are defined.

4.2. **On the continuity assumption.** The uniform equicontinuity assumption, is too strong for some important applications.

Example 4. ***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, the only informed player, is told what the realized state of nature is. Thus, we may think of him as having two types,* a *and* b*. All the other players are of one type that knows nothing.*

Player 1 choosing the action corresponding to his type and everybody else randomizing with equal probability between the two choices, is an equilibrium. But obviously it is not extensively robust, not even ex-post Nash. Actually, it becomes less ex-post stable as the number of players n *increases. For example, for large* n*, with probability close to one, approximately 1/2 of the players would want to revise after observing the choice of player 1.*

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 the state* b*, and a non expert. Assign player 1 equal probability of being one of the first two types and every other player probability one of being of the third type, then the payoff of every player may be defined anonymously as follows. Player 1's payoff is one if he chooses according to his type and zero otherwise, every other player's payoff is one if his choice matches the information given to a positive fraction of the realized experts, zero otherwise.*

The failing of extensive robustness is due to the failing of uniform equicontinuity. To see this, consider two sequences of 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* $\overline{c}^n$ *he 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 fact that* $c_2^n = \overline{c}_2^n$ *and that the empirical distributions of player 2's opponents in the two character profiles become arbitrarily close as* $n \to \infty$*.*

While full uniform equicontinuity may fail, still an equilibrium will be extensively robust if there is continuity near its expected play. For example, in voting by majority for one of two candidates, continuity fails when the number of votes is near the .50-.50 distribution. But for other distribution, for example when the expected number of votes is .60-.40 and the voting population is large, extensive robustness holds.

Recall that the proof of the main result was composed of two parts: (1) to show that an equilibrium is (ϵ, ρ) ex-post Nash, with ρ decreasing to zero at an exponential rate, and (2) that such an equilibrium must be $(\epsilon, n\rho/\epsilon)$ robust. The assumptions of anonymity and continuity were used only for the first part. So our main result will continue to hold if we can weaken these assumptions while preserving the conclusion of the first part.

Indeed, in Kalai (2000 & 2003), where the first part is proven, there are two such possibilities. First, it is shown that the continuity assumption does not have to hold globally, but only near the expected play of an equilibrium, as in the .60-.40 voting example. Moreover, less than local continuity is needed. All that one needs is that the strategic dependence of a player on his opponents goes to zero, as the empirical distribution of opponents type and actions becomes a known constant.

4.3. **Independence, rational expectations and learning.** One of the major restrictions of our model is to independent types. The next example shows that such a restriction is necessary.

Example 5. ***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 players payoff is one if he chooses the action that corresponds to the realized (unknown) 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. Conditional on the realized state, independently of each other, every player is told the correct 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. It is also clear that extensive robustness, and even ex-post Nash 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, extensive robustness fails due to the type dependency in the prior probabilities (unconditioned on the state).

The example above shows a sever limitation of our model. In many games, and especially in the rational expectations literature, we have type dependencies due to some unknown state of nature. This is a topic worthy of further research, and we conjecture that while weaker than in the current paper, still meaningful results may be obtained. For example, one may hope to obtain results about fast rate of learning to play ex-post or rational expectations equilibria when we have a repeated Bayesian game with many anonymous players. For an elaboration on this conjecture, see Kalai (2000 & 2003)

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

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

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

Example 6. ***n-Person Battle of the sexes:*** *Simultaneously, each of* n *male players and* n *female players have to choose computer I or computer M. A male's payoff equals the proportion of the total population that he matches if he chooses I, but only* 0.9 *times the proportion he matches if he chooses M. For a female the opposite is true, she is paid the full proportion that she matches when she chooses M and* 0.9 *of the proportion she matches when she chooses I.*

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

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

But let us imagine the same scenario, 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 from the point of view of the first female.

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

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

5. References.

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

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

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

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

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

Green, 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...

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

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

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

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

Kreps,D., and R. Wilson (1982): "Sequential Equilibria," *Econometrica* 50, pp863-894.

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.

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

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

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

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

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

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.

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