Ex-Post Stability of Bayes-Nash Equilibria of Large Games*

Guilherme Carmona† Konrad Podczeck‡

University of Cambridge Universität Wien

June 22, 2011

Abstract

We present a result on approximate ex-post stability of Bayes-Nash equilibria in semi-anonymous Bayesian games with a large finite number of players. The result allows players' action and type spaces to be general compact metric spaces, thus extending a result by Kalai (2004).

Journal of Economic Literature Classification Numbers: C72

Keywords: Bayes-Nash equilibrium; large games; ex-post stability; concentration.

*We thank two anonymous referees for very helpful comments. Financial support from Fundação para a Ciência e a Tecnologia is gratefully acknowledged. Any remaining errors are, of course, ours.

†Faculty of Economics and King's College, University of Cambridge and Nova School of Business and Economics. Address: University of Cambridge, Faculty of Economics, Sidgwick Avenue, Cambridge, CB3 9DD, UK; Phone: (44) 122 333 5225; email: gmc42@cam.ac.uk.

‡Address: Institut für Wirtschaftswisenschaften, Universität Wien, Hohenstaufengrasse 9, A-1010 Wien, Austria. Email: konrad.podczeck@univie.ac.at

1 Introduction

A Bayes-Nash equilibrium of a Bayesian game is ex-post Nash, also called ex-post stable, if with probability one, the realization of types and actions are such that the profile of actions is a Nash equilibrium of the complete information version of the game determined by the realized profile of types. Similarly, a Bayes-Nash equilibrium is called approximately ex-post Nash, or approximately ex-post stable, if, with a probability approximately equal to one, the profile of actions is an approximate equilibrium of the complete information version of the game determined by the profile of types. Thus, ex-post stability, both in its exact and in its approximate versions, can be interpreted as a no-regret condition since rarely will there be a player who can gain more than a small amount by changing his action after observing the type and the action of the other players.

Approximate ex-post stability has been shown by Kalai (2004) to be a property of Bayes-Nash equilibria of sufficiently large semi-anonymous Bayesian games, where semi-anonymity means that each player's payoff is affected by the profiles of type-action characters of the other players only through the distributions of these profiles. More precisely, it is shown in Kalai (2004) that if a family of semi-anonymous Bayesian games is fixed such that the set of payoff functions in this family of games satisfies some uniform equicontinuity condition with respect to the distributions of the profiles of feasible type-action characters of the rivals of a player, then Bayes-Nash equilibria of games in this family are approximately ex-post Nash if the number of players in the game is large enough. This result by Kalai (2004) assumes finite action and type spaces.

In this paper, we extend Kalai's ex-post stability result by allowing action and type spaces of the players to be general compact metric spaces. The main motivation of our work is that several game theoretic applications consider models where players' types or actions may vary continuously in infinite sets. Standard examples are location games or Cournot games in which quantities can be adjusted continuously. The framework of general compact metric spaces also covers action and type spaces that are infinite dimensional subsets of function spaces. This case also arises naturally in some models. We illustrate this in two examples given in Section 3. The first example is a Cournot oligopoly game where the types of the players are given by their cost functions and these functions may vary in an infinite dimensional compact subset of the space of all real-valued functions on the action space. The advantage of this formalization is that it avoids restricting firms' cost functions to be specified by a finite number of parameters (e.g., a constant marginal cost and a fixed cost), thus allowing for a rich set of possibilities for firms' cost functions.

Our second example is a variation of the Village versus Beach example in Kalai (2004). The original example by Kalai (2004) considers a situation where each individual in a group of people, with an equal number of men and women, is asked whether or not he or she would like to go the beach, the alternative being staying in the village. Each man's payoff is the proportion of women that his choice matches and each woman's payoff is the proportion of men that her choice mismatches. Our example is similar, but instead of a single question, there are countably infinitely many questions to which players can answer "yes" or "no". This leads naturally to an infinite dimensional action space, namely the space of all sequences of numbers zero or one, where one stand for answering "yes", and zero for answering "no".

Another example (not worked out in detail in this note) which leads to compact metric action spaces that are not given as subsets of Euclidean spaces would be a Cournot oligopoly game placed into a context of commodity differentiation modeled as in Jones (1984). In the model of Jones (1984), there is an infinite compact metric space Ω of differentiated commodities, which are assumed to be divisible, and the commodity bundles are the elements of the space $M^b(\Omega)$ of bounded Borel measures on Ω. Commodities that are close as points in Ω are interpreted as similar. The space $M^b(\Omega)$ is endowed with the weak*-topology[1], so that two commodity bundles are topologically close not only if they contain the same commodities in similar quantities, but also if they contain similar commodities in similar quantities, a notion of closeness

[1]This topology on $M^b(\Omega)$ is the topology of pointwise convergence on the continuous functions on Ω, evaluation being given by integration.

of commodity bundles that is natural in a context of commodity differentiation. Now in a Cournot game in such a setting of differentiated commodities, the action sets of the firms, i.e., the sets of their potential output vectors, would be closed subsets of $M^b(\Omega)$, and assuming capacity constraints, these sets would also be bounded. As Ω is a compact metric space, closed and bounded subsets of $M^b(\Omega)$ are compact and metrizable, so firms would have compact metric action spaces. Since the differentiated commodities are assumed to be divisible, it would also be natural to allow for the possibility that there is no finite upper bound on the number of different commodities a firm actually produces, so that, even when the set Ω of commodities is given as a subset of some Euclidean space, the action sets of the firms would not need to have a finite dimensional representation.

An ex-post stability result that, like ours, allows for infinite action and type spaces has been obtained by Deb and Kalai (2010) (see Deb (2008) for an earlier version of that paper). Our setting differs from that of Deb and Kalai (2010) in several important aspects. On the one hand, our setting is more restrictive since, unlike Deb and Kalai (2010), we assume semi-anonymous payoff functions. On the other hand, in the framework of Deb and Kalai (2010) the action and type spaces of players are restricted to be compact subsets of some Euclidean space and it is assumed that, for the $\|\cdot\|_1$-metric and some constant M, each player's payoff function is uniformly M-Lipschitz continuous in his type-action character and uniformly $M/(N-1)$-Lipschitz continuous in the profiles of the type-action characters of the other players, where N is the number of players in the game. In contrast, our result allows for general compact metric action and type spaces, and the set of payoff functions in a given family of games has to satisfy a uniform equicontinuity condition, but these functions need not satisfy Lipschitz conditions. In fact, in our result, a uniform equicontinuity condition is imposed only with respect to changes of a player's payoff resulting from changes of type-action characters of the other players; in particular, payoff functions are allowed to be discontinuous in the owner's type-action character.

Allowing for payoff functions that are discontinuous in the owner's type-action character is an advantage of our result over that of Deb and Kalai (2010) in addition to allowing for infinite dimensional type and action spaces. An example where this additional flexibility may be important is provided again by the Cournot oligopoly framework. Indeed, a typical feature of industrial production are set-up costs, i.e., when the level of production is expanded, then at some levels, costs may jump upwards since increasing the level of production may require, at those levels, additional machines to be acquired or the production facility to be expanded. To cover this properly in a model in which the level of production may vary continuously, one has to allow for discontinuous cost functions, and hence for payoff functions that are discontinuous in the owners' actions.[2]

On the technical side, our paper shares with Kalai (2004) and Deb and Kalai (2010) the use of results on the concentration of probability measures. This aspect also makes our paper related with Azrieli and Shmaya (2011) where such results are used to address the existence of approximate pure strategy equilibria in games with complete information.

It is reasonable to expect that our result can be extended to a more general framework. In particular, we expect that it is possible to dispense with semi-anonymity using the techniques of Deb and Kalai (2010) and Azrieli and Shmaya (2011), and to dispense with the assumption that players' types are independent using the approach of Gradwohl and Reingold (2010). We leave these issues for further research.

2 Main Result

Let A and T be non-empty compact metric spaces, fixed for the rest of this section. Let $S = T \times A$ be endowed with some product metric, and let $M(S)$ be the space of

[2]Allowing for discontinuous payoff functions does not imply that the set of Bayes-Nash equilibria is empty, a case in which our result would be trivial. In fact, in the context of a Cournot oligopoly example, if all cost functions are lower semi-continuous (as they would be if discontinuities arise due to set-up costs of new machines at different production levels), then, if revenue functions are continuous in all players actions, each player's payoff function is upper semi-continuous in his action and continuous in the action of the other players. For this case, it follows from Dasgupta and Maskin (1986, Corollary) that a mixed Bayes-Nash equilibrium exists.

all Borel probability measures on S, endowed with the narrow topology.[3] Note that $M(S)$ is compact and metrizable. Let ρ be any metric on $M(S)$ which induces the narrow topology.

