

Partially-Specified Large Games

Ehud Kalai

Kellogg School of Management, Northwestern University,
Evanston IL 60208, USA

kalai@kellogg.northwestern.edu

[http://www.kellogg.northwestern.edu/
faculty/kalai/htm/personal-page/index.htm](http://www.kellogg.northwestern.edu/faculty/kalai/htm/personal-page/index.htm)

Abstract. The sensitivity of Nash equilibrium to strategic and informational details presents a difficulty in applying it to games which are not fully specified. Structurally-robust Nash equilibria are less sensitive to such details. Moreover, they arise naturally in important classes of games that have many semi-anonymous players. The paper describes this condition and its implications.

1 Introduction

Deciding whether to attend WINE 2005 is a participation game, where players' payoffs depend on the participation choices of others. But like many other games, it is only partially specified. It is known that players may go to a web site and click in their choice before some deadline, but it is not known in what order they will move, what information they will have, who may make agreements with whom, etc., etc.

Equilibrium analysis of a partially-specified game forces the analyst to make-up all the missing details. But since equilibrium predictions are often sensitive to such details, the analyst's predictions are often unreliable.¹

Equilibria that are structurally robust, as described in this paper, are less sensitive to many game specifics.² As such, they offer a partial resolution to the difficulty above. The analyst can compute the equilibrium of a game with a minimal structure (simultaneous-one-move play), and be assured that it remains equilibrium no matter what the missing details are.

Structural robustness is a strong property that fails in the equilibria of most games. But as discussed in this paper, in games with many semi-anonymous players all the equilibria are structurally robust.³

¹ This high level of sensitivity is well known to researchers and users of game theory. Even small modifications, such as allowing the players to make meaningless cheap-talk announcements, may drastically alter the equilibrium of the game, see for example Crawford and Sobel (1982).

² Other notions of robustness and structural robustness were proposed in economics. We refer the reader to Hansen and Sargent (2001) for some examples. An earlier weaker notion of structural stability was introduced in Kalai (2004) under the name of extensive robustness.

³ Positive results about large games were obtained earlier in the cooperative-games literature, see Aumann and Shapley (1974) for a survey.

In addition to dealing with partially specified games, structural robustness has other important implications, even in fully specified games. For better or worse, an equilibrium with this property is more persistent, it self-purifies, has a strong ex-post Nash property, can be used to model a small-world game embedded in a larger game, and has a strong rational-expectations property in market games.

2 Definition of Structural Robustness

Starting with an n -person simultaneous-one-move game, we first describe a large number of variations on how the game may be played. Such variations allow for sequential (instead of simultaneous) moves, multiple opportunities to revise earlier choices, information transmission, commitments, delegations, etc. The eventual purpose is to identify equilibrium as being structurally robust, if it survives under all such variations.

2.1 Metagames

To describe a variation of the simultaneous move game, G , we introduce the notion of a metagame of G . The concept is natural and straight forward but the formal notations are cumbersome.⁴ For this reason, it is better to explain it through a couple of simple examples.

Consider first a 2-person Match Pennies game (MP for short) with player 1 being a male who wishes to match his opponent, and player 2 being a female who wishes to mismatch her opponent.⁵ The play of the game has the four possible H-P pairs, (H,H), (H,T), (T,H), (T,T), with the corresponding payoffs, (1,-1), (-1,1), (-1,1), (1,-1). A *metagame* of MP is any m -person perfect-recall extensive game \mathcal{M} that has exactly one of the four H-P pairs above associated with every one of its final nodes. In other words, every play of \mathcal{M} yields a play of MP.

Condition 1 *Preservation of strategies and payoffs:* *We restrict ourselves to metagames that satisfy three properties. First, the players of the metagame include all the original players of the underlying game MP.*

Second, any of the original MP players has metagame strategies that guarantee him, in the metagame, the same pure choices that he has in MP. For example, since player 1 can choose H in MP, he should have a metagame strategy that guarantees him (no matter what strategies are used by the opponents) to end up (with probability one) at a final node in which his label is H (the associate pair being (H,H) or (H,T)).

Third, the metagame must preserve the original payoffs in the sense that at every final node of \mathcal{M} , the payoffs of the original players are the same as their payoffs in MP. For example, if the metagame ends at a final node with the associate pair (H,H) then the payoffs of players 1 and 2 are (1,-1).

⁴ As is the case for most formal models that involve extensive form games.

