Stability in Large Bayesian Games with Heterogeneous Players*

Joyee Deb† and Ehud Kalai‡

Abstract

Bayesian Nash equilibria that fail to be hindsight (or alternatively ex-post) stable do not provide reliable predictions of outcomes of games in many applications. We characterize a family of Bayesian games in which all equilibria are asymptotically hindsight-stable, and discuss the consequences of this robustness property. In contrast to earlier literature, we establish hindsight stability in a class of games in which players are not anonymous and where type spaces and action spaces can be infinite.

* **Acknowledgements**. The authors would like to thank Michael Kearns, Colin McDiarmid and Eran Shmaya for useful comments and suggestions. We also thank seminar participants at GAMES 2008, Northwestern University and SAET meetings in Singapore for helpful comments. This is a substantially revised version of Chapter 3 of the first author's Ph.D. dissertation submitted to Northwestern University in June 2008.

† New York University, Stern School of Business, Department of Economics; joyee.deb@nyu.edu

‡ Northwestern University, Kellogg School of Management, Managerial Economics and Decision Sciences Department; kalai@kellogg.northwestern.edu

1 Introduction

Hindsight stability, also referred to as ex-post robustness, is an important property of equilibria of one-shot simultaneous-move Bayesian games. An equilibrium of a Bayesian game is hindsight-stable if no player has an incentive to change her action after she learns the realized types and the realized actions of all the players. Equilibria that are not hindsight-stable may not provide useful predictions of outcomes in games in which players can revise choices based on hindsight information, and indeed, a significant number of social interactions are of this type (see Section 1.1 for examples).

There is a substantial literature devoted to issues of stability in games with many players, or large games. Kalai (2004) shows that all equilibria in a certain class of large games are *structurally robust*, a stronger property that implies hindsight stability.[1] Unfortunately, however, the restrictions on the class of games in Kalai (2004) leave out many important real-life situations. The first restriction is that the number of possible types and actions of the players be finite. Second, when a player evaluates her payoffs, she is restricted to viewing all her opponents as being anonymous or indistinguishable. Both these assumptions severely limit the applicability of the results.

The main objective of this paper is to characterize an important family of large games not restricted by assumptions of finiteness and anonymity and, yet, in which all equilibria are still hindsight-stable.

We consider a general family of Bayesian games where players' types and action spaces are compact subsets of finite dimensional Euclidean spaces, and in which players' types are drawn independently. We assume that payoffs can depend on types and actions of specific rivals in asymmetric ways, but subject to two regularity conditions: "Uniform Lipschitz Continuity in one's own types and actions", and "Uniform Scaled Lipschitz Continuity in rivals' types and actions." In the main result of this paper, we show that these two conditions are sufficient to guarantee that all equilibria of these games become approximately hindsight-stable at an exponential rate as the number of players increases.

The regularity conditions we impose on the payoff functions are interesting in their own right. In particular, the Uniform *Scaled* Lipschitz continuity with respect to rivals implies that a player has a *limited individual impact* on her opponents, and that the impact decreases with the total number of players in the game.[2] Limited individual

[1] Kalai (2004) illustrated that a property called extensive stability holds for a certain family of large games. But Kalai (2005 and 2008) argues that these games satisfy a stronger property, called structural stability.

[2] Different bounds on the informational size or the influence of players have been presented in earlier literature. See, for example, Al-Najjar and Smorodinsky (2000), McLean and Postlewaite (2002) and references

impact does not imply anonymity and, as we illustrate in examples below, is far from assuming that players become approximately symmetric.

From a mathematical point of view, the limited individual impact condition allows us to apply a powerful law of large numbers: McDiarmid's Bounded Differences Inequality. This law and its variants may be useful in a variety of economic settings.[3]

The proof of the main result is done in two steps. Using McDiarmid's inequality, in the first step, we show that the convergence to hindsight stability is true for games with a *finite* number of possible types and actions. Then, approximating a general game by its *finite grid reduction game*, we show that the convergence holds in general.

Hindsight stability relates to other important issues in economic applications: It implies that mixed strategies self-purify (see Kalai (2004) and Cartwright and Wooders (2006)); it implies a strong rational expectations property in certain market games (see Kalai (2004, 2008)); and it implies that the revelation principle holds in implementation problems (see Green and Laffont (1987)). Robustness properties of equilibria are of interest in computer science, in the areas of distributed communications and computing (see Halpern (2003)). In particular, robust equilibria perform better in systems that involve asynchronous communications (see Kalai (2005)), and in protocols that involve faulty behavior (see Gradwohl and Reingold (2008)).

Related current work on equilibrium robustness in large games includes Gradwohl and Reingold (2010), who study robustness of Bayesian equilibria in games that allow certain correlations between players' types, and Carmona and Podczeck (2010), who study large games with infinite type and action spaces but with anonymity. Other recent work on large games includes Azrieli and Shmaya (2010), who study purification (not self-purification) in large non-anonymous complete information games, and Bodoh-Creed (2010), who studies implementation in large games.

1.1 Illustrative Examples

Before we go into the formal model and results, we first present some examples that illustrate the importance of hindsight stability in games where revisions are possible.

Example 1. ***A Location game with men and women:*** *Simultaneously, n male players, $m_1, m_2, ...m_n$, and n female players, $f_1, f_2, ..., f_n$, choose locations, l_{m_j} and l_{f_i}, in the interval $[0, 1]$. Men want to be close to the women, but women want to be far from*

therein.

[3]Unlike classical laws of large numbers which deal with the expected value of the *average* of random variables, this law deals with the expected value of any function of the variables in which the impact of individual variables is limited.

the men. Specifically, each female player's utility is $u_{f_i} = \frac{1}{n}\sum_{j=1}^{n}|l_{f_i} - l_{m_j}|$*, whereas each male player's utility is* $u_{m_i} = 1 - \frac{1}{n}\sum_{j=1}^{n}|l_{m_i} - l_{f_j}|$.[4]

At a symmetric equilibrium of this game, all the men choose to locate at the point $\frac{1}{2}$, and each woman chooses between 0 and 1 with equal probability.[5] Clearly, with a small number of players, say $n = 3$, such an equilibrium is not hindsight-stable. Ex-post, after seeing the locations of the women, every man would have a non-negligible incentive to move to the side with the majority of the women. In fact, if there is a small odd number of players and moving is not costly, no matter what equilibrium strategies are played, the situation is "never stable." For example, after the men move to the side where the most women are, the women would move to the opposite side, to which the men would respond, and so on.