We write 1_s for the Dirac measure at any $s \in S$. If $\mu_1, \ldots, \mu_n$ are any elements of $M(S)$, then

$$\tilde{\mu} = \mu_1 \otimes \cdots \otimes \mu_n,$$

i.e., $\tilde{\mu}$ is the product probability measure on S^n defined from the measures $\mu_1, \ldots, \mu_n$, and, given $i \in I$ and $s \in S$,

$$\tilde{\mu}_s^i = \mu_1 \otimes \cdots \otimes \mu_{i-1} \otimes 1_s \otimes \mu_{i+1} \otimes \cdots \otimes \mu_n,$$

i.e., $\tilde{\mu}_s^i$ is the product measure on S^n defined from $\mu_1, \ldots, \mu_{i-1}, 1_s, \mu_{i+1}, \ldots, \mu_n$. By $\tilde{s}$ we denote a generic element of S^n, and by $\tilde{s}_k$ its kth component.

Let $\mathcal{C}$ be the set of all real-valued functions v such that

(a) $\mathrm{dom}(v)$ is a closed subset of $S \times M(S)$, denoting by $\mathrm{dom}(v)$ the domain of v;

(b) v is bounded and measurable for the subspace σ-algebra on $\mathrm{dom}(v)$ defined from the Borel σ-algebra of $S \times M(S)$.

We say that a subset K of $\mathcal{C}$ is *bounded* if $\sup\{||v||_\infty : v \in K\} < \infty$, where $||v||_\infty = \sup\{|v(x)| : x \in \mathrm{dom}(v)\}$. We say that a subset K of $\mathcal{C}$ is *uniformly equicontinuous in the $M(S)$-component* if given $\varepsilon > 0$ there is a $\delta > 0$ such that for all $v \in K$, $|v(s,\mu) - v(s',\mu')| < \varepsilon$ whenever $(s,\mu), (s',\mu') \in \mathrm{dom}(v)$ satisfy $s = s'$ and $\rho(\mu,\mu') < \delta$.

Note that the property of being a subset of $\mathcal{C}$ that is uniformly equicontinuous in the $M(S)$-component is topological, i.e., does not depend on the particular choice of the metric ρ. This is so because $M(S)$ with the narrow topology is compact and therefore all metrics on $M(S)$ that induce this topology are uniformly equivalent.[4]

[3] Recall that the narrow topology on $M(S)$ is the smallest topology on $M(S)$ for which all sets of the form $\{\mu \in M(S) \colon \mu(G) > \alpha\}$ are open, where G is an open subset of S, and α a real number.

[4] Recall that two metrics ρ_1, ρ_2 on a set X are uniformly equivalent if for every $\varepsilon > 0$ there are $\delta_1 > 0$ and $\delta_2 > 0$ such that for any $x, y \in X$, $\rho_2(x,y) < \varepsilon$ whenever $\rho_1(x,y) < \delta_1$ and $\rho_1(x,y) < \varepsilon$ whenever $\rho_2(x,y) < \delta_2$.

We consider *semi-anonymous Bayesian games* defined as follows. There is a finite set $I = \{1, \ldots, n\}$ of players. For each $i \in I$, player i's action space A_i is a closed subset of A, and player i's type space T_i is a closed subset of T. Let $S_i = A_i \times T_i$ for each $i \in I$.

For each $i \in I$, ν_i denotes the individual type distribution on T_i. Type realizations are independent among players. A mixed strategy (in distributional form) of player i is an element σ_i of $M(S)$ with $\sigma_i(S_i) = 1$ such that the marginal on T_i is ν_i.[5] Given a profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ of mixed strategies (a mixed strategy, for short), and some mixed strategy σ'_i of player $i \in I$, we write (σ'_i, σ_{-i}) for $(\sigma_1, \ldots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \ldots, \sigma_n)$.

For all $i \in I$, player i's payoff function is given by an element $v_i \in \mathcal{C}$ such that $\mathrm{dom}(v_i) = S_i \times H_i$, where

$$H_i = \left\{ \mu \in M(S) \colon \mu = \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j} \text{ where } s_j \in S_j \text{ for all } j \in I \setminus \{i\} \right\}.$$

The interpretation is that $v_i(t, a, \mu)$ is player i's payoff when he is of type t, plays action a and faces the distribution $\mu = \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}$ on S induced by the types and actions of all other players. Thus, given a mixed strategy $\sigma = (\sigma_1, \ldots, \sigma_n)$, the expected payoff of player i is

$$U_i(\sigma) = \int_{\prod_{j \in I} S_j} v_i \left(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j} \right) d\tilde{\sigma}(\tilde{s}) \,.$$

This concludes the definition of a semi-anonymous Bayesian game. Thus, a semi-anonymous Bayesian game G can be described as a list $G = (T_i, A_i, \nu_i, v_i)_{i \in I}$.

A mixed strategy $\sigma = (\sigma_1, \ldots, \sigma_n)$ of a semi-anonymous game G is a *Bayes-Nash equilibrium of* G if for each $i \in I$ and mixed strategy σ'_i of player i, $U_i(\sigma) \geq U_i(\sigma'_i, \sigma_{-i})$.

For all $s = (t, a) \in S$, let $t(s) = t$. Given a Bayesian semi-anonymous game G and a number $\varepsilon > 0$, an element $(s_1, \ldots, s_n) \in \prod_{i \in I} S_i$ is said to be ε*-Nash in* G if

$$v_i \left(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j} \right) \geq v_i \left(s, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j} \right) - \varepsilon$$

[5]This way of representing mixed or behavioral strategies was introduced into the literature by Milgrom and Weber (1985); see this paper for a discussion.

for all $i \in I$ and $s \in S_i$ with $t(s) = t(s_i)$. That is, $(s_1, \dots, s_n) = ((t_1, a_1), \dots, (t_n, a_n))$ is ε-Nash in G if $(a_1, \dots, a_n)$ is an ε-equilibrium of the complete information version $G(t_1, \dots, t_n) = (T_i, A_i, 1_{t_i}, v_i)_{i=1}^n$ of G where players' types are $t_1, \dots, t_n$.

Given a mixed strategy $\sigma = (\sigma_1, \dots, \sigma_n)$, for each player $i \in I$ and type $t \in T_i$ we let

$$m_i(t, \sigma) = \sup\left\{ \int_{\prod_{j \in I} S_j} v_i\left(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}_s^i(\tilde{s}) \colon s = (t, a), a \in A_i \right\}.$$

Note that if $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Bayes-Nash equilibrium, then for each $i \in I$ this supremum is attained for ν_i-almost all $t \in T_i$ (see Lemma 6 in the next section). Thus, in a Bayes-Nash equilibrium, for ν_i-almost all $t \in T_i$, $m_i(t, \sigma)$ is the maximal conditional expected payoff of player i given that his type is t and the other players play according to $(\sigma_1, \dots, \sigma_n)$.

For real numbers $\varepsilon > 0$ and $0 < \eta < 1$, a mixed strategy $\sigma = (\sigma_1, \dots, \sigma_n)$ of G is said to be (ε, η)-*ex post Nash* if

$$\tilde{\sigma}\Bigg(\bigg\{(s_1, \dots, s_n) \in \prod_{i \in I} S_i : (s_1, \dots, s_n) \text{ is } \varepsilon\text{-Nash in } G \text{ and}$$
$$\left| v_i\left(s_i, \frac{1}{1-n} \sum_{j \in I \setminus \{i\}} 1_{s_j}\right) - m_i\big(t(s_i), \sigma\big) \right| \leq \varepsilon \text{ for all } i \in I \bigg\}\Bigg) \geq 1 - \eta.$$

In this case, we also say informally that σ is approximately ex-post Nash. Thus, a mixed strategy is approximately ex-post Nash if there is a high probability that the realization of types and actions is such that the action profile is an approximate equilibrium of the complete information version of the game determined by the realized types. Furthermore, with the same probability, each player receives a payoff in the complete information version close to the one he receives in the incomplete information version.

Our main result states that in all sufficiently large semi-anonymous Bayesian games with payoff functions selected from a subset of $\mathcal{C}$ that is uniformly equicontinuous in the $M(S)$-component and bounded, all Bayes-Nash equilibria are approximately ex-post Nash.

Theorem. *Let K be a subset of $\mathcal{C}$ and suppose K is uniformly equicontinuous in the $M(S)$-component and bounded. Then, given $\varepsilon > 0$ and $0 < \eta < 1$, there is an $N \in \mathbb{N}$ such that the following holds for every $n \geq N$: If $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a Bayes-Nash equilibrium of a semi-anonymous Bayesian game G with n players whose payoff functions belong to K, then σ is (ε, η)-ex post Nash.*

Remark 1. Using the narrow topology on $M(S)$ for the notion of continuity in the theorem may be interpreted as follows. The narrow topology on $M(S)$ reflects the topology of the underlying space S of possible type-action characters in the sense of treating two distributions on S as close whenever they involve similar type-action characters with similar frequencies; e.g., if $\mu = \sum_{k=1}^{m} \alpha_k 1_{s_k}$ and $\mu' = \sum_{k=1}^{m} \alpha'_k 1_{s'_k}$ are such that the type-action characters s_k and s'_k are close in S for each k, and the frequencies α_k and α'_k with which they occur are close for each k, then the distributions μ and μ' are treated as close by the narrow topology. Thus the narrow topology on $M(S)$ in the continuity requirement in the theorem means that players, given their own type-action character, get similar payoffs whenever profiles of type-action characters of the respective rival players have distributions involving similar type-action characters with similar frequencies.