⁵ Simultaneously, each one of them chooses either H or T. If the choices match, she pays him 1; if they mismatch, he pays 1 to her.

Three simple examples of metagames of MP that preserve strategies and payoffs are the following.

Example 1. Sequential-play metagame: player 1 moves first and chooses H or T, player 2 moves second and chooses H or T after being informed of the choice of player 1. The associated pair at a final node consists of the pair of sequential choices made on the path leading to it, and the payoffs are the MP payoffs defined at this pair.

Example 2. Metagame with revisions: simultaneously each of the two players chooses H or T; then, after being informed of the pair of choices, simultaneously they each choose H or T again. The associated pair at every final node of the metagame consists of the choices made in the second round, and the payoffs are the MP payoffs defined at this pair.

Example 3. Metagame with an outsider setting the rules: player 3, the rule setter, moves first and decides whether the game will be played simultaneously, sequentially, or simultaneously with revisions. After being informed of his choice, the original players play according to the rule chosen by him and receive the appropriate payoffs. The metagame preserves strategies and payoffs, no matter what payoffs are assigned to the outsider (the rule setter) at the final nodes.

As the reader should be able to see, one can construct a large number of interesting metagames that preserves strategies and payoffs.

2.2 Metaversions of Strategies

To require that an equilibrium "survive in all metagames," we first must clarify how a strategy of the simultaneous move game is played in a given metagame. This can be done easily due to the requirement that the metagame preserve strategies.

For example, consider the .50-.50 H-T mixed strategy of player 1 (he chooses H or T with equal probability in MP) and consider the metagame with a rule setter.

Let ψ_1 be the metagame mixed strategy in which player 1 chooses with equal probabilities one of the following two pure strategies: (1) he chooses H in every one of his information sets in the metagame, or (2) he chooses T in every one of his information sets in the metagame. This strategy guarantees that he ends up with an H (with probability .50) or guarantees that he ends up with a T (with probability .50).⁶

⁶ This is stronger than simply guaranteeing a .50-.50 outcome on H or T. For example, consider a metagame that starts with nature moving left or right with equal probabilities, and players 1 and 2 follow with a simultaneous choice of H or T after observing nature's choice. To choose H after nature moves left and T after nature moves right is a strategy that guarantees player 1 .50-.50 probabilities over a final label of H or T. But this strategy cannot be decomposed into a .50-.50 lottery over one strategy that guarantees him H and one strategy that guarantees him T.

When such is the case we say that ψ_1 is a *metaversion* of the .50-.50 H-T mixed strategy. However, there are additional metaversions of the .50-.50 H-T mixed strategy. For example, player 1 may mix over the following two pure metagame strategies: (1) he chooses H in every one of his metagame information sets, or (2) he chooses T in every one of his metagame information sets except for the first round of the subgame with revisions where he chooses H.

Having defined metaversions of individual strategies, we can define metaversions of strategy profiles. For example, the strategy profile (ψ_1, ψ_2, ψ_3) in the metagame with a ruler is a *metaversion* of the strategy profile (σ_1, σ_2) in MP, if ψ_1 is a metaversion of σ_1 and ψ_2 is a metaversion of σ_2 (outsiders' strategies are not restricted).

2.3 Structural Robustness

An equilibrium σ of the simultaneous move game is *structurally robust*, if it is an equilibrium in every metagame that preserves strategies and payoffs. This means that no matter what metagame is played, in every metaversion of σ the strategies of the original players are best response.

For example, for an equilibrium $\sigma = (\sigma_1, \sigma_2)$ of MP to be structurally robust, the metaversions of σ must be optimal in every metagame of MP. If in the metagame with a ruler, for example, $\psi = (\psi_1, \psi_2, \psi_3)$ is a metaversion of σ , then ψ_1 and ψ_2 must be a best response strategies of player 1 and player 2 respectively (ψ_3 is not restricted).

We may think of σ as being a uniform Nash equilibrium in every metagame, since any metaversions of (σ_1, σ_2) are optimal no matter what the strategies of the outsiders are.⁷

It is easy to see that MP has no structurally robust equilibrium. Since the underlying game is always a metagame of itself, any structurally robust equilibrium must be equilibrium of the underlying game itself. Thus the only candidate for structural robustness in MP is the profile consisting of the pair of .50-.50 H-T mixed strategies. But this pair fails to be an equilibrium in the sequential play metagame for example. Randomly choosing H or T is no longer a best reply of the follower since she is better off mismatching player 1's observed choice.

