EQUILIBRIUM EXISTENCE IN GAMES WITH A CONCAVE BAYESIAN POTENTIAL

Ezra Einy and Ori Haimanko

Discussion Paper No. 20-02

June 2020

Monaster Center for
Economic Research
Ben-Gurion University of the Negev
P.O. Box 653
Beer Sheva, Israel

Fax: 972-8-6472941
Tel: 972-8-6472286

Equilibrium Existence in Games With a Concave Bayesian Potential*

Ezra Einy† and Ori Haimanko‡

March 2020

Abstract

We establish existence of a pure-strategy Bayesian Nash equilibrium in Bayesian games with convex and compact action sets that have an upper semi-continuous and concave potential (or a weighted version thereof) at any state of nature. No assumptions are made on the information structure in these games; in particular, there may be uncountably many states of nature or information types, and in the latter case the common prior need not be absolutely continuous w.r.t. the product of its marginals. As an application, we show that Bayesian Nash equilibrium exists in many well-known games and their generalizations that have semi-quadratic payoffs, including Bertrand and Cournot oligopolies with linear demand.

Journal of Economic Literature classification numbers: C62, C72, D82.

Key words: Bayesian games, (weighted) Bayesian potential, equilibrium existence, concave payoffs, absolute continuity, information structures.

1 Introduction

The extensive use of Bayesian games in economic theory has been made possible by the fact that quite general categories of games with incomplete information possess a

*This is an expanded version of DP 19-11 of Monaster Center for Research in Economics, Ben-Gurion University of the Negev.

†Department of Economics, Ben-Gurion University of the Negev, Israel. e-mail: einy@bgu.ac.il

‡Department of Economics, Ben-Gurion University of the Negev, Israel. e-mail: orih@bgu.ac.il.

Bayesian Nash equilibrium (henceforth, BNE). The path-breaking work of Milgrom and Weber (1985) was the first to prove BNE existence beyond the finite framework of Harsanyi (1967), and it did so with remarkable generality: players' action and type sets were allowed to be (possibly uncountable) metric spaces.[1] The restriction imposed on the Bayesian game in their BNE existence result combined two conditions. One is the continuity of the players' payoff functions on the set of action profiles for any realization of the players' types; the other requires absolute continuity of information, meaning that the joint distribution of the players' types must be absolutely continuous with respect to the product of its marginal distributions.

The usefulness of the absolute continuity condition is demonstrated by its applicability in many benchmark cases considered in economic theory, such as those where the players' types are independently distributed, or merely have joint density, and also when the type sets are finite or countable. Most of the literature devoted to extensions of the Milgrom and Weber result has, too, assumed absolute continuity of information or its variants, while focusing on a relaxation of the payoff continuity assumption (see, e.g., Carbonell-Nicolau and McLean (2018), He and Yannelis (2016) and the references therein). Restricting attention to absolutely continuous information is definitely not a matter of convenience, however. That is because BNE may fail to exist without that restriction even if each player has finitely many actions, as was shown by Simon (2003). (This notwithstanding, there are classes of games where information is not absolutely continuous, in a particularly flagrant way, that possess a BNE.[2])

A natural question that arises is whether there are plausible requirements on games that are played at every *state of nature*[3] that would guarantee existence of a BNE in the corresponding Bayesian game for *any* underlying information structure,

[1] While action sets need to be compact, all topological assumptions on the type sets in Milgrom and Weber's BNE existence result were removed in Balder (1988).

[2] Hellman and Levy (2017) characterize such a class in the domain of Bayesian games with finitely many actions that are "purely atomic," i.e., such that all type-sets are a non-atomic continuum, but, given his type, each player knows with certainty that the others' types belong to a finite or countable set.

[3] Henceforth, we will consistently refer to the *states of nature* model of incomplete information, which supersedes a simpler model of Harsanyi types. In the context of the latter model, "state" is simply a realization of all players' information types.

absolutely continuous or not. One set of such requirements was suggested by Mamer and Schilling (1986) and Einy et al. (2008): they have shown that when state-payoffs are zero-sum and continuous in each action separately, then a BNE exists for general information structures in a two-player setting. In this work, we will show that another condition on the state-game – having a Bayesian potential with certain properties – guarantees BNE existence for any information structure.