Remark 2. In Kalai (2004), a semi-anonymous Bayesian game $G = (T_i, A_i, \nu_i, v_i)_{i \in I}$ is defined as above except that, for each $i \in I$, v_i is given as a real-valued function defined on $\prod_{i \in I} S_i$, and semi-anonymity is brought in by stipulating that $v_i(\tilde{s}) = v_i(\tilde{s}')$ whenever $\tilde{s}_i = \tilde{s}'_i$ and $\frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j} = \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}'_j}$. Player i's payoff of a mixed strategy $\sigma = (\sigma_1, \ldots, \sigma_n)$ is $\int_{\prod_{j \in I} S_j} v_i(\tilde{s})\, d\tilde{\sigma}(\tilde{s})$.

Note that given a game specified in this way, there is, for each $i \in I$, a uniquely determined function $\hat{v}_i \colon S_i \times H_i \to \mathbb{R}$ such that $v_i(\tilde{s}) = \hat{v}_i(\tilde{s}_i, \frac{1}{n-1} \sum_{j \neq i} 1_{\tilde{s}_j})$ for all $\tilde{s} \in \prod_{j \in I} S_j$ (with H_i as defined above). In particular, given a mixed strategy σ,

$$\int_{\prod_{j \in I} S_j} v_i(\tilde{s})\, d\tilde{\sigma}(\tilde{s}) = \int_{\prod_{j \in I} S_j} \hat{v}_i\left(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j} \right) d\tilde{\sigma}(\tilde{s}),$$

i.e., computing expected payoff in terms of v_i or $\hat{v}_i$ amounts to the same. Thus, specification of a semi-anonymous game as in this note and specification of a semi-anonymous game as in Kalai (2004) amounts to the same.

3 Examples

In this section we present two examples to illustrate our main result. The examples show, in particular, that infinite dimensional type and action spaces may arise naturally and that the measurability and equicontinuity conditions in our result can be satisfied in the context of such type and action spaces.

3.1 Cournot Oligopoly

An advantage of allowing for general compact metric type and action spaces in our formalization is that one can easily incorporate rich payoff diversity. For instance, in a semi-anonymous game where the payoff of each single player depends only on his type-action character and the distribution of the actions of the other players, but not on the types of these players, this can be obtained as follows: Let $M(A)$ be the set of all Borel probability measures on the universal action space A, endowed with the narrow topology, and let the universal type space T be a non-empty set of bounded measurable functions on $A \times M(A)$ which is compact for the sup-norm. Then for each player i define the payoff function $v_i \colon T_i \times A_i \times H_i \to \mathbb{R}$ by setting $v_i(t, a, \mu) = t(a, \mu_A)$ where μ_A denotes the marginal distribution of μ on A.

A concrete economic example where an approach like this is useful is provided by the following Cournot oligopoly game (a similar example has been considered in Carmona (2008)). The particular feature of this example is that each firm is uncertain about the cost functions (and hence about the strategies) of the other firms and there is a large set of possible cost functions that firms can have. In particular, the possible cost functions are not restricted to be characterized by a finite number of parameter (say, a constant marginal cost and a fixed cost). Furthermore, these costs functions need not be continuous and thus may exhibit set-up costs.

Specifically, we formulate the Cournot oligopoly game as a Bayesian game $G = (T_i, A_i, \nu_i, v_i)_{i \in I}$ as follows. The set I of players (i.e., firms) is $I = \{1, \ldots, n\}$. Each firm produces a non-negative quantity of a homogeneous good. Due to capacity constraints, production cannot exceed a given level $m > 0$. Thus, for each $i \in I$, player i's action set is $A_i = [0, m]$. Let T be a set of bounded measurable functions on $[0, m]$, endowed with the (relativized) sup-norm topology, such that T is compact for this topology. The space T is the set of possible cost function of the firms; it may be infinite dimensional and its members may exhibit set-up costs. Types of firms are identified with cost function; thus, for each $i \in I$, player i's type space is $T_i = T$. Let the type distributions ν_i be any Borel probability measures on T. Note that even though all firms in the game G have the same type space, there may be asymmetries concerning costs as the type distributions ν_i may vary across firms.

Payoffs of firms are given as follows. If firms in the game G produce quantities $a_1, \dots, a_n$, then the market price of the good is $P(\sum_{i=1}^n a_i/n)$, where $P : [0, m] \to \mathbb{R}$ is continuous.[6] Thus, the revenue of firm i is $P(\sum_{i=1}^n a_i/n)a_i$. Therefore, firm i's payoff when it is of type t_i, plays action a_i and the other firms play $(a_j)_{j \in I \setminus \{i\}}$ is

$$P\left(\frac{a_i}{n} + \frac{\sum_{j \in I \setminus \{i\}} a_j}{n}\right) a_i - t_i(a_i).$$

Now to formulate payoff functions in the notation of Section 2, let $S = T \times [0, m]$ and recall from Section 2 that for each $i \in I$,

$$H_i = \left\{\mu \in M(S) \colon \mu = \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{(t_j, a_j)},\ (t_j, a_j) \in S_j \text{ for all } j \in I \setminus \{i\}\right\},$$

i.e., H_i is the set of the type-action distributions that can be generated by the players $j \in I \setminus \{i\}$. (Of course, the sets S_i, $i \in I$, are the same here, and this is so for the sets H_i as well.) Note that if firm $i \in I$ chooses action $a \in [0, m]$ and the distribution over types and actions generated by the other firms in I is $\mu = \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{(t_j, a_j)} \in H_i$, then the total output of all the firms in I can be written as $a + (n-1) \int_S \mathrm{proj}_{[0,m]}\, d\mu$, where $\mathrm{proj}_{[0,m]}$ denotes the projection of S onto $[0, m]$. Thus, in the notation of Section 2, the payoff function $v_i \colon T \times [0, m] \times H_i \to \mathbb{R}$ of firm i is given by setting

$$v_i(t, a, \mu) = P\left(\frac{1}{n} a + \left(1 - \frac{1}{n}\right) \int_S \mathrm{proj}_{[0,m]}\, d\mu\right) a - t(a)$$

for $(t, a, \mu) \in S \times H_i$.

[6] This formulation assumes that if the number n of firms is increased, then the set of buyers is appropriately replicated.

Now let

$$K = \Big\{ v \colon T \times [0,m] \times M(S) \to \mathbb{R} \colon \text{for some } n \in \mathbb{N} \backslash \{0\},$$
$$v(t,a,\mu) = P\Big(\frac{1}{n}a + \Big(1 - \frac{1}{n}\Big) \int_S \mathrm{proj}_{[0,m]} \, d\mu\Big) a - t(a)$$
$$\text{for all } (t,a,\mu) \in T \times [0,m] \times M(S) \Big\}$$

where, as in Section 2, $M(S)$ is endowed with the narrow topology. Note that in the game G defined above, the payoff function of any firm $i \in I$ is just the restriction of some element v of K to the set $T \times [0,m] \times H_i$, this v being determined by the number n of players. It is easily seen that all elements of K are measurable and that K is uniformly equicontinuous in the $M(S)$-component and bounded (see the Appendix for details). Thus our main result applies, showing that all Bayes-Nash equilibria of the Cournot oligopoly game G are approximately ex-post Nash provided that the number of players is sufficiently large.

Closing this example, we remark that there are several aspects in which the example is not covered by the framework of Deb and Kalai (2010). First, the example allows for an infinite dimensional type space. Second, the example allows for the realistic case of production with set-up costs and hence condition (LC1) of Deb and Kalai (2010) may fail. Finally, condition (LC2) of Deb and Kalai (2010) need not hold either. Indeed, that condition would require the inverse demand function P to be Lipschitz, where in our example this function may be any continuous function from $[0,m]$ to $\mathbb{R}$.

3.2 A Matching Game with Infinite Questions

The example presented in this section can be seen as an infinite-dimensional version of the Village versus Beach example in Kalai (2004). We construct a matching game between men and women where each player is confronted with a list of “yes” or “no” questions, and the action of a player is to give a list of answers. Each man wants to match the answers given by women to the questions that are important to him and, in contrast, each woman wants to mismatch the answers of men given to the questions that are important to her.

The case where there is only one question in the list corresponds to the Village versus Beach example in Kalai (2004). In our example, players may care about countably infinitely many questions, the importance given by each player to each question being determined by the player's type. Thus our example addresses situations where there is no finite bound on the numbers issues players may be interested in. For instance, the list of question could include: "Do you like tennis?", "Do you like opera?", "Would you like to go to the beach at 10am?" and so on. Or, more simply, the questions could just ask for each integer k: "Do you like the number k?"