2.4 Approximate Structural Robustness

The existence difficulty becomes less severe when the number of players is large. Consider a generalized Match Pennies game that consists of n males and n females, $2n$ MP. Simultaneously, they each choose H or T. A male's payoff is the proportion of females his choice matches and a female's payoffs is the proportion of males she mismatches. The .50-.50 H-T profile of mixed strategies is an equilibrium that turns out to be highly structurally robust when n is large. In the sequential-play metagame with all the males moving first, for example,

⁷ Using the standard fixed point method, for every structurally robust σ and every metagame, every metaversion of σ can be completed to a metagame profile in which all the players (including the outsiders) best respond.

when the females turn to choose comes up, the distribution of male choices will be nearly one half H's and one half T's with probability close to one. Thus randomly choosing between H and T is nearly optimal for every female. (Using Chernoff bounds, for any $\varepsilon > 0$ the probability that the proportion of H-choosing males be outside the $.50 \pm \varepsilon$ range goes down to zero at an exponential rate in the number of males).

The above observation motivates an approximate definition of being (ε, ρ) structurally robust. For a given pair of such non negative numbers we require that the equilibrium of the simultaneous move game be (ε, ρ) Nash equilibrium in every metagame that preserves strategies and payoffs. Being (ε, ρ) Nash equilibrium means that the event "following a play path along which some player can improve his payoff at some information set by more than ε " has probability of at most ρ .⁸

2.5 Generalization to Bayesian Games

Due to their wide applicability, it is important to generalize the above notions to simultaneous-one-move Bayesian games.⁹ The generalizations are straight forward, as illustrated by the following example of an n-person Bayesian MP game (BMP for short).

Every player i is randomly drawn to be a *male* or a *female* type according to commonly known individual prior probabilities. Before they make any choices, everyone is informed (only) of his own realized type. Next, simultaneously, each of the n players chooses H or T. Every player's payoff is a function of his type, his choice, and the distribution of types and choices of his opponents. For example, a *male's* (type) payoff may be the proportion of *females* (female types) his choice matches and a *female's* payoff may be the proportion of *males* she mismatches.

A metagame of BMP is again any m-person perfect recall extensive form game with an n-tuple of H-T's associated with each of its final nodes. But to be compatible with the underlying Bayesian game, it must start with a move of nature, where the types of the individual players are drawn with the same distribution as in BMP, and with every player being informed of his own realized type prior to the start of play.

An equilibrium of BMP is (ε, ρ) structurally robust if it remains an (ε, ρ) Nash equilibrium in every metagame of BMP that preserves strategies and payoffs.

3 Main Result

To keep the presentation simple, we restrict ourselves to games in which player types are restricted to be called *male* or *female* and player actions are restricted to be called H or T. These particular names of types and actions are not important, nor is the fact that there are only two types and two actions. It is important however, that there is a finite number of each.

⁸ See Kalai (2004) for elaboration.

⁹ See Harsanyi (1967/68).

Consider a family of games \mathcal{F} that includes for every $n = 1, 2, \dots$ many n person games (could even be uncountably many games for every n). Every n -person game $G_n \in \mathcal{F}$ is described by a collection of n pairs $G_n = (\tau_i, u_i)_{i=1}^n$ having the following interpretation. The non negative number τ_i is the probability that player i is a *male* type (male for short) and $1 - \tau_i$ is the probability that player i is a *female* type (female for short). It is assumed that types are drawn independently across players.

The payoff function of every player i , u_i , is anonymous in the sense that it may depend only on aggregate data of player i 's opponents. More formally $u_i : (t_i, a_i, e_{-i}) \mapsto r \in [0, 1]$, where $t_i = \textit{male}$ or *female*, $a_i = H$ or T , and $e_{-i} = (e_{m,H}, e_{m,T}, e_{f,H}, e_{f,T})$ describes the proportions of opponents type-choice combinations. For example $e_{m,H}$ describes the proportion of player i 's opponents who are *males* who choose H .

We assume also that the collection of all payoff functions (consisting of all the functions u_i 's from all the games in \mathcal{F}) is uniformly equicontinuous.

Theorem 1. *Structural Robustness.* *Given the family \mathcal{F} above and an $\varepsilon > 0$, there are positive constants α and β , $\beta < 1$, such that for every n person game in the family all the equilibria are $(\varepsilon, \alpha\beta^n)$ structurally robust.*

This theorem is a generalization of a theorem in Kalai (2004). Even though the result is substantially stronger, the proof is similar and is therefore omitted (it requires that the number of player types and possible actions be uniformly bounded by finite numbers. This could possibly be replaced by assumptions of continuity and compactness).