To overcome this difficulty, we may try to model the situation as a continuous-time location game in which each player's location at time t is a function of the history of all past locations. But general continuous-time models are difficult to formulate and often require unreasonably strong rationality on the part of players. For these reasons, it may be better to identify *myopic stationary equilibria* in which no player wants to change her choice, assuming that her opponents will stay with their choices (this is especially appealing when the number of players is large). In general, as in the above example, such stationary equilibria may fail to exist. However, by definition, any equilibrium that is (approximately) hindsight-stable in the simultaneous-move one-shot game yields (approximate) myopic stationary equilibria of the continuous-time game. Thus, the issue of existence and identification of (approximate) myopic stationary equilibrium is partially resolved by the existence and identification of an (approximately) hindsight-stable equilibrium.

The main result of this paper implies that asymptotically, as n becomes large, all the equilibria of the location game above become hindsight-stable. Specifically, for an arbitrarily small positive number ε, the probability of the event that "some player can gain more than ε by unilaterally revising her action ex-post" decreases to zero at an exponential rate as n becomes large.

In the next example, we consider a game with heterogeneous players. In particular, this example illustrates that a game with heterogeneous players cannot be approximated by games with anonymous players, even asymptotically, as the impact of the individual players vanishes.

---

[4]This is a generalization of the Village vs Beach game (Kalai (2004)) with a continuum of actions.

[5]At other equilibria, all the men choose $\frac{1}{2}$, and every woman chooses, purely or randomly, one of the two extremes. For odd n, some females must randomize.

Example 2. ***Heterogeneous payoff functions.*** *Consider a location game as above, but with players g_i, $g = 1, 2, ..., G$, for some fixed integer G, and $i = 1, 2, ..., n$ (nG players in total). The index g denotes groups e.g., genders, tribes, nationalities, races or combinations of such characteristics and i names the player within a group. Each player chooses a point l_{g_i} in a set C_{g_i}, subset of some fixed compact set in $\mathbb{R}^m$.*

Players have heterogeneous payoff functions that depend on the identity of their specific opponents. A simple example may be a player 5_3 who values the proximity to other players in a exponentially decreasing order in group similarity: $u_{5_3} = 3 - \sum_{g=1}^{G} (\frac{1}{2})^{|g-5|} \sum_{i=1}^{n} \frac{1}{n} d(l_{5_3}, l_{g_i})$, where $d(x, y)$ is the distance between x and y.

The payoff function u_{5_3} satisfies the regularity conditions required in our setting (the two Lipschitz Continuity conditions discussed earlier). In particular, the rate of change of u_{5_3}, as one changes the location of any one opponent, decreases to zero as the number of players increases. Thus, players have limited individual impact. But it is also easy to see that, though individual opponents become of negligible importance, they are far from being equally unimportant. No matter how large n is, for player 5_3, the average location of the players in her group (group 5) is twice as important as the average location of players in group 6, which is twice as important as the average location of players in group 7, etc. Regardless of n, when considering her choice, she is interested mainly in being close to players of her own and similar groups and has little interest in the location of players of highly dissimilar groups.

The main result of this paper implies that in this example, as the number of players increases, all the equilibria become hindsight-stable.[6]

The heterogeneity of types in the general model of this paper goes significantly further than in the example above. As already noted, different players may have different payoff functions. So, while player 5_3 above prefers to be close to members of her group, player 5_4 may prefer to be far away from them. Also, while player 5_3 here views all opponents in any group as anonymous, in general, this is not required: Her payoff function does not have to depend only on her distance from other opponents, and does not have to treat opponents within each group in a symmetric fashion.

Example 3. ***Bayesian Cournot game****: The players are n sellers, each having the capacity to produce up to k units of an identical divisible product. The price-quantity relationship is described by the demand function: $p = 1 - q/n$.*[7] *Each seller i knows*

[6] For the model and results of this paper, the number of groups in Example 2 may also be made arbitrarily large. This would show, formally, that the earlier finite anonymous model of Kalai (2004) cannot accommodate the games in this paper (even if the number of actions were finite).

[7] A simple interpretation is that there are n buyers, each with a demand function $q = 1 - p$.

his per-unit production costs c_i i.e., his type, which is randomly drawn by a commonly known prior probability distribution μ_i over some fixed interval of real numbers I_i. With knowledge of his own type, each player decides on a production level $0 \leq x_i \leq k$, resulting in a profile (c, x), of individual costs and quantities. The resulting payoff of player i is $u_i(c, x) = x_i(1 - \sum_{j=1}^{n} x_j/n - c_i)$.

In Bayesian Cournot games, pure strategy equilibria may not exist.[8] But even when an equilibrium exists, if the number of players is small, then the produced quantities are likely to be hindsight unstable. Initially, a player does not know the opponents' production costs and the resulting opponents' production levels; therefore, her decision is based on the distribution of these variables. Yet, with hindsight, after seeing the produced quantities of all the opponents, she would most likely regret her choice.

If the production process were ongoing, with the x_is representing production rates, and if changes in the production rates were relatively inexpensive, then most producers would adjust the production rates after observing the opponents' choices. Thus, the Bayesian Cournot prediction of the simultaneous-move one-shot market game would fail to reflect the actual outcomes.

Our result establishes that if there is a large number of producers (and buyers), any Bayesian equilibrium is approximately hindsight-stable. Thus, the production levels obtained at equilibria of the simultaneous-move one-shot Cournot game will give no player a significant reason to revise her choice.

Two additional properties are direct consequences of our result. First, being hindsight-stable means that the realized pure actions will constitute an (approximate) equilibrium of the complete information game determined by the realized types. In this sense, the Bayesian equilibrium is (asymptotically) self-purifying.[9] Second, in every Bayesian equilibrium, the realized actions will satisfy a rational expectations property to yield realized prices that are competitive: The individual quantities chosen initially will be optimal at the realized prices. While obtaining competitive equilibrium as the limit of Cournot games is not new,[10] here, it is an immediate consequence of our main result in the more general context of Bayesian games. Moreover, it is not obtained at a particular equilibrium, but is the (asymptotic) property of every Bayesian equilibrium.

[8] See discussion and references in Einy et al. (2010).

[9] This property is significantly stronger than (simple) purification, which originated the study of large strategic games. See Schmeidler (1973) and the follow-up literature. Under simple purification, one simply establishes the existence of a pure strategy equilibrium.

[10] See Mas-Colell (1983), Novshek and Sonnenschein (1983) and follow-up literature.