To formally define a game $G = (T_i, A_i, \nu_i, v_i)_{i \in I}$ in this regard, for each player $i \in I$ let the action set A_i be a non-empty subset of $\{0,1\}^{\mathbb{N}}$. The interpretation is that an element $a = (a_0, a_1, \ldots) \in A_i$ is a vector of answers, a_k being the answer to question k, and if $a_k = 1$ (resp. $a_k = 0$) this means that the answer to question k is "yes" (resp. "no"). A natural special case is $A_i = \{0,1\}^{\mathbb{N}}$. However, the case where A_i is a proper subset of $\{0,1\}^{\mathbb{N}}$ is also of interest. For instance, let $\{B_j\}_{j=0}^{\infty}$ be a countable partition of $\mathbb{N}$ such that B_j is finite and non-empty for all $j \in \mathbb{N}$ and let A_i be the set of $a \in \{0,1\}^{\mathbb{N}}$ such that $|\{k \in B_j : a_k = 1\}| \leq 1$ for all $j \in \mathbb{N}$. Such a specification of A_i corresponds to a situation where there is a countably infinite number of issues (or activities), each of them having a finite number of (mutually exclusive) alternatives. More concretely, one of the issues, say issue 4, might concern the choice of going to the beach at 10am, going to the beach at 10:05am or staying in the village; letting $B_4 = \{8, 9\}$, any $a \in A_i$ satisfies $(a_8, a_9) \in \{(1,0), (0,1), (0,0)\}$ under the above specification of A_i, with the interpretation that $(a_8, a_9) = (1, 0)$ means going to the beach at 10am, $(a_8, a_9) = (0, 1)$ means going to the beach at 10:05am and $(a_8, a_9) = (0, 0)$ means staying in the village.

The type of a player identifies his/her gender and how important each question is for him/her. Thus the type of a player can be denoted as a pair (θ, ω) where $\theta \in \{-1, 1\}$ with the interpretation that -1 stands for male and 1 for female, and where $\omega = (\omega_k)_{k \in \mathbb{N}}$ is a vector of weights, with $0 \leq \omega_k \leq 1$ for each $k \in \mathbb{N}$, the component ω_k being the weight given to question k, expressing how important this question is for the player. In order for the notion of payoff stated in the next paragraph to be well-defined, we assume that any such vector ω of weights has the property that the sum $\sum_{k=0}^{\infty} \omega_k$ converges.

Suppose that the number n of players is even and that there are $n/2$ males and $n/2$ females. We then define the payoff for player $i \in I$ when he/she is of type $(\theta_i, \omega_i) = (\theta_i, (\omega_{i,k})_{k \in \mathbb{N}})$, plays $a_i = (a_{i,k})_{k=0}^{\infty}$ and the other players play $(a_j)_{j \in I \setminus \{i\}}$ by

$$\theta_i \sum_{k=0}^{\infty} \omega_{i,k} \left| a_{i,k} - \frac{|\{j \in I \setminus \{i\} : \theta_j = -\theta_i \text{ and } a_{j,k} = 1\}|}{n/2} \right|.$$

Thus for any player i, each question k matters for his/her payoff (a) by the importance of this question for him/her as measured by $\omega_{i,k}$, (b) by how close his/her answer to this question is to the fraction of players of the opposite gender who have answered "yes" (i.e., those players $j \neq i$ with $\theta_j = -\theta_i$ and $a_{j,k} = 1$), and (c) by whether he or she wants to match or to mismatch that fraction (i.e., by the player's gender θ_i).

We will assume in this example that there is a common type space T for all players (i.e., $T_i = T$ for all $i \in I$). Thus let $T = \Theta \times \Omega$, where $\Theta = \{-1, 1\}$ and Ω is a set of weight vectors. To make the example fit to our main result, we have to make sure that the equicontinuity hypothesis on payoff functions in this result can be satisfied when the number of players varies. For this purpose, we have to assume not just that for each $\omega \in \Omega$ the sum $\sum_{k=0}^{\infty} \omega_k$ converges, but actually that this convergence is uniform over the possible types of players. In view of this, for some fixed $\gamma \in [0,1]^{\mathbb{N}}$ with $\sum_{k=0}^{\infty} \gamma_k < \infty$ we let $\Omega = \{\omega \in [0,1]^{\mathbb{N}} : \omega_k \leq \gamma_k \text{ for all } k \in \mathbb{N}\}$.

Towards an application of our main result, we also need to make the common type space T of the players to be a compact metric space, and we need to make the action spaces A_i to be closed subsets of a common compact metric space, so that, in particular, the measurability and equicontinuity hypotheses made in this result on the payoff functions are satisfied. To meet this requirement, we proceed as follows. Let $\{0,1\}^{\mathbb{N}}$ be endowed with a metric that induces the product topology defined from the discrete topology on $\{0,1\}$, which makes $\{0,1\}^{\mathbb{N}}$ a compact metric space, and assume that A_i is closed in $\{0,1\}^{\mathbb{N}}$ for all $i \in I$.[7] Further, letting $[0, \gamma_k]$ for all $k \in \mathbb{N}$

[7]For instance, the set A_i defined above relative to the partition $\{B_j\}_{j=0}^{\infty}$ of $\mathbb{N}$ satisfies this closedness assumption.

be endowed with the subspace topology defined from the usual topology of $\mathbb{R}$, and Θ with the discrete topology, we endow $\Omega = \prod_{k=0}^{\infty}[0, \gamma_k]$ and then $T = \Theta \times \Omega$ with the product topology and a corresponding metric, so that Ω and T become compact metric spaces.

Next, concerning type distribution, let ν_i, $i \in I$, be any Borel probability measures on T such that the marginal $\nu_{i,\Theta}$ of ν_i on $\Theta = \{-1, 1\}$ is degenerate for each $i \in I$, and such that $|\{i \in I : \nu_{i,\Theta} = 1_1\}| = |\{i \in I : \nu_{i,\Theta} = 1_{-1}\}| = n/2$. Note that this ensures that the gender of a player is non-random and that the numbers of females and males in the game are equal.

Finally, we formulate payoff functions in the notation of Section 2. To this end, let $S = T \times \{0, 1\}^{\mathbb{N}}$, and for any $t = (\theta, \omega) \in T$, let $\theta(t) = \theta$ and $\omega(t) = \omega$. Further, for any $k \in \mathbb{N}$ and $\theta \in \Theta$, let $B_k^{\theta} = \{(t, a) \in S : \theta(t) = \theta \text{ and } a_k = 1\}$. Then, in the notation of Section 2, for each player i the payoff as defined above can be described by the payoff function $v_i : T_i \times A_i \times H_i \to \mathbb{R}$ given by setting

$$v_i(t, a, \mu) = \theta(t) \sum_{k=0}^{\infty} \omega(t)_k \left| a_k - \frac{2(n-1)\mu(B_k^{-\theta(t)})}{n} \right|$$

for $(t, a, \mu) \in T_i \times A_i \times H_i$. This completes the specification of the game G.

Now let

$$K = \Big\{ v \colon T \times \{0,1\}^{\mathbb{N}} \times M(S) \to \mathbb{R} \colon \text{for some } n \in \mathbb{N} \text{ with } n \geq 2, \\ v(t, a, \mu) = \theta(t) \sum_{k=0}^{\infty} \omega(t)_k \left| a_k - \frac{2(n-1)\mu(B_k^{-\theta(t)})}{n} \right| \\ \text{for all } (t, a, \mu) \in T \times \{0,1\}^{\mathbb{N}} \times M(S) \Big\}$$

and consider $M(S)$ as endowed with the narrow topology. Note that in the game G defined above, the payoff function of any player i is just the restriction of some element v of K to the set $T \times A_i \times H_i$, this v being determined by the number n of players. It is easily seen that all elements of K are measurable and that K is uniformly equicontinuous in the $M(S)$-component and bounded (see the Appendix for details). Thus our main result implies that all Bayes-Nash equilibria of the matching game G with infinite questions are approximately ex-post Nash provided that the number of players is sufficiently large.

4 Proof of the Theorem

The proof of our main result is based on several lemmas. Lemma 1 is a consequence of 492X(d) in Fremlin (2003) on combinatorial concentration of measures. Recall that if $X_1, \ldots, X_n$ are finitely many sets, the normalized Hamming metric d_H on the product $\prod_{i=1}^n X_i$ is given by $d_H(x,y) = \frac{1}{n}\bigl|\{i : 1 \leq i \leq n \text{ and } x_i \neq y_i\}\bigr|$ for all $x, y \in \prod_{i=1}^n X_i$.