3.1 Some Clarifications

The anonymity imposed on the individual payoff functions is far from being full anonymity, since players may have individual payoff functions and individual prior distributions over types. While this already means that the players are not anonymous individuals, it even breaks some of the anonymity in the payoff functions.

For example, we may have a formulation that includes a type whose name is *Mr.Jones*, and in a certain game only player 1 is the *Mr.Jones* type (he has prior probability one of being so and every other player has probability zero of being so). Now, while the players payoffs are anonymous in the technical sense formulated above, the payoffs of the players may depend on the action of player 1 in a non symmetric way. For example player 2 payoff may depend entirely on the distribution of actions chosen by *Mr.Jones* types, thus making player 2 payoff be a function of player 1's action. This method of breaking anonymity is only partial, however, since we can accommodate only finitely many types and the number of players becomes infinite.

Another consequence of the individualized payoff functions and prior probabilities is that the family of games above may contain many games that are drastically different from the Match Pennies game that we started with. For example, it may contain two players, both female type with probability one, and

each has a payoff function that is one if and only if she matches the opponent. This is a pure coordination game.

The Bayesian aspect of the formulation allows for interesting games. For example player 1 may be a female who wishes to match, but player 2 may have a positive prior operability of being a male who wishes to mismatch and a positive prior probability of being a female who wishes to match. Thus, player 1 may not know exactly whether he is in a coordination game or a match pennies game.

In a similar way, the family may contain prisoners' dilemma games of various sizes. When we allow n types, all n person payoff functions can be accommodated.

4 Implications

4.1 Partially Specified Games

In what sense does the above theorem help the modeling of a partially specified game? The analyst may write such a game with a minimal structure, as a simultaneous one-move game, and compute its equilibrium. This equilibrium will survive even if the structure is made richer.

For example, in a large anonymous participation game, where the choice of every player is whether or not to participate in an event, the analyst may write the game where simultaneously every player decide whether to participate or not, and compute its equilibrium. The structural robustness theorem implies that this equilibrium will be sustained no matter how the simultaneous move and informational assumptions are changed. So even if choices are made sequentially, private and public messages being transmitted according to any dynamic stochastic process, players are allowed to repeatedly revise earlier choices, players are allowed to interact and make use of outsiders, under all such possibilities the computed equilibrium is not destroyed.

However, the resolution is only partial. While all the equilibria of the simultaneous-one-move game are sustained in all extensive metagames, each metagame may have additional equilibria that are not present in the simultaneous one-move game. This puts us into another typical difficulty of game theory, namely multiplicity of equilibria. Whether one equilibrium will be sustained over others depends on focal point and related considerations.

The simultaneous-one-move equilibria have an added advantage over the others: they are the only equilibria present in the intersection of all metagames, since the simultaneous one-move game is one of the metagames of itself.

4.2 Ex-post Stability

The structural robustness property implies other game theoretic properties of the equilibrium. One important consequence is *strong ex-post stability* (also known as ex-post Nash).¹⁰ This means that after the game is over, even with

¹⁰ See Cremer and McLean (1985), Green and Laffont (1987) and follow up literatures, for earlier (weaker) versions of the concept.

(full or partial) hind-sight information about the types and choices of the others, no player has incentive to change (or regrets) the choice he made.

A vector of types and actions is ε ex-post Nash if for every player and for every outcome of the game, even with perfect hind-sight information about the realized types and selected actions of all his opponents, the player cannot gain an ε -improvement in his payoff by a unilateral change of his own action.

An equilibrium is (ε, ρ) *ex-post Nash* if the probability of ending up with a vector of types and actions which is ε ex-post Nash is at least $1 - \rho$.

To see that (ε, ρ) structural robustness implies (ε, ρ) ex-post Nash consider the metagame with revision discussed earlier (they play the one shot game once, observe everybody's realized types and choices, and get to revise their choices in a second round). For any equilibrium of the simultaneous move game, consider the following metaversion. Play the equilibrium in the first round, with no revisions in the second round. The structural robustness property implies that this metaversion of the equilibrium is an equilibrium of the metagame with revision. This, however, is equivalent to being ex-post Nash (also in the (ε, ρ) adjusted senses of both).