2 Model

Let $\mathscr{T}$ denote the space of feasible types of players, and let $\mathscr{A}$ denote the space of all feasible actions of players. We assume that $\mathscr{T}$ and $\mathscr{A}$ are fixed subsets of some Euclidean spaces. We will consider a family $\Gamma = \Gamma(\mathscr{T}, \mathscr{A})$ of Bayesian games $G(N, T, A, \tau, \{u_i\})$ that can each be described as follows.

- There are N players $\{i = 1, ..., N\}$.
- The type of each player i is drawn *independently* from a type-space T_i, where T_i is a compact subset of $\mathscr{T}$. Players are informed about their own type. Let T denote the type space of all players i.e. $T := \times_{i=1,...,N} T_i$. Let τ be a probability measure on the Borel subsets of T. Under the independence assumption (assumed throughout the paper), τ is the product measure of its marginal distributions on each T_i, denoted by τ_i. The type distributions are common knowledge.
- Each player i chooses actions from her action space A_i which is a compact subset of $\mathscr{A}$. Let A denote the space of action profiles of all players i.e. $A := \times_{i=1...,N} A_i$.

  We refer to a pair $(t_i, a_i) := c_i$ as a *type-action character* of player i (often abbreviated to *character*). Denote the space of a player's *type-action characters* $T_i \times A_i$ as $\mathscr{C}_i$. Denote the space of *type-action character profiles* as $\mathscr{C} = \times_i \mathscr{C}_i$. For any player i, we denote the space of type-action characters of her rivals as $\mathscr{C}_{-i} := \times_{j \neq i} T_j \times A_j$.
- Players' payoffs are given by bounded measurable functions $u_i : \mathscr{C} \mapsto \mathbb{R}$. For convenience, we sometimes use the derived functions $u_i^{c_i} : \mathscr{C}_{-i} \mapsto \mathbb{R}$ where $u_i^{c_i} = u_i(c_i, c_{-i})$.

Definition 1 (**Uniformly Bounded Strategy Space**). *A family of games* $\Gamma = \Gamma(\mathscr{T}, \mathscr{A})$ *is said to have a uniformly bounded strategy space if* $\mathscr{T}$ *is a compact subset of* $\mathbb{R}^{k_T}$ *and* $\mathscr{A}$ *is a compact subset of* $\mathbb{R}^{k_A}$ *for some fixed integers* $k_T, k_A > 0$.

We define the standard $L1$-metric to be $d(x, y) = \sum_{m=1}^{M} |x_m - y_m|$ for any $x, y \in \mathbb{R}^M$. We consider families of games with uniformly bounded strategy spaces, and with payoff functions that satisfy the following two regularity conditions.

Definition 2 (**LC1: Uniform K-Lipschitz Continuity in One's Own Character**). *Given* $K \geq 0$*, the payoff functions* u_i *in a family of games* $\Gamma(\mathscr{T}, \mathscr{A})$ *are said to be uniformly* K*-Lipschitz continuous in one's own character, if for every player* i*, any* *character profile c and any type-action character of player i, c'_i,*

$$|u_i(c_i, c_{-i}) - u_i(c'_i, c_{-i})| < K\, d(c_i, c'_i),$$

where $d(.,.)$ is the L1 metric.

Condition LC1 is a form of Lipschitz Continuity in one's own types and actions. As an aside, note that the equivalence of the L1 and Euclidean norms implies that the Lipschitz condition with respect to the L1 metric implies Lipschitz Continuity in the Euclidean metric.

Definition 3 (**LC2: Uniform Scaled L-Lipschitz Continuity in Rival Character Profile**). *Given $L \geq 0$, the payoff functions u_i in a family of games $\Gamma(\mathscr{T}, \mathscr{A})$ are said to be Uniformly Scaled L-Lipschitz Continuous in the rivals' type-action character, if for every N-player game in $\Gamma(\mathscr{T}, \mathscr{A})$, for every player i, for any c, c'_{-i},*

$$|u_i(c_i, c_{-i}) - u_i(c_i, c'_{-i})| < \frac{L}{N-1}\, d(c_{-i}, c'_{-i})$$

where $d(.,.)$ is the L1 metric.

Notice that the Lipschitz bound L is uniform for all N in the family of games.

Definition 4 (**Pure Strategy**). *A pure strategy of a player i is a measurable function $s_i : T_i \mapsto A_i$.*

Defining mixed strategies as maps from types to mixtures over pure strategies has the drawback that it is not well defined in games with a continuum of types (see Aumann (1964)). We use the notion of distributional strategies as introduced by Milgrom and Weber (1985). As Milgrom and Weber establish, a distributional strategy is simply another way of representing mixed and/or behavioral strategies.[11] Please see Milgrom and Weber (1985) for more on the correspondence between behavioral strategies, mixed strategies and distributional strategies.

Definition 5 (**Distributional Strategy**). *A distributional strategy for player i is a probability measure σ_i on the Borel subsets of $T_i \times A_i$ for which the marginal distribution of T_i coincides with τ_i. Formally, for any $S \subset T_i$, $\sigma_i(S \times A_i) = \tau_i(S)$.*

[11] There is a simple correspondence between behavior strategies and distributional strategies. A behavior strategy can be defined as function $\beta_i : A_i \times T_i \mapsto [0,1]$ such that (i) for every $B \subset A_i$, the function $\beta_i(B, \cdot) : T_i \mapsto [0,1]$ is measurable and (ii) for every $t_i \in T_i$, the function $\beta_i(\cdot, t_i) : A_i \mapsto [0,1]$ is a probability measure. For any distributional strategy σ_i, the regular conditional distributions $\sigma_i(B|t_i)$ defined over the Borel subsets B of A_i can be understood as behavior strategies $\beta_i(B, t_i) = \sigma_i(B|t_i)$. Conversely, for any behavioral strategy β_i, the corresponding distributional strategy σ_i is defined for all Borel subsets of $T_i \times A_i$ by $\sigma_i(S \times B) = \int_S \beta_i(B, t_i \tau_i(dt_i)$

When players use distributional strategies, the expected payoff of player i is defined as follows:

$$U_i(\sigma) = \int u_i(c) d\sigma(c),$$

where, for a profile of distribution strategies, we let $\sigma = (\sigma_1, \ldots, \sigma_N)$ denote the product distribution over $\times(T_i \times A_i) = C$.