Lemma 1. *Let* $\langle (X_i, \Sigma_i, \mu_i) \rangle_{i=1}^n$ *be a non-empty finite family of probability spaces and* (X, Λ, λ) *the corresponding product probability space. Suppose that* $f \colon X \to \mathbb{R}$ *is a bounded* λ*-integrable function which is* 2*-Lipschitz for the normalized Hamming metric on* X*. Then* $\lambda(\{x \in X \colon \bigl|f(x) - \int f\, d\lambda\bigr| \geq \gamma\}) \leq 2e^{-n\gamma^2/16}$ *for every* $\gamma \geq 0$*.*

Proof. Set $g = \frac{1}{2} f$. Then g is 1-Lipschitz for the normalized Hamming metric on X, and therefore, according to Fremlin (2003, 492X(d)), we must have

$$\lambda\biggl(\biggl\{x \in X \colon g(x) - \int g\, d\lambda \geq \gamma\biggr\}\biggr) \leq e^{-n\gamma^2/4}$$

for every $\gamma \geq 0$. It follows that

$$\lambda\biggl(\biggl\{x \in X \colon f(x) - \int f\, d\lambda \geq \gamma\biggr\}\biggr) \leq e^{-n\gamma^2/16} \tag{1}$$

for every $\gamma \geq 0$. Since the function $-f$ is also 2-Lipschitz for the normalized Hamming metric on X, we must also have

$$\lambda\biggl(\biggl\{x \in X \colon -f(x) - \int -f\, d\lambda \geq \gamma\biggr\}\biggr) \leq e^{-n\gamma^2/16}$$

for every $\gamma \geq 0$, or, in other words,

$$\lambda\biggl(\biggl\{x \in X \colon -\biggl(f(x) - \int f\, d\lambda\biggr) \geq \gamma\biggr\}\biggr) \leq e^{-n\gamma^2/16} \tag{2}$$

for every $\gamma \geq 0$. The conclusion of the lemma follows by combining (1) and (2). □

As a consequence of Lemma 1, the next lemma shows that if μ_i, $i = 1, \ldots, n$, are any Borel probability measures on S, then with high $\mu_1 \otimes \cdots \otimes \mu_n$-probability, the distribution $\frac{1}{n}\sum_{i=1}^n 1_{\tilde{s}_i}$ on S is close to the average distribution $\frac{1}{n}\sum_{i=1}^n \mu_i$ whenever n is sufficiently large.

Lemma 2. *For all $\delta > 0$ and $\eta > 0$, there is $N \in \mathbb{N}$ such that the following holds for every $n \geq N$: If $\mu_1, \ldots, \mu_n$ are any Borel probability measures on S, then $\tilde{\mu}(\{\tilde{s} \in S^n \colon \rho(\frac{1}{n}\sum_{i=1}^n 1_{\tilde{s}_i}, \frac{1}{n}\sum_{i=1}^n \mu_i) > \delta\}) \leq \eta$.*

Proof. Let $C(S)$ be the set of all real-valued continuous functions on S, and let

$$L = \{h \in C(S) \colon \|h\|_\infty \leq 1 \text{ and } h \text{ is 1-Lipschitz}\}.$$

Let ρ_H denote Huntingdon's metric on $M(S)$. Recall that ρ_H is defined by setting $\rho_H(\mu, \mu') = \sup\{|\int h\, d\mu - \int h\, d\mu'| \colon h \in L\}$ for all $\mu, \mu' \in M(S)$, and recall that ρ_H induces the narrow topology of $M(S)$ (see Fremlin (2003, 437L and 437Y(i))). As noted in Section 2, all metrics on $M(S)$ that induce the narrow topology are uniformly equivalent, so it suffices to show that the lemma holds with ρ_H as a particular choice of the metric ρ.

Fix any δ, $\eta > 0$ and set $\gamma = \min\{\delta/3, \eta\}$. Note that the subset L of $C(S)$ is bounded, closed and equicontinuous, hence compact. We may therefore select finitely many elements $h_1, \ldots, h_k$ of L so that for any $h \in L$ we have $\|h - h_j\|_\infty \leq \gamma$ for some $h_j \in \{h_1, \ldots, h_k\}$. It is straightforward to check that

$$\text{(3)} \qquad \rho_H(\mu, \mu') \leq \delta \text{ for all } \mu, \mu' \in M(S) \text{ satisfying} \quad \left|\int h_j\, d\mu - \int h_j\, d\mu'\right| \leq \gamma \text{ for each } j = 1, \ldots, k.$$

Choose $N \in \mathbb{N}$ such that $2e^{-N\gamma^2/16} < \gamma/k$. Fix any $n \in \mathbb{N}$ with $n \geq N$, and any Borel probability measures $\mu_1, \ldots, \mu_n$ on S. Recall that $\tilde{\mu}$ denotes the corresponding product measure on S^n. For each $j = 1, \ldots, k$, define a function $\tilde{h}_j \colon S^n \to \mathbb{R}$ by setting

$$\tilde{h}_j(\tilde{s}) = \frac{1}{n}\sum_{i=1}^n h_j(\tilde{s}_i) \quad \text{for all } \tilde{s} = (\tilde{s}_1, \ldots, \tilde{s}_n) \in S^n.$$

Since $\|h_j\|_\infty \leq 1$ for each $j = 1, \ldots, k$, each of the functions $\tilde{h}_j$ is 2-Lipschitz for the normalized Hamming metric on S^n. It is also clear that each $\tilde{h}_j$ is bounded and $\tilde{\mu}$-integrable. Using Lemma 1 it follows that

$$\tilde{\mu}\left(\left\{\tilde{s} \in S^n \colon \left|\int \tilde{h}_j\, d\tilde{\mu} - \tilde{h}_j(\tilde{s})\right| \geq \gamma\right\}\right) \leq 2e^{-n\gamma^2/16} \leq 2e^{-N\gamma^2/16} < \gamma/k$$

for each $j = 1, \ldots, k$, and hence that

$$\tilde{\mu}\left(\left\{\tilde{s} \in S^n \colon \left|\int \tilde{h}_j \, d\tilde{\mu} - \tilde{h}_j(\tilde{s})\right| < \gamma \text{ for each } j = 1, \ldots, k\right\}\right) > 1 - \gamma. \tag{4}$$

Note that for each $\tilde{s} = (\tilde{s}_1, \ldots, \tilde{s}_n) \in S^n$, and each $j = 1, \ldots, k$,

$$\tilde{h}_j(\tilde{s}) = \frac{1}{n} \sum_{i=1}^{n} h_j(\tilde{s}_i) = \int_S h_j \, d\Big(\frac{1}{n} \sum_{i=1}^{n} 1_{\tilde{s}_i}\Big). \tag{5}$$

Also, for each $j = 1, \ldots, k$, using Fubini's theorem,

$$\begin{aligned}
\int_{S^n} \tilde{h}_j(\tilde{s}_1, \ldots, \tilde{s}_n) \, d\tilde{\mu}(\tilde{s}_1, \ldots, \tilde{s}_n) &= \frac{1}{n} \int_{S^n} \sum_{i=1}^{n} h_j(\tilde{s}_i) \, d\tilde{\mu}(\tilde{s}_1, \ldots, \tilde{s}_n) \\
&= \frac{1}{n} \int_S \Big(\ldots \Big(\int_S \Big(\int_S \sum_{i=1}^{n} h_j(\tilde{s}_i) \, d\mu_1(\tilde{s}_1) \Big) \, d\mu_2(\tilde{s}_2) \Big) \ldots \Big) \, d\mu_n(\tilde{s}_n) \\
&= \frac{1}{n} \sum_{i=1}^{n} \int_S h_j(s) \, d\mu_i(s) \\
&= \int_S h_j \, d\Big(\frac{1}{n} \sum_{i=1}^{n} \mu_i\Big).
\end{aligned}$$

Hence, from (4) and (5),

$$\tilde{\mu}\left(\left\{\tilde{s} \in S^n \colon \left|\int_S h_j \, d\Big(\frac{1}{n} \sum_{i=1}^{n} \mu_i\Big) - \int_S h_j \, d\Big(\frac{1}{n} \sum_{i=1}^{n} 1_{\tilde{s}_i}\Big)\right| < \gamma, \forall j = 1, \ldots, k\right\}\right) > 1 - \gamma.$$

By (3), it follows that

$$\tilde{\mu}\left(\left\{\tilde{s} \in S^n \colon \rho_H\Big(\frac{1}{n} \sum_{i=1}^{n} \mu_i, \frac{1}{n} \sum_{i=1}^{n} 1_{\tilde{s}_i}\Big) \le \delta\right\}\right) > 1 - \gamma.$$

Since $\gamma \le \eta$, this implies that

$$\tilde{\mu}\left(\left\{\tilde{s} \in S^n \colon \rho_H\Big(\frac{1}{n} \sum_{i=1}^{n} \mu_i, \frac{1}{n} \sum_{i=1}^{n} 1_{\tilde{s}_i}\Big) \le \delta\right\}\right) > 1 - \eta.$$

This completes the proof. □

The next lemma shows that the conclusion of Lemma 2 remains valid when only $n - 1$ of the measures $\mu_1, \ldots, \mu_n$ are averaged out. This is a consequence of the fact that when n is large, removing one element from the average makes little difference.