The ex-post stability is stronger than just described, because it holds even for partial information. It may be that no player can improve his payoff by more than ε when given complete information about the types and choices of his opponents, but he can improve his *expected* payoff by more than ε when he conditions on partial information about his opponents. Being strongly (ε, ρ) ex-post Nash means that this is not the case. No matter what hind-sight (perfect or imperfect) information is given to individual players, with probability at least $1 - \rho$ no one can improve his expected payoff by more than ε by a unilateral change of his selected action.

4.3 Self Purification

The ex-post Nash property may alternatively be viewed as a property of self purification. Starting with Schmeidler (1973), there is a large literature showing that large anonymous games have pure strategy Nash equilibrium. Schmeidler's paper starts by proving the existence of "mixed strategies" equilibrium, and then proving that he can purify it, i.e., replace the mixed strategies by pure ones without destroying the equilibrium property.

But when an equilibrium is (fully) ex-post Nash (in the sense defined in this paper), the profiles of pure strategies that can be generated by its play must all be Nash equilibria of the game. In other words, there is no need to purify the equilibrium since it purifies itself (through laws of large numbers).

And going beyond the purification literature of Schmeidler (which studies only simultaneous-move normal-form games), self purification holds for simultaneous-move Bayesian games. The play of any Bayesian equilibrium which is ex-post Nash must generate, for every vector of realized types, a profile of pure actions which is a Nash equilibrium of the normal-form game determined by the realized types.

There are some technical differences between the two approaches. Schmeidler and follow up papers model large games by assuming a continuum of players, while the current paper models it asymptotically by letting the number of players grow to infinity. A difficulty in Schmeidler's model is that mixed strategy profiles are hard (border line impossible) to define, since they involves a continuum of independent mixed strategies. A limitation of the current paper is that the full ex-post Nash property holds only in the limit. The reader may verify though, that the approximate notion of (ε, ρ) ex-post Nash gives rise to a natural approximate notion of self purification which becomes fully so in the limit.

4.4 Small World in a Bigger Game

Game theorists model games in isolation, ignoring strategic and informational spillovers between the game and the outside world in which the game is played. This often leads to incorrect analysis since such spillovers can easily destroy equilibria of the isolated game. It is therefore important to identify conditions under which an equilibrium persist, even when embedded in a bigger world.

Since metagames allow additional players, additional actions, back and forth communication and the like, they provide a model of a bigger world in which the underlying game may be embedded. Moreover, the condition of structural robustness is precisely the one that guarantees the survival of the equilibrium under such embedding. This makes the sufficient-conditions for structural robustness important in the discussion of the embedding issue.

Clearly, the requirement that the metagame preserve strategies and payoffs limits the scope of possible embeddings. But it is important to note that if the outside world in which the isolated game is played does not preserve its strategies and payoffs, then the isolated game should not be studied in the first place. It is simply not the game that exists in reality.

Assuming, therefore, the metagame does preserve strategies and payoffs, as a general rule it still does not preserve the equilibria of the isolated game, unless the isolated game is large as assumed in our model.

We may conclude that the equilibria of large semi-anonymous games are persevered when embedded in the outside world, provided that the embedding preserve strategies and payoffs. We should keep in mind that while the equilibria are preserved, the presence of the outside world may introduce new additional equilibria, not identified by the analysis of the isolated game.

4.5 Implementation

Mechanism designers deal with games in which the natural equilibrium may be inefficient, and aim to create new related games in which equilibria are efficient. The new game has the same set of possible types outcomes as the original game, but it is redesigned to have "good incentives." In the language of the current paper, mechanism designers often replace an underlying "bad game" by a metagame, and by doing so they *implement* a better equilibrium.

To a significant extent the implementation literature studies the possible improvements that a mechanism designer may attain under various restrictions

on the possible or permissible implementing metagame. The condition that a metagame preserves strategies and payoffs is interesting in this regard. It means that the designer cannot invade the options of the participants, by disallowing them some choices or by modifying the consequences of the original choices.

Under this restriction, the structural robustness property has an interesting implication, since it limits the ability of the mechanism designer to eliminate the bad equilibrium. This means that he must resort to more subtle means of implementation. He can only hope to find a metagame in which the original bad equilibrium becomes non appealing. For example, he may be able to construct a metagame in which the original bad equilibrium involves the use of dominated strategies, or has some other non appealing aspects from focal point considerations.