Definition 6 (Equilibrium). *A profile of (distributional) strategies σ^* is an equilibrium*[12] *if*

$$\text{For all } i, \sigma_i', \qquad U_i(\sigma_1^*, \ldots \sigma_i^*, \ldots \sigma_N^*) \geq U_i(\sigma_1^*, \ldots \sigma_i', \ldots, \sigma_N^*).$$

Since we prove robustness properties of equilibria asymptotically, as the number of players increases, we need to define a notion of approximate equilibrium.

Definition 7 (ε-Best Response). *Let $\varepsilon > 0$ (small). A strategy σ_i^* is an ε-best response for player i to σ_{-i} if for every positive-probability set of player i's types, $\tilde{T}_i \subset T_i$ (with $\tau_i(\tilde{T}_i) > 0$), the following condition holds: $U_i((\sigma_i', \sigma_-i)|\tilde{T}_i) - U_i((\sigma_i^*, \sigma_{-i})|\tilde{T}_i) \leq \varepsilon$ for all distributional strategies σ_i' of player i.*

Definition 8 (ε-Bayesian Equilibrium). *A strategy profile σ^* is an ε-Bayesian equilibrium if each σ_i^* is an ε-best response to σ_{-i}^*.*

Finally, we introduce the appropriate notion of robustness, which we call *hindsight stability*. A type-action profile is hindsight-stable for a player if she does not want to change her action even after she observes the realized types and actions of her rivals. Formally, we define approximate hindsight stability as follows.

Definition 9 (Approximate Hindsight Stability). *A type-action character profile $c = (t, a)$ is ε-hindsight-stable for player i if $u_i((t_i, \tilde{a}_i), c_{-i}) - u_i(c) \leq \varepsilon$ for all $\tilde{a}_i \in A_i$. The type-action character profile is ε-hindsight-stable, if it is ε-hindsight-stable for all players. A strategy profile σ^* is (ε, ρ)-hindsight-stable if it yields hindsight-stable type-action profiles with probability at least $1 - \rho$.*

Note that an equilibrium is approximately hindsight-stable if the realized actions, not the mixed actions, constitute an approximate Nash equilibrium of the realized complete information game.

[12] In the context of this paper, we are not concerned about the existence of equilibria, since our objective is to establish stability properties of equilibria where they exist. As as aside, note that independence of types and uniform continuity of payoff functions in a game deliver existence of an equilibrium in distributional strategies.

3 Games with Limited Individual Impact

It will be useful now to examine more carefully the regularity conditions LC1 and LC2 and clarify their role in our setting. Condition LC1 is a technical assumption. It ensures that we can approximate the games we study (with infinite type-action spaces) by an appropriate finite game. Condition LC2 is a more interesting assumption. It turns out that LC2 and uniform boundedness of the strategy space, together, imply an important strategic property of our class of games that we call "limited individual impact (LII)." The LII property means that the effect that any player can unilaterally have on an opponent's payoff is uniformly bounded and decreases with the number of players in the game.

The LII property is used to establish convergence to hindsight stability of equilibria as the number of players increases. It turns out that LII is, in fact, technically quite useful as it implies that a certain law of large numbers (Mcdiarmid's Bounded Differences Inequality) holds, which is used to establish convergence. To the best of our knowledge, this law of large numbers has not been used before in the economics literature, and we believe that it can be extremely useful in the analysis of large games.[13] We define the LII property formally, and establish its relationship with LC2.[14]

Definition 10 (**λ-Limited Individual Impact**). *Given $\lambda \geq 0$, players in games in $\Gamma(\mathscr{T}, \mathscr{A})$ are said to have λ-limited individual impact if the set of payoff functions $\{u_i\}$ satisfies the following condition:*

For all i, for all type-action characters c, c', $|u_i(c_i, c_{-i}) - u_i(c_i, c'_{-i})| \leq \dfrac{\lambda}{N-1}$,

whenever c_{-i}, c'_{-i} differ only in one coordinate.

As discussed in Example 2 earlier, the LII property does not imply anonymity, as a player's utility is not affected just by the aggregate actions of her rivals. The condition requires only that as the number of players increases, the impact of any individual player should vanish. However, the impact of any player on another can still be potentially asymmetric and dependent on the type and identity of the rival, even with a large number of players. Further, no restrictions are imposed on the difference

[13] Recently, some related concentration results like Hoeffding's inequality have been used in the learning literature, see for instance, Cesa-Bianchi and Lugosi (2006). McDiarmid's inequality is a further generalization.

[14] A similar property called "maximal impact" is defined by Azrieli and Shmaya (2010) in recent work related to purification of equilibria in games with complete information. They establish the existence of pure approximate equilibrium in games with sufficiently small maximal impact.

that a player can make to her own payoffs by changing her action choice unilaterally.

The lemma below shows that LC2 and uniformly bounded strategy spaces imply that the payoff functions satisfy the limited individual impact condition.

Lemma 1. *Let $\Gamma(\mathscr{T}, \mathscr{A})$ be a family of games with a uniformly bounded strategy space and satisfying the Uniform Scaled L-Lipschitz Continuity in rival character profiles. Then, there exists a constant $\lambda \geq 0$ such that players in $\Gamma(\mathscr{T}, \mathscr{A})$ satisfy the λ-limited individual impact condition.*

Consider any game in $\Gamma(\mathscr{T}, \mathscr{A})$, and fix any player i. For any two type-action character profiles of player i's rivals c_{-i} and c'_{-i} that differ only in one coordinate, the distance between the two profiles is less than some upper bound B. The existence of such an upper bound follows from the uniform bounded strategy space of $\Gamma(\mathscr{T}, \mathscr{A}, L)$. Now the Uniform Scaled L-Lipschitz Continuity (LC2) implies that for any type-action character of player i, the difference in player i's payoffs at c_{-i} and c'_{-i} can be, at most, $\frac{LB}{N-1}$. This implies that the λ-limited individual impact condition holds with $\lambda = LB$.