Lemma 3. *For all $\delta > 0$ and $\eta > 0$, there is $N \in \mathbb{N}$ such that the following holds for every $n \geq N$: If $\mu_1, \ldots, \mu_n$ are Borel probability measures on S, then there is a measurable subset E of S^n, with $\tilde{\mu}(E) \geq 1 - \eta$, such that for each $\tilde{s} \in E$ and each $i \in I = \{1, \ldots, n\}$, $\rho\big(\frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \mu_j\big) \leq \delta$.*

Proof. Fix $\delta > 0$. We can find $N_1 \in \mathbb{N}$ such that for every $n \geq N_1$ the following is true: Whenever $\gamma_1, \ldots, \gamma_n$ are Borel probability measures on S, then for each $i \in I = \{1, \ldots, n\}$,

$$\rho\bigg(\frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \gamma_j, \frac{1}{n}\sum_{j\in I} \gamma_j\bigg) \leq \delta/3.$$

To see this, recall that the Prohorov metric ρ_P on $M(S)$ is defined by setting, for any $\gamma, \gamma' \in M(S)$,

$$\begin{aligned}\rho_P(\gamma, \gamma') = \inf\{\varepsilon > 0 \colon \gamma(E) \leq \gamma'(B_\varepsilon(E)) + \varepsilon \text{ and} \\ \gamma'(E) \leq \gamma(B_\varepsilon(E)) + \varepsilon \text{ for all Borel sets } E \subseteq S\},\end{aligned}$$

where $B_\varepsilon(E)$ is the open ε-ball around E for the metric of S, and recall that ρ_P generates the narrow topology. Note that for any $\gamma_1, \ldots, \gamma_n$ as above, with $n \geq 2$, and any $i \in I = \{1, \ldots, n\}$, we have

$$\bigg|\frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \gamma_j(E) - \frac{1}{n}\sum_{j\in I} \gamma_j(E)\bigg| \leq \frac{1}{n}$$

for all Borel sets $E \subseteq S$, and therefore $\rho_P\big(\frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \gamma_j, \frac{1}{n}\sum_{j\in I} \gamma_j\big) \leq 1/n$. Again using the fact that all metrics on $M(S)$ that generate the narrow topology are uniformly equivalent, the assertion follows.

Now by Lemma 2, we can find $N_2 \in \mathbb{N}$ such that for every $n \geq N_2$ the following holds: Whenever $\mu_1, \ldots, \mu_n$ are Borel probability measures on S, then there is a measurable subset E of S^n with $\tilde{\mu}(E) \geq 1 - \eta$ such that for each $\tilde{s} \in E$,

$$\rho\bigg(\frac{1}{n}\sum_{j\in I} 1_{\tilde{s}_j}, \frac{1}{n}\sum_{j\in I} \mu_j\bigg) \leq \delta/3.$$

Setting $N = \max\{N_1, N_2\}$ gives us the desired N. □

The previous lemma shows that if n is large and $\mu_1, \ldots, \mu_n$ are any Borel probability measures on S, then, for all $i \in \{1, \ldots, n\}$, the distribution $\frac{1}{n-1}\sum_{j\neq i} 1_{\tilde{s}_j}$ on S

is close to the average distribution $\frac{1}{n-1}\sum_{j\neq i}\mu_i$ with high $\mu_1\otimes\cdots\otimes\mu_n$-probability. This fact will be used below (see Lemma 5) to show that, for all types and actions of any player i, his expected payoff against the mixed strategies $(\mu_j)_{j\neq i}$ of the other players is close to the payoff against the average distribution $\frac{1}{n-1}\sum_{j\neq i}\mu_j$, provided this average distribution belongs to the domain of his payoff function v_i. However, in a semi-anonymous Bayesian game $G=(T_i,A_i,\nu_i,v_i)_{i\in I}$, for each $i\in I$, player i's payoff function is defined on $S_i\times H_i$ and there is no guarantee that the average distribution $\frac{1}{n-1}\sum_{j\neq i}\mu_j$ belongs to H_i. We therefore extend players payoff functions from $S_i\times H_i$ to $S_i\times M(S)$.

Lemma 4. *Let $K\subseteq\mathcal{C}$ be uniformly equicontinuous in the $M(S)$-component and bounded. Suppose that for each $v\in K$, $\mathrm{dom}(v)=S_v\times H_v$, where S_v and H_v are closed subsets of S and $M(S)$ respectively. Then each $v\in K$ has an extension to an element $v'\in\mathcal{C}$ with $\mathrm{dom}(v')=S_v\times M(S)$ such that the set $K'\subseteq\mathcal{C}$ of all these extensions is again uniformly equicontinuous in the $M(S)$-component and bounded.*

Proof. Indexing the elements of K, we can write $\{v_i : i\in I\}$ for K (here, I is an abstract index set) and then, for each $i\in I$, $S_i\times H_i$ for $\mathrm{dom}(v_i)$. We may assume that each v_i takes all of its values in $[1,2]$. Following Mandelkern (1990), for each $i\in I$ define $v_i'\colon S_i\times M(S)\to[1,2]$ by setting $v_i'(s,\mu)=v_i(s,\mu)$ if $\mu\in H_i$, and

$$v_i'(s,\mu)=\inf_{u\in H_i}\frac{v_i(s,u)\rho(\mu,u)}{\rho(\mu,H_i)}$$

if $\mu\notin H_i$. Then each v_i' takes all of its values in $[1,2]$, and because $\{v_i\colon i\in I\}$ is uniformly equicontinuous in the $M(S)$-component, the proof of the theorem in Mandelkern (1990) shows that $\{v_i'\colon i\in I\}$ must be uniformly equicontinuous in the $M(S)$-component. Thus it remains to show that for each $i\in I$, v_i' is measurable on its domain $S_i\times M(S)$.

Pick any $i\in I$. Since for each fixed $s\in S_i$, $v_i'(s,\cdot)$ is continuous on $M(S)$, by Aliprantis and Border (2006, Lemma 4.51, p. 153) it suffices to show that for each fixed $\mu\in M(S)$, $v_i'(\cdot,\mu)$ is measurable on S_i. If $\mu\in H_i$, this holds since in this case $v_i'(\cdot,\mu)$ coincides with $v_i(\cdot,\mu)$. Thus consider $\mu\notin H_i$. Let D be a countable dense subset of H_i. Since the function $u \mapsto v_i(s,u)\rho(\mu,u)/\rho(\mu,H_i)\colon H_i \to \mathbb{R}$ is continuous for each $s \in S_i$, we have

$$v_i'(s,\mu) = \inf_{u\in H_i} \frac{v_i(s,u)\rho(\mu,u)}{\rho(\mu,H_i)} = \inf_{u\in D} \frac{v_i(s,u)\rho(\mu,u)}{\rho(\mu,H_i)}$$

for each $s \in S_i$. Since $s \mapsto v_i(s,u)$ is measurable on S_i for each $u \in D$, it follows that $s \mapsto v_i'(s,\mu)$ is measurable as well. □

The next lemma implies, in particular, that the following holds if n is sufficiently large: Whenever $(\sigma_1,\ldots,\sigma_n)$ is a mixed strategy in a semi-anonymous Bayesian game, then, for all players $i \in I$, types $t \in T_i$ and actions $a \in A_i$, player i's expected payoff when of type t and playing a against the mixed strategies of the other players is, for his extended payoff function, close to the certain payoff against the average distribution $\frac{1}{n-1}\sum_{j\neq i}\sigma_j$. The following notation is used in the sequel: $\mathrm{dom}_S(v)$ denotes the projection of $\mathrm{dom}(v)$ on S.

Lemma 5. *Let K be a subset of $\mathcal{C}$. Suppose that K is uniformly equicontinuous in the $M(S)$-component and bounded, and that for each $v \in K$, $\mathrm{dom}(v) = \mathrm{dom}_S(v) \times M(S)$. Then, for all $\varepsilon > 0$, there is an $N \in \mathbb{N}$ such that the following holds for every $n \geq N$: If $\sigma_1,\ldots,\sigma_n$ are Borel probability measures on S, then, for each $v \in K$, each $s \in \mathrm{dom}_S(v)$, and each $i \in I = \{1,\ldots,n\}$,*

$$\left| v\left(s, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \sigma_j\right) - \int_{S^n} v\left(\tilde{s}_i, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}_s^i(\tilde{s}) \right| < \varepsilon.$$

Proof. Note first that the integral $\int_{S^n} v\big(\tilde{s}_i, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}\big)\, d\tilde{\sigma}_s^i(\tilde{s})$ is indeed defined if $s \in \mathrm{dom}_S(v)$. To see this, observe that we have $\tilde{\sigma}_s^i(\{\tilde{s} \in S^n\colon \tilde{s}_i = s\}) = 1$ by definition of $\tilde{\sigma}_s^i$. Hence, if $s \in \mathrm{dom}_S(v)$, then $\tilde{\sigma}_s^i\big(\{\tilde{s} \in S^n\colon (\tilde{s}_i, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}) \in \mathrm{dom}(v)\}\big) = 1$, because $\mathrm{dom}(v) = \mathrm{dom}_S(v) \times M(S)$ by hypothesis.

Fix any $\varepsilon > 0$. By hypothesis, we can choose a numbers $\delta, \eta > 0$ such that $2\eta \sup\{\|v\|_\infty : v \in K\} < \varepsilon/2$ and such that for all $v \in K$, $|v(s,\mu) - v(s,\mu')| < \varepsilon/2$ whenever (s,μ), $(s,\mu') \in \mathrm{dom}(v)$ satisfy $\rho(\mu,\mu') < \delta$.