The observations above may be important to real-life policy makers and system designers. Languages, measurement systems, and keyboard choices are examples illustrating that highly inefficient social equilibria are likely to persist even despite social attempts to replace them by more efficient ones.

4.6 Rational Expectations

When restricting ourselves to market games, the strong ex-post Nash property of structurally robust equilibrium implies that the equilibrium has a strong rational expectations property.¹¹ This may be illustrated by considering a Bayesian version of a Shapley-Shubik game, where players of random types (representing possible initial endowments, information and preferences) submit, in a one-simultaneous-move game, portions of their initial endowments to be traded in the market. Final market prices are computed as a function of the submitted quantities and net trades are executed by these final prices.¹²

In general, Nash equilibrium of such a game fails to satisfy the rational expectations property economists expect from an equilibrium. This property requires that every agent's trade be optimal at the market prices, given his individual type and the inference (bases on the observed prices) he may make about the unknown types of the others. This cannot be expected to hold for the Shapley-Shubik players since their strategies are based entirely on knowledge of their individual parameters, without knowledge of the choices made by the others (and the consequential market prices).

But an equilibrium which is strongly ex-post Nash turns out to automatically satisfy the rational expectations property. Prices, in the Shapley-Shubik game, are ex-post information. Thus, under a strong ex-post Nash condition, knowing the prices gives no player incentives to revise his choices. This means that his trade is optimal given knowledge of the prices in the same sense required by rational expectations.

The intuition is fairly straight forward. When the game has many traders, with unknown types but known prior distributions over types, laws of large numbers allow the traders to correctly anticipate the final market prices which

¹¹ For examples of economic and game theoretic implications see Lucas (1972), Grossman and Stiglitz (1980), Jordan and Radner (1982) and Forge and Minelli (1998).

¹² See Shapley and Shubik (1977) and Peck (2003).

depends on aggregated data of the opponents. Therefore the ex-post Nash property and the rational expectation property both hold.¹³

The fact that the ex-post Nash property is strong implies a strong rational expectations property. The trade every player ends up with is optimal given his information, any inference he may make through the realized prices, but also any inference he may make from any other information he may acquire.

References

- Aumann, R.J., and L.S. Shapley (1974): *Values of Nonatomic Games*. Princeton: Princeton University Press.
- Crawford, V., and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431-51.
- Cremer, J., and R.P. McLean (1985): "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53, 345-61.
- Forges, F., and E. Minelli (1998): "Self-fulfilling mechanisms in Bayesian games," *Games and Economic Behavior*, 25, 292-310.
- Green, J.R., and J.J. Laffont (1987): "Posterior Implementability in a Two-person Decision Problem," *Econometrica*, 55, 69-94.
- Hansen, L.P. and T. Sargent (2001): "Robust Control and Model Uncertainty," *American Economic Review*; 91, 60-66.
- Harsanyi, J. (1967/68): "Games with Incomplete Information Played by Bayesian Agents," *Management Science*, 14, 159-182, 320-334, 486-502.
- Grossman, S.J. and J.E. Stiglitz (1980): "On the Impossibility of Informationally Efficient Markets," *American Economic Review*, 70, 393-407.
- Jordan, J.S., and R. Radner (1982): "Rational Expectations in Microeconomic Models: An Overview," *Journal of Economic Theory*, 26, 201-223.
- Kalai, E., (2004): "Large Robust Games," *Econometrica*, 72, 1631-66.
- Lucas, R.E. (1972): "Expectations and the neutrality of money," *Journal of Economic Theory*, 4, 101-124.
- McLean, R., and A. Postlewaite (2002): "Informational Size and Incentive Compatibility," *Econometrica*, 70, 2421-54.
- Peck, J., (2003): "Large Market Games with Demand Uncertainty," *Journal of Economic Theory*, 109, 283-299.
- Schmeidler, D., (1973): "Equilibrium Points of Nonatomic Games," *Journal of Statistical Physics*, 17, 295-300.
- Shapley, L.S. and M. Shubik (1977): "Trade Using One Commodity as a Means of Payment," *Journal of Political Economy*, 85, 937-68.

¹³ The above arguments require continuity of the payoff functions under small changes in the distribution of opponent characteristics. This is so because in the game above players are small in their strategic influence and in the information they possess (see for example McLean and Postlewaite (2002)). When individual players possess substantial inside information, continuity fails and we do not obtain the required ex-post Nash and the consequential rational expectations properties.