4 Hindsight Stability in Large Games

The main result of the paper is a robustness result. If we consider a family of games satisfying our regularity conditions, then any Bayesian equilibrium in this family is approximately hindsight-stable if the game is played by a large number of players. The formal statement of the result is as follows:

Theorem 1 (**Hindsight Stability in Large Games**). *Consider a family of games $\Gamma(\mathscr{T}, \mathscr{A})$ with a uniformly bounded strategy space and satisfying regularity conditions LC1 and LC2. Given $\varepsilon > 0$, there exist constants $\alpha = \alpha(\Gamma, \varepsilon)$, $\beta = \beta(\Gamma, \varepsilon) < 1$ such that if σ^* is an equilibrium of a game $G(N, T, A, \tau, \{u_i\}) \in \Gamma$ then σ^* is $(\varepsilon, \alpha\beta^{N-1})$ hindsight-stable.*

The rest of this section is devoted to proving this theorem. We first prove a slightly stronger version of the theorem for the restricted case of a finite number of actions and types: all approximate equilibria become hindsight-stable as the number of players increases. Then we show that a game with an infinite number of actions and types may be approximated by a *finite grid reduction* game, in which the number of types and actions is finite. The regularity conditions LC1 and LC2 are used to show this, and also to show that the equilibria of the infinite game are approximate equilibria in the finite grid reduction games. We then use the fact that approximate equilibria of finite games are approximately hindsight-stable to deduce that the result is also true for equilibria of the original infinite game.

4.1 Games with Finite Type-Action Spaces

In this subsection, we restrict attention to games with finite type and action spaces. Consider a family of games $\bar{\Gamma}(\bar{\mathscr{T}}, \bar{\mathscr{A}})$ comprising games $G(N, \bar{T}, \bar{A}, \tau, \{u_i\})$ where $\bar{\mathscr{T}}$ and $\bar{\mathscr{A}}$ are finite type and action spaces and $\bar{T}$ and $\bar{A}$ are subsets of $\bar{\mathscr{T}}$ and $\bar{\mathscr{A}}$ respectively. Suppose that $\bar{\Gamma}(\bar{\mathscr{T}}, \bar{\mathscr{A}})$ satisfies LC2. We establish the following result.

Theorem 2 (Hindsight Stability in Finite Games). *Consider a family of finite Bayesian games $\bar{\Gamma}(\bar{\mathscr{T}}, \bar{\mathscr{A}})$ that satisfies LC2. Given $\varepsilon > 0$, there exist positive constants $\alpha = \alpha(\bar{\Gamma})$, $\beta = \beta(\bar{\Gamma}, \varepsilon) < 1$ such that if σ^* is any η-equilibrium of any game $G \in \bar{\Gamma}$, then σ^* is $(2\eta + \varepsilon, \alpha\beta^{N-1})$ hindsight-stable.*

The above theorem is a robustness result for Bayesian equilibria in finite games, and, to the best of our knowledge, the first such robustness result in an environment without anonymity.[15] To prove this theorem, we use some intermediate results.[16]

Lemma 2 (McDiarmid's Independent Bounded Differences Inequality). *Let $X = (X_1, X_2, \ldots, X_M)$ be a family of independent random vectors with X_k taking values in $\mathscr{X}_k$ for each k. Suppose that there exist constants l_k $(k = 1, \ldots, M)$ such that the real-valued function g defined on $\times \mathscr{X}_k$ satisfies*

$$|g(x) - g(x')| \le l_k \text{ whenever } x \text{ and } x' \text{ differ only in the } k^{th} \text{ coordinate}.$$

Let μ be the expected value of the random variable $g(x)$. Then, for any $t \ge 0$,

$$Pr(g(x) - \mu \ge t) \le e^{\frac{-2t^2}{\sum_{k=1}^{M} l_k^2}}.$$

[15]We show that LC2 and uniformly bounded strategy spaces are sufficient conditions for hindsight stability in finite games. We leave unanswered the question of whether the conditions are necessary. We believe however, that for a hindsight stability result, some form of "continuity" cannot be dispensed with. The rough intuition is that a continuity condition is what can tie together different games in the general family that we consider. We do not restrict attention to replicating games: The family Γ may contain different types of games for each possible N, and we need a continuity condition to guarantee the asymptotic result regardless of the type of game.

[16]It is worthwhile to note that Theorem 2 can be proved by assuming the Limited Individual Impact condition, instead of LC2. The proof would also be identical to the one presented here.

Proof. This result and the proof can be found in McDiarmid (1989).[17] $\square$

A straightforward application of the result in our framework yields the next lemma. Recall some notation: for any player i and character $c_i = (a_i, t_i)$, we defined the function $u_i^{c_i}(c_{-i}) := u_i(c_i, c_{-i})$. Further, we denote player i's expected payoff given her type-action character by $\mu_i^{c_i}$ i.e., $\mu_i^{c_i} = \mathbb{E}[u_i^{c_i}(c_{-i})]$.

Lemma 3. *Let σ be a strategy profile of $G(N, \bar{T}, \bar{A}, \tau, \{u_i\}) \in \bar{\Gamma}$ that satisfies condition LC2. Then, there exists $\lambda > 0$ such that*

$$\textit{For all } i, c_i, \qquad Pr_{\sigma_{-i}}(|u_i(c) - \mu_i^{c_i}| > \alpha) \leq 2e^{\frac{-2(N-1)\alpha^2}{\lambda^2}}.$$

Proof. For any player i, a realized type-action character profile of her opponents is a sequence of independent random vectors $(C_1, ..., C_{i-1}, C_{i+1}, ..., C_N)$. By Lemma 1, we know that the family of games satisfies the λ-limited individual impact condition i.e., we can find a constant λ such that

for all $i,$ for all $c_{-i} \in \mathscr{C}_{-i}$, whenever c_{-i}, c'_{-i} differ only in one coordinate,

$$|u_i(c_i, c_{-i}) - u_i(c_i, c'_{-i})| \leq \frac{\lambda}{N-1}.$$

Fix $\alpha > 0$ and any c_i. Applying Lemma 2 to the function $u_i^{c_i}$, we get the result. $\square$

Lemmas 2 and 3 make the role of "limited individual impact" transparent. The Independent Bounded Differences inequality tells us that the probability that a function of independent random variables deviates from its mean by any quantity t is inversely proportional to the maximum impact that each random variable has on the value of the function. In our setup, limited individual impact means that any player's impact on a rival's payoffs is bounded and, moreover, decreases with N. This implies, in turn, that the probability of deviations from the mean vanishes exponentially fast with N.

Proof of Theorem 2: Suppose that σ^* is an η-equilibrium of an N-player game G in $\bar{\Gamma}(\bar{\mathscr{T}}, \bar{\mathscr{A}})$ satisfying LC2. Let $\varepsilon > 0$ be given. Fix any player i and a type-action character $c_i = (a_i, t_i)$. Define $D_{i,c_i} = \{c_{-i} : |u_i(\tilde{c}_i, c_{-i}) - u_i(c)| > 2\eta + \varepsilon$ for some $\tilde{c}_i$ with $\tilde{t}_i = t_i\}$. In other words, D_{i,c_i} denotes the rival type-action character profiles where player i with character $c_i = (a_i, t_i)$ can unilaterally change her action choice and gain more than $2\eta + \varepsilon$. Let D_i denote $\bigcup_{c_i} D_{i,c_i}$, and let $D = \bigcup_i D_i$.