Relative to these δ and η, let N be chosen according to Lemma 3. Fix any $n \geq N$ and let $\sigma_1,\ldots,\sigma_n$ be any Borel probability measures on S. Pick any $v \in K$, any $s \in \mathrm{dom}_S(v)$, and any $i \in I = \{1, \dots, n\}$. Lemma 3 applied to the measures $\sigma_1, \dots, \sigma_{i-1}, 1_s, \sigma_{i+1}, \dots, \sigma_n$ gives a measurable set $E \subseteq S^n$ with $\tilde{\sigma}_s^i(E) \geq 1 - \eta$ such that $\rho\big(\frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \sigma_j\big) \leq \delta$ for each $\tilde{s} \in E$. Hence, by choice of δ, since $\tilde{\sigma}_s^i(\{\tilde{s} \in S^n \colon \tilde{s}_i = s\}) = 1$, it follows that

$$\tilde{\sigma}_s^i\left(\left\{\tilde{s} \in S^n \colon \left|v\left(s, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \sigma_j\right) - v\left(\tilde{s}_i, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}\right)\right| \leq \varepsilon/2\right\}\right) \geq 1 - \eta.$$

Since $2\eta\|v\|_\infty < \varepsilon/2$, it follows from this that

$$\left|v\left(s, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} \sigma_j\right) - \int_{S^n} v\left(\tilde{s}_i, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}_s^i(\tilde{s})\right| < \varepsilon.$$

As $i \in I$, $v \in K$, and $s \in \mathrm{dom}_S(v)$ were arbitrary, the lemma is proved. □

The next lemma shows that if $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Bayes-Nash equilibrium, then for each $i \in I$, the maximal conditional expected payoff of player i given that his type is t and the other players play according to $(\sigma_1, \dots, \sigma_n)$ is attained for ν_i-almost all $t \in T_i$.

Lemma 6. *Let $G = (T_i, A_i, \nu_i, v_i)_{i\in I}$ be a semi-anonymous Bayesian game and suppose that $\sigma = (\sigma_1, \dots, \sigma_n)$ is a Bayes-Nash equilibrium of G. Then, for each $i \in I$, there is a Borel set $F_i \subseteq T_i \times A_i$, with $\sigma_i(F_i) = 1$, such that if $s = (t, a) \in F_i$, then $\int_{\prod_{j\in I} S_j} v_i\left(\tilde{s}_i, \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}_s^i(\tilde{s}) = m_i(t, \sigma)$.*

Proof. Consider any $i \in I$. Let $\mathcal{A}$ denote the Borel σ-algebra of A_i, $\mathcal{B}$ that of T_i, $\bar{\nu}_i$ the completion of ν_i, and $\mathcal{B}_{\nu_i}$ the domain of $\bar{\nu}_i$. Furthermore, let $\tilde{\sigma}_{-i}$ denote the product measure $\sigma_1 \otimes \cdots \otimes \sigma_{i-1} \otimes \sigma_{i+1} \otimes \cdots \otimes \sigma_n$ on $\prod_{j\in I\setminus i} S_j$, let $\tilde{s}_{-i}$ denote a generic element of $\prod_{j\in I\setminus i} S_j$, and let $g \colon A_i \times T_i \to \mathbb{R}$ be the function defined by setting

$$g(t, a) = \int_{\prod_{j\in I\setminus\{i\}} S_j} v_i\left((t, a), \frac{1}{n-1}\sum_{j\in I\setminus\{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}_{-i}(\tilde{s}_{-i}) \text{ for all } (t, a) \in T_i \times A_i.$$

Since v_i is bounded, using Fubini's theorem it follows that g is $\mathcal{B} \otimes \mathcal{A}$-measurable. Moreover, for any mixed strategy σ_i' of player i, the expected payoff of player i when the other players play σ_{-i} is $\int_{S_i} g \, d\sigma_i'$.

Note that $\sup\{g(t, a) \colon a \in A_i\} = m_i(t, \sigma)$ for each $t \in T_i$. Hence, since g is $\mathcal{B} \otimes \mathcal{A}$-measurable it follows from Castaing and Valadier (1977, Lemma III.39 and remarks in the sequel) that $t \mapsto m_i(t,\sigma)$ is $\mathcal{B}_{\nu_i}$-measurable, and that for any $\varepsilon > 0$ there is a $(\mathcal{B}_{\nu_i}, \mathcal{A}_i)$-measurable $h^\varepsilon\colon T_i \to A_i$ such that $\int g(t, h^\varepsilon(t))\, d\bar{\nu}_i(t) \geq \int m_i(t,\sigma)\, d\bar{\nu}_i(t) - \varepsilon$. Given such a function h^ε, let σ_i^ε be the Borel probability measure on $T_i \times A_i$ defined by setting $\sigma_i^\varepsilon(B) = \bar{\nu}_i(\{t \in T_i\colon (t, h^\varepsilon(t)) \in B\})$ for each $B \in \mathcal{B} \otimes \mathcal{A}$. Then the marginal of σ_i^ε on T_i is ν_i, and thus σ_i^ε is a mixed strategy of player i. Also, $\int g\, d\sigma_i^\varepsilon = \int g(t, h^\varepsilon(t))\, d\bar{\nu}_i(t)$ and hence $\int g\, d\sigma_i^\varepsilon \geq \int m_i(t,\sigma)\, d\bar{\nu}_i(t) - \varepsilon$. Since σ is a Bayes-Nash equilibrium, $\int g\, d\sigma_i \geq \int g\, d\sigma_i^\varepsilon$. It follows that $\int g\, d\sigma_i \geq \int m_i(t,\sigma)\, d\bar{\nu}_i(t)$.

The property of $m_i(\cdot,\sigma)$ being $\mathcal{B}_{\nu_i}$-measurable means that there is $B \in \mathcal{B}$ with $\nu_i(B) = 0$ such that the restriction of $m_i(\cdot,\sigma)$ to $T_i \backslash B$ is Borel measurable. Choose such B and let

$$F_i = \{(t,a) \in (T_i \backslash B) \times A_i\colon g(t,a) = m_i(t,\sigma)\}.$$

Then F_i is a Borel set in $T_i \times A_i$. By the facts that the marginal of σ_i on T_i is ν_i, which in particular implies that $\sigma_i(B \times A_i) = 0$, and that $g(t,a) \leq m_i(t,\sigma)$ for all $(t,a) \in T_i \times A_i$, we see that were $\sigma_i(F_i) < 1$, we would have

$$\int g\, d\sigma_i < \int m_i(t,\sigma)\, d\bar{\nu}_i(t)\,,$$

contradicting what has been established in the previous paragraph. Thus, $\sigma_i(F_i) = 1$. To complete the proof, it only remains to note that for each $s = (t,a) \in T_i \times A_i$,

$$g(t,a) = \int_{\prod_{j\in I} S_j} v_i\left(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I\backslash\{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}_s^i(\tilde{s})\,. \qquad \square$$

We finally turn to the proof of our main result.

Proof of the Theorem. Fix $\varepsilon > 0$ and $0 < \eta < 1$. Let $K' \subseteq \mathcal{C}$ be chosen for K according to Lemma 4. Using Lemma 5 and Lemma 3, we can select $N \in \mathbb{N}$ such that the following hold for all $n \geq N$ and any given Borel probability measures $\mu_1, \ldots, \mu_n$ on S, setting $I = \{1, \ldots, n\}$:

(I) For each $v \in K'$, each $s \in \mathrm{dom}_S(v)$, and each $i \in I$,

$$\left| v\left(s, \frac{1}{n-1} \sum_{j \in I\backslash\{i\}} \mu_j\right) - \int_{S^n} v\left(\tilde{s}_i, \frac{1}{n-1} \sum_{j\in I\backslash\{i\}} 1_{\tilde{s}_j}\right) d\tilde{\mu}_s^i(\tilde{s}) \right| < \varepsilon/4$$

(the integral being defined, see the beginning of the proof of Lemma 5).

(II) There is a measurable $E \subseteq S^n$, with $\tilde{\mu}(E) > 1-\eta$, such that if $(s_1, \ldots, s_n) \in E$, then for each $v \in K'$, $s \in \mathrm{dom}_S(v)$, and $i \in I$,

$$\left| v\left(s, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} \mu_j\right) - v\left(s, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\right) \right| < \varepsilon/4.$$

Fix $n \geq N$, and suppose $(\sigma_1, \ldots, \sigma_n)$ is a Bayes-Nash equilibrium of a game G with a finite set $I = \{1, \ldots, n\}$ of players whose payoff functions v_i belong to K. By Lemma 6, for each $i \in I$ we can choose a measurable $F_i \subseteq S_i$ with $\sigma_i(F_i) = 1$ such that:

(III) If $s = (t, a) \in F_i$, then $\int_{\prod_{j \in I} S_j} v_i\big(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j}\big)\, d\tilde{\sigma}^i_s(\tilde{s}) = m_i(t, \sigma)$.