In a Bayesian potential game, a notion that was introduced by Heumen et al. (1996), a potential game[4] is played at any state of nature.[5] The state-dependent potential function for the state games, called Bayesian potential, is a natural tool as far as proving BNE existence is concerned: it has been well understood that any maximizer of the *expectation* of a Bayesian potential over the set of all *pure* Bayesian strategy profiles is a BNE of the Bayesian potential game.[6] We will use a generalized version of Bayesian potential games, in which the state-dependent potential is not (necessarily) exact but weighted, in the way that players' (possibly state-dependent) weights w are measurable w.r.t. their private information.[7] Although, unlike in the weight-free context, the expected w-Bayesian potential is not (necessarily) a normal-form weighted potential for the w-Bayesian potential game, we will show that the expected potential's maximizer, *if one exists*, is a BNE.

When the space of states of nature in a game is uncountable, what stands in the way of proving the existence of such a maximizer is the fact that topologies in which the expected w-Bayesian potential would normally be continuous are in general too strong to make the set of pure Bayesian strategy profiles compact.[8] However, it turns out that, with action sets being compact and convex subsets of a Euclidean space, the topological tension between continuity and compactness does not arise when a w-Bayesian potential is continuous and *concave* in all states of nature. The proof is

[4]Here a standard, Monderer and Shapley (1996) potential game, is meant.

[5]While Heumen et al. (1996) confined themselves to finitely many states of nature, Ui (2009) extended the concept of a Bayesian potential game to infinite information structures.

[6]That is because the expected Bayesian potential is a Monderer and Shapley (1996) potential for the normal form of the Bayesian potential game.

[7]This also generalizes the notion of a weighted Bayesian potential game of Ui (2009), in which the weights are state-independent.

[8]For a detailed discussion in the context of general utility functions, see, e.g., p. 626 in Balder and Yannelis (1993).

based on results of Balder and Yannelis (1993). They characterized *weakly compact* sets of contingent consumption plans, and showed that, for concave and upper semi-continuous state-utilities, the expected utility is *upper semi-continuous* in the *weak topology* on such sets of contingent plans. In the Bayesian games context, their results can be directly applied to the set of profiles of "pseudo-strategies" (which fictitiously endow all players have full information), showing that this set is weakly compact and, assuming the w-Bayesian potential to be upper semi-continuous and concave at each state, that the expected w-Bayesian potential is weakly upper semi-continuous. Our proof will show that constraints of measurability w.r.t. private information can be added in pseudo-strategy profiles (turning them into true strategy profiles) without affecting the conclusions. This combination of compactness and upper semi-continuity implies the existence of a Bayesian potential maximizer, and hence of a BNE in pure strategies. Importantly, this argument for BNE existence does not exploit any particular attributes of the information structure, and thus our existence result holds in fullest possible generality in that respect.[9]

The method of finding a BNE as a maximizer of a common, real or fictitious, expected payoff function has been considered previously, in a strand of literature that grew out of the work of Radner (1962). Radner considered "team games," where the players have a common payoff (hence, a potential) that is a concave quadratic polynomial in the players' actions at each state of nature. He showed existence of a maximizer of the expected payoff in pure Bayesian strategies, under the assumption that uncertainty affects only the linear term of the payoff, and that the players' signals and the coefficients of the linear term have a joint normal distribution.[10] The games of Radner (1962) were found to be useful in studying information effects in linear Cournot and Bertrand oligopoly models, as it was implicitly recognized that some specifications of Radner's quadratic payoff function can serve as concave Bayesian

[9] There is a considerable strand of literature that studies existence of *pure strategy* BNE in Bayesian games, that began with Radner and Rosenthal (1982). Their framework and results have been since significantly extended (see, e.g., He and Sun (2019) and the references therein), but the assumptions consistently exclude non-absolutely continuous information and type sets with atoms.

[10] When all parameters of the quadratic payoff are uncertain and have a general distribution, Radner offered a sufficient condition for the maximum existence. It is summarized in Footnote 24 of our Section 4.2.

potentials for such oligopolies with incomplete information on various parameters (see Raith (1996) for a unifying approach and a survey). The first explicit use of Radner's game as a concave Bayesian potential was in Ui (2009), who applied Radner's BNE characterization in a study of efficient information use in a class of Bayesian games with quadratic payoffs.[11]

Our result on BNE existence applies in the above-mentioned contexts because, for those Bayesian games, Radner's payoff function constitutes (an exact) Bayesian potential that is continuous and concave at each state. But our result also extends the scope of the earlier findings in two important respects. First, since it asserts BNE existence without any restriction on the information structure, players' signals and the game parameters need not