[17] We are grateful to Michael Kearns and Colin McDiarmid for pointing us to this result.

We will show that the probability of D goes to zero. Now,

$$|u_i^{\tilde{c}_i}(c_{-i}) - u_i^{c_i}(c_{-i})| \leq |u_i^{\tilde{c}_i}(c_{-i}) - \mu_i^{\tilde{c}_i}| + |\mu_i^{\tilde{c}_i} - \mu_i^{c_i}| + |\mu_i^{c_i} - u_i^{c_i}(c_{-i})|.$$

If the first and last term are both less than $\frac{\eta+\varepsilon}{2}$ and the second term is less than η, then the left-hand side is less than $2\eta + \varepsilon$.

We know that, by the limited individual impact property and Lemma 3, for any characters c_i and $\tilde{c}_i$, we have

$$Pr_{\sigma_{-i}}\left(|u_i^{\tilde{c}_i}(c_{-i}) - \mu_i^{\tilde{c}_i}| \leq \frac{\eta+\varepsilon}{2}\right) > 1 - 2e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}},$$

and similarly,

$$Pr_{\sigma_{-i}}\left(|u_i^{c_i}(c_{-i}) - \mu_i^{c_i}| \leq \frac{\eta+\varepsilon}{2}\right) > 1 - 2e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}.$$

Further, since σ is an η-equilibrium, we also have $|\mu_i^{\tilde{c}_i} - \mu_i^{c_i}| \leq \eta$. It follows that,

$$Pr_{\sigma_{-i}}\left(|u_i^{\tilde{c}_i}(c_{-i}) - u_i^{c_i}(c_{-i})| \leq 2\eta + \varepsilon\right) > 1 - 4e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}.$$

Put differently, $Pr_{\sigma_{-i}}\left(|u_i^{\tilde{c}_i}(c_{-i}) - u_i^{c_i}(c_{-i})| > 2\eta + \varepsilon\right) \leq 4e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}$. It follows that, at a given realized character of player i, the probability that any deviation yields more than $2\eta + \varepsilon$ gain for player i is, at most, $4|C|e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}$, where $|C|$ is the total number of possible type-action character profiles i.e., $Pr_{\sigma_{-i}}(D_i) \leq 4|C|e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}$. Finally, the probability that any player can profit by more than $2\eta + \varepsilon$ by deviating i.e., $Pr(D)$, is at most, $4N|C|e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}$. Equivalently, if the number of players is N, the strategies are $\left(2\eta + \varepsilon, 4N|C|e^{-\frac{\frac{(\eta+\varepsilon)^2}{2}(N-1)}{\lambda^2}}\right)$ hindsight-stable. Note that the second term vanishes as $N \to \infty$. So, we have found constants $\alpha' = 4|C|$ and $\beta' = e^{-\frac{(\eta+\varepsilon)^2}{2\lambda^2}} < 1$ such that all equilibria are $\left(2\eta + \varepsilon, N\alpha'\beta'^{N-1}\right)$ hindsight-stable. Notice that we can replace the constants α' and $\beta' < 1$ by bigger constants α and $\beta < 1$ so that $N\alpha'\beta'^{N-1} < \alpha\beta^N$ for all N. This completes the proof. $\square$

4.2 Finite Grid Reductions

Now we show that Theorem 2, which establishes robustness of Bayesian equilibria in games of limited individual impact with finite type and action spaces, indeed generalizes to the infinite case. Consider a Bayesian game $G(N, T, A, \tau, \{u_i\})$ where the type and action spaces T and A are compact subsets of uniformly bounded type and action spaces $\mathscr{T}$ and $\mathscr{A}$. Fix $\Delta > 0$ small. We define a new game $\hat{G}$ that we call the finite grid reduction. We need some additional definitions and notation.

For every $x \in \mathbb{R}$, define $r(x) := \hat{x} = \Delta \max\{k : k \in \mathbb{Z} \text{ such that } k\Delta \leq x\}$. For any $x \in \mathbb{R}^m$, define $r(x) := \hat{x} = (\hat{x}_1, \ldots, \hat{x}_m)$. .

Associated with the compact type space T, we can define a finite type-space $\hat{T}_i = \{r(t_i) : t_i \in T_i\}$, and $\hat{T} = \times T_i$. Similarly, associated with the compact action space A, we define a finite action space $\hat{A}_i = \{r(a_i) : a_i \in A_i\}$, and $\hat{A} = \times A_i$. Define the corresponding type-action character spaces $\hat{C}_i = \hat{T}_i \times \hat{A}_i$ and $\hat{C} = \times \hat{C}_i$. Corresponding definitions can be applied to $\mathscr{T}$ and $\mathscr{A}$ as well.

Associated with the prior marginal distributions of types τ_i, define distributions $\hat{\tau}_i$ over $\hat{T}_i$ such that $\hat{\tau}_i(\hat{t}_i) = \tau_i(r^{-1}(\hat{t}_i))$. As usual, $\hat{\tau} = (\hat{\tau}_1, \ldots, \hat{\tau}_n)$. For any character profile $\hat{c}$, define player i's payoff $\hat{u}_i(\hat{c}) = \frac{1}{2}\left[\sup_{c \in r^{-1}(\hat{c})} u_i(c) + \inf_{c \in r^{-1}(\hat{c})} u_i(c)\right]$. As before, let $\hat{U}_i(\sigma)$ denote the expected payoff of player i from strategy profile σ.