Choose $E \subseteq S^n$ relative to $(\sigma_1, \ldots, \sigma_n)$ according to (II) above, and then let

$$F = E \cap (F_1 \times \cdots \times F_n).$$

Note that $\tilde{\sigma}(F) = \tilde{\sigma}(E)$. Thus, $\tilde{\sigma}(F) > 1 - \eta$.

Consider any $(s_1, \ldots, s_n) \in F$ and any player $i \in I$. Let $t_i = t(s_i)$. Identify v_i with its extension to an element of K'. Then, by choice of E and by (I),

$$\begin{aligned} v_i\left(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\right) &\geq v_i\left(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} \sigma_j\right) - \varepsilon/4 \\ &\geq \int_{\prod_{j \in I} S_j} v_i\left(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}^i_{s_i}(\tilde{s}) - \varepsilon/2. \end{aligned}$$

Since $s_i \in F_i$, it follows from this and (III) that

$$v_i\left(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\right) \geq m_i(t_i, \sigma) - \varepsilon/2.$$

On the other hand, for any $s \in S_i$ such that $t(s) = t_i$, using the choice of E, (I) and the definition of $m_i(t_i, \sigma)$ in this order, we obtain that

$$\begin{aligned} v_i\left(s, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\right) &\leq v_i\left(s, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} \sigma_j\right) + \varepsilon/4 \\ &\leq \int_{\prod_{j \in I} S_j} v_i\left(\tilde{s}_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{\tilde{s}_j}\right) d\tilde{\sigma}^i_s(\tilde{s}) + \varepsilon/2 \leq m_i(t_i, \sigma) + \varepsilon/2. \end{aligned}$$

In particular, we must have

$$v_i\Bigg(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\Bigg) \leq m_i(t_i, \sigma) + \varepsilon/2.$$

Summing up, it follows that $\big|v_i\big(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\big) - m_i(t_i, \sigma)\big| \leq \varepsilon/2$, and that for any $s \in S_i$ with $t(s) = t_i$,

$$\begin{aligned} v_i\Bigg(s, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\Bigg) &\leq m_i(t_i, \sigma) + \varepsilon/2 \\ &\leq v_i\Bigg(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\Bigg) + \varepsilon. \end{aligned}$$

Thus, as $i \in I$ was arbitrary, $(s_1, \ldots, s_n)$ is ε-Nash in G and $\big|v_i\big(s_i, \frac{1}{n-1} \sum_{j \in I \setminus \{i\}} 1_{s_j}\big) - m_i(t(s_i), \sigma)\big| \leq \varepsilon$ for each player i. As $(s_1, \ldots, s_n)$ was an arbitrary element of F and $\tilde{\sigma}(F) > 1 - \eta$, we may conclude that $(\sigma_1, \ldots, \sigma_n)$ is an (ε, η)-ex post equilibrium of G. This completes the proof. □

A Appendix

In this appendix we prove the assertions made at the ends of Sections 3.1 and 3.2 about the sets K defined there.

A.1 Cournot Oligopoly

To see that the set K defined at the end of Section 3.1 has the properties claimed there, note first that the mapping $\mu \mapsto \int_S \mathrm{proj}_{[0,m]}\, d\mu$ from $M(S)$ to $\mathbb{R}$ is continuous because S and hence $M(S)$ are compact. Combining this fact with the fact that the function P is continuous, and that $[0, m]$ and $M(S)$ are compact, it is plain that the set K of functions is uniformly equicontinuous in the $M(S)$-component.

Next note that as the set T of functions is endowed with the (relativized) sup-norm topology, the mapping $t \mapsto t(a)$ from T to $\mathbb{R}$ is continuous for each $a \in [0, m]$. Thus since T, being compact, is separable, and since each element of T is measurable, the mapping $(t, a) \mapsto t(a)$ from $T \times [0, m]$ to $\mathbb{R}$ is measurable. Combining this fact with the fact that $(a, \mu) \mapsto P\Big(\frac{1}{n}a + \big(1 - \frac{1}{n}\big) \int_S \mathrm{proj}_{[0,m]}\, d\mu\Big)$ is continuous and hence measurable, it follows that each function $v \in K$ is measurable.

Finally, since the function P is bounded (being a continuous function on $[0, m]$), and since the set T is bounded (being a set of bounded functions on $[0, m]$ that is compact for the sup-norm), the set K of functions is bounded as well.

A.2 A Matching Game with Infinite Questions

We show that the set K of functions defined in Section 3.2 is bounded and uniformly equicontinuous in the $M(S)$-component, and that all of its elements are measurable.

To this end, note first that $\big|a_k - \lambda\mu(B_k^{-\theta(t)})\big| \leq 2$ for each $\lambda \in [1, 2]$ and each $k \in \mathbb{N}$. We may therefore define a function $\tilde{v}\colon [1, 2] \times S \times M(S) \to \mathbb{R}$ by

$$\tilde{v}(\lambda, (t, a), \mu) = \theta(t) \sum_{k=0}^{\infty} \omega(t)_k \big|a_k - \lambda\mu(B_k^{-\theta(t)})\big|.$$

The function $\tilde{v}$ is continuous. To see this, just note the following facts, which are immediate by the definitions and the topologies of the spaces involved:

1. The function $(\lambda, (t, a), \mu) \mapsto \theta(t)\big\langle\big|a_k - \lambda\mu(B_k^{-\theta(t)})\big|\big\rangle_{k\in\mathbb{N}}$ from $[1, 2] \times S \times M(S)$ into $[-2, 2]^{\mathbb{N}}$ is continuous when $[-2, 2]^{\mathbb{N}}$ is endowed with the product topology. (Note that for all $k \in \mathbb{N}$ and $\theta \in \Theta$, the set $B_k^{-\theta}$ is open and closed in S, which implies that $\mu \mapsto \mu(B_k^{-\theta})$ is continuous. Also note that if $t_i \to t$ in T, then $\theta(t_i) \in \{-1, -1\}$ must eventually be constant.)

2. The function $(\lambda, (t, a), \mu) \mapsto \langle\omega(t)_k\rangle_{k\in\mathbb{N}}$ from $[1, 2] \times S \times M(S)$ into Ω is continuous.

3. The function $(\omega, u) \mapsto \sum_{k=0}^{\infty} \omega_k \cdot u_k$ from $\Omega \times [-2, 2]^{\mathbb{N}}$ into $\mathbb{R}$ is continuous when $[-2, 2]^{\mathbb{N}}$ is endowed with the product topology.

Now since $\tilde{v}$ is continuous and its domain is compact, the set $\{\tilde{v}(\lambda, \cdot, \cdot)\colon \lambda \in [1, 2]\}$ is a bounded set of continuous and therefore measurable functions on $S \times M(S)$. The fact that $\tilde{v}$ is continuous and its domain $[1, 2] \times S \times M(S)$ is compact also implies that the set $\{\tilde{v}(\lambda, (t, a), \cdot)\colon \lambda \in [1, 2], (t, a) \in S\}$ is a uniformly equicontinuous set of functions on $M(S)$. Evidently for the set K defined in Section 3.2 we have $K \subseteq \{\tilde{v}(\lambda, \cdot, \cdot) \colon \lambda \in [1,2]\}$, showing that K is uniformly equicontinuous in the $M(S)$-component and bounded, with all of its elements being measurable.

References

ALIPRANTIS, C., AND K. BORDER (2006): *Infinite Dimensional Analysis.* Springer, Berlin, 3rd edn.

AZRIELI, Y., AND E. SHMAYA (2011): "Lipschitz Games," The Ohio State University and Northwestern University.

CARMONA, G. (2008): "Purification of Bayesian-Nash Equilibria in Large Games with Compact Type and Action Spaces," *Journal of Mathematical Economics*, 44, 1302–1311.

CASTAING, C., AND M. VALADIER (1977): *Convex Analysis and Measurable Multifunctions*, vol. 580 of *Lect. Notes Math.* Springer-Verlag, New York.

DASGUPTA, P., AND E. MASKIN (1986): "The Existence of Equilibrium in Discontinuous Economic Games, I: Theory," *Review of Economics Studies*, 53, 1–26.

DEB, J. (2008): "Essays in Economic Theory," Ph.D. thesis, Northwestern University.

DEB, J., AND E. KALAI (2010): "Stability in Large Bayesian Games with Heterogeneous Players," New York Univerty and Northwestern University.

FREMLIN, D. (2003): *Measure Theory, vol. 4: Topological Measure Spaces.* Torres Fremlin, Colchester.

GRADWOHL, R., AND O. REINGOLD (2010): "Partial Exposure in Large Games," *Games and Economic Behavior*, 68, 602–613.

JONES, L. E. (1984): "A Competitive Model of Commodity Differentiation," *Econometrica*, 52, 507–530.

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

Mandelkern, M. (1990): "On the Uniform Continuity of Tietze Extensions," *Archiv der Mathematik*, 55, 387–388.

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