We call the new game defined by $\hat{G} = (N, \hat{T}, \hat{A}, \hat{\tau}, \hat{u})$ the finite Δ-grid reduction of the original game $G = (N, T, A, \tau, u)$. Compactness implies that $\hat{G}$ is a game with finite type and action spaces. For any strategy σ_i in the original game G, we can define the associated strategy $\hat{\sigma}_i$ in the Δ-finite reduction $\hat{G}$ by $\hat{\sigma}_i(\hat{c}_i) = \sigma_i(r^{-1}(\hat{c}_i))$. For a strategy profile σ of G, $\hat{\sigma} = (\hat{\sigma}_1, \ldots, \hat{\sigma}_n)$. Below, we establish some results about the relationship between a game G and its finite grid reduction $\hat{G}$.[18]

Lemma 4. *Let $G(N, T, A, \tau, \{u_i\})$ be a Bayesian game where T and A are compact subsets of uniformly bounded type and action spaces $\mathscr{T}$ and $\mathscr{A}$ (with k_T and k_A being the bounds on the dimensionality of $\mathscr{T}$ and $\mathscr{A}$ respectively), and suppose that the payoff functions u_i satisfy LC1 and LC2 (with constants K and L, respectively). Fix $\Delta > 0$ small. Then, for any strategy profile σ of game G and the associated strategy profile $\hat{\sigma}$ in the Δ-finite grid reduction $\hat{G}$ and for any positive probability type $\hat{t}_i$ of player i, the following is true: $|\hat{U}_i(\hat{\sigma}|\hat{t}_i) - U_i(\sigma|r^{-1}(\hat{t}_i))| \leq (k_T + k_A)(K + L)\Delta$.*

[18] We prove the main result, Theorem 1, using the notion of finite grid reductions, as we believe this makes transparent the role of each of the assumptions LC1 and LC2. However, a more direct proof can also be obtained using LC1 and LC2.

Proof. Let $\hat{c} \in \hat{C}$. For any $c, c' \in r^{-1}(\hat{c})$, the conditions LC1 and LC2 imply that

$$\begin{aligned}
&|u_i(c_1, \ldots, c_n) - u_i(c'_1, \ldots c'_n)| \\
&\leq |u_i(c_1, \ldots, c_n) - u_i(c'_1, c_2 \ldots c_n)| + |u_i(c'_1, c_2 \ldots, c_n) - u_i(c'_1, c'_2 \ldots c_n)| + \ldots \\
&\qquad\qquad\qquad\qquad \ldots + |u_i(c'_1, \ldots, c_n) - u_i(c'_1, c'_2 \ldots c'_n)| \\
&\leq K(k_T + k_A)\Delta + \tfrac{L}{N-1}(k_T + k_A)\Delta + \ldots + \tfrac{L}{N-1}(k_T + k_A)\Delta \\
&\leq (k_T + k_A)(K + L)\Delta.
\end{aligned}$$

Then, the result follows from the definition of $\hat{u}_i$ and $\hat{U}_i$. ☐

Proposition 1. *Let $G(N, T, A, \tau, \{u_i\})$ be a Bayesian game where T and A are compact subsets of uniformly bounded type and action spaces $\mathscr{T}$ and $\mathscr{A}$ (with k_T and k_A being the bounds on the dimensionality of $\mathscr{T}$ and $\mathscr{A}$ respectively), and suppose that payoff functions u_i satisfy LC1 and LC2 (with constants K and L, respectively). Then, for every equilibrium σ of G, the associated strategy $\hat{\sigma}$ is a $2(k_T + k_A)(K + L)\Delta$-equilibrium of the Δ-finite reduction game $\hat{G}$.*

Proof. This follows almost immediately from Lemma 4 above. For every player i, type $\hat{t}_i$ and strategy $\hat{\sigma}'_i$, we have

$$\begin{aligned}
&\hat{U}_i(\hat{\sigma}'_i, \hat{\sigma}_{-i}|\hat{t}_i) - \hat{U}_i(\hat{\sigma}|\hat{t}_i) \\
\leq\ & U_i(\sigma'_i, \sigma_{-i}|r^{-1}(\hat{t}_i)) + (k_T + k_A)(K + L)\Delta - U_i(\sigma|r^{-1}(\hat{t}_i)) + (k_T + k_A)(K + L)\Delta \\
\leq\ & 2(k_T + k_A)(K + L)\Delta,
\end{aligned}$$

where the second inequality follows from the fact that σ is an equilibrium of G. ☐

Proposition 2. *Let $G(N, T, A, \tau, \{u_i\})$ be a Bayesian game where T and A are compact subsets of uniformly bounded type and action spaces $\mathscr{T}$ and $\mathscr{A}$, and payoff functions satisfy conditions LC1 and LC2 (with constants K and L respectively). Let $\hat{G}$ be the Δ-finite reduction of G. If $\hat{\sigma}$ is a (η, ρ) hindsight-stable equilibrium of the Δ-finite reduction game $\hat{G}$, then the associated strategy σ of the game G is $(\eta + (k_T + k_A)(K + 2L)\Delta, \rho)$ hindsight-stable in G.*

Proof. Let $\hat{S} = \{\hat{c} :$ no player can gain more than η by deviating from $\hat{c}\}$. Let $S = r^{-1}(\hat{S})$. By definition of the finite grid reduction, $Pr_{\hat{\sigma}}(\hat{S}) = Pr_{\sigma}(S)$. Further, since $\hat{\sigma}$ is (η, ρ) hindsight-stable, we know that $Pr_{\hat{\sigma}}(\hat{S}) \geq 1 - \rho$. Now, if a player cannot gain by more than η by deviating at $\hat{c}$, then at any $c \in r^{-1}(\hat{c})$, she cannot improve by more than $\eta + (k_T + k_A)(K + 2L)\Delta$. This follows from Lemma 4. Therefore, the probability that some player can deviate and make a gain of more than $\eta + (k_T + k_A)(K + 2L)\Delta$ must be less than ρ. ☐

4.3 Proof of Main Result

The proof of Theorem 1 is now straightforward, using the above results. Consider a family of games $\Gamma(\mathscr{T}, \mathscr{A})$ where the strategy space is uniformly bounded and payoff functions satisfy LC1 and LC2 with Lipschitz constants K and L, respectively. Fix $\varepsilon > 0$, small. Consider an equilibrium σ^* of a game G in Γ. Now, consider $\hat{G}$ a Δ-finite reduction of G where $\Delta = \frac{\varepsilon}{2(k_T+k_A)(5K+6L)}$. We know from Proposition 1 that the associated equilibrium $\hat{\sigma}^*$ of $\hat{G}$ must be a $2(k_T + k_A)(K + L)\Delta$-equilibrium of $\hat{G}$ i.e., $\hat{\sigma}^*$ is an $\frac{\varepsilon(K+L)}{(5K+6L)}$-equilibrium of $\hat{G}$. We can then apply our result about hindsight stability in finite games to the finite grid reduction. Theorem 2 implies (for $\frac{\varepsilon}{2} > 0$) that there are constants $\alpha(\Gamma, \varepsilon)$ and $\beta(\Gamma, \varepsilon) < 1$ such that $\hat{\sigma}^*$ is $(\frac{2\varepsilon(K+L)}{5K+6L} + \frac{\varepsilon}{2}, \alpha\beta^{N-1})$ hindsight-stable in $\hat{G}$.[19] Finally, we apply Proposition 2 to show that the original equilibrium σ^* is $(\varepsilon, \alpha\beta^{N-1})$ hindsight-stable. This establishes the main theorem.

5 Concluding Remarks

Little was known about the stability properties of Bayesian equilibria in games with infinite types and actions, or in games that are not anonymous. In this note, we study a class of Bayesian games in which the type and action spaces are infinite and players are not anonymous. We impose two regularity conditions on the payoff functions–variants of Lipschitz Continuity–and show that this is enough to guarantee hindsight stability of Bayesian Nash equilibria if the number of players is large. Notice that provide sufficient conditions for hindsight stability. It would be interesting to investigate if, and in what sense, these conditions are necessary as well.[20]

A promising line of research may be to study other ways of modifying the notions of stability and relaxing the restrictions on the family of games. For instance, we may consider an alternate notion of hindsight stability. The notion of hindsight stability in this note requires that no player have an incentive to revise her action after learning *all* the information about the types and actions of the opponents. For full hindsight stability, this means also that players would not have an incentive to revise their choices after learning *partial* information about their opponent's types and actions. However, as illustrated by an example in Kalai (2004), this is not the case for approximate (ε, ρ)-

[19] Note that the value of the constants α and β depend on the particular finite grid reduction game, and therefore depend on ε.

[20] For instance, we use condition LC1 to extend the robustness result from finite games to infinite games, using the finite grid approximation method. It may be interesting to ask if there is a more direct proof that perhaps does not use LC1, but some alternate weaker condition.

hindsight-stability. Even if an equilibrium is (ε, ρ)-hindsight-stable in the sense of this paper, there may still be partial information that reaches a player ex-post with probability greater than ρ, and based on this information, the player could increase her payoff by more than ε by a unilateral ex-post deviation. Strong (ε, ρ)-hindsight-stability of an equilibrium (see Kalai (2004)) requires that there be no significant probability for ex-post unilateral revisions that lead to meaningful gains by individual players even after they obtain partial hindsight information about the play of the game. It is not clear whether this stronger condition holds for the family of games presented in this paper, or if other stronger properties such as structural robustness hold. We leave these issues out of the current paper, as addressing them would require the formulation of extensive versions of the game, which is significantly more difficult with a continuum of types and actions. It is also not clear if weaker robustness properties of Bayesian Nash equilibrium would hold if types were not independent or if we had payoff functions that displayed discontinuities (see Gradwhol and Reingold (2010), for example). This would be a line of investigation that is particularly important for applications.

References

Al-Najjar, N. and Smorodinsky, R.: 2000, Pivotal Players and the Characterization of Influence, *Journal of Economic Theory* **92**(2), 318–342.

Aumann, R.: 1964, Mixed vs. Behavior Strategies in Infinite Extensive Games, *Annals of Mathematics Studies* **52**, 627–650.

Azrieli, Y. and Shmaya, E.: 2010, Lipschitz Games. Ohio State University Mimeo.

Bodoh-Creed, A.: 2010, Approximation of Large Games. Cornell University Mimeo.

Carmona, G. and Podczeck, K.: 2010, Ex-post Stability of Bayes-Nash Equilibria of Large Games. Mimeo.

Cartwright, E. and Wooders, M.: 2009, On Purification of Equilibrium in Bayesian Games and Ex-post Nash Equilibrium, *International Journal of Game Theory* **38**(1), 127–136.

Cesa-Bianchi, N. and Lugosi, G.: 2006, *Prediction, Learning and Games*, Cambridge University Press.

Deb, J.: 2008, *Essays in Economic Theory Chapter 3: Large Games with Limited Individial Impact*, PhD thesis, Northwestern University, Evanston, Illinois.

Dembo, A. and Zeitouni, O.: 1998, *Large Deviations Techniques and Applications*, Vol. 38 of *Applications of Mathematics*, Springer.

E. Einy, O. Haimanko, D. M. and Shitovitz, B.: 2010, On the Existence of Bayesian Cournot Equilibrium, *Games and Economic Behavior* **68**(1), 77–94.

Fudenberg, D. and Tirole, J.: 1991, *Game Theory*, MIT Press.

Gradwohl, R. and Reingold, O.: 2008, Fault Tolerance in Large Games, *Proceedings of the 9th ACM Conference on Electronic Commerce*, pp. 274–283.

Gradwohl, R. and Reingold, O.: 2010, Partial Exposure in Large Games, *Games and Economic Behavior* **68**(2), 602–613.

Green, J. R. and Laffont, J. J.: 1987, Posterior Implementability in a Two-Person Decision Problem, *Econometrica* **55**, 69–94.

Halpern, J.: 2003, A Computer Scientist Looks at Game Theory, *Games and Economic Behavior* **45**(1), 114–131.

Kalai, E.: 2004, Large Robust Games, *Econometrica* **72**(6), 1631–1665.

Kalai, E.: 2005, Partially-Specified Large Games., *Lecture Notes in Computer Science*, Vol. 3828, pp. 3–13.

Kalai, E.: 2008, Large Games: Structural Robustness, *in* S. Derlauf and L. Blume (eds), *The New Palgrave Dictionary of Economics*, 2nd edn, Palgrave Macmillan.

Mas-Colell, A.: 1983, Walrasian Equilibria as Limits of Noncooperative Equilibria, Part i: Mixed Strategies, *Journal of Economic Theory* **30**, 153–170.

McDiarmid, C.: 1989, On the Method of Bounded Differences, *in* J. Siemons (ed.), *Surveys in Combinatorics*, Vol. London Mathematical Society Lecture Note Series 141, pp. 148–188.

McDiarmid, C.: 1998, *Probabilistic Methods for Algorithmic Discrete Math.*, Springer Verlag, chapter Concentration, pp. 1–54.

McLean, R. and Postlewaite, A.: 2002, Informational Size and Incentive Compatibility, *Econometrica* **70**, 2421–2453.

Milgrom, P. and Weber, R.: 1985, Distributional Strategies for Games with Incomplete Information, *Mathematics of Operations Research* **10**(4), 619–632.

Novshek, W. and Sonnenschein, H.: 1978, Cournot and Walras equilibrium, *Journal of Economic Theory* **19**, 223–266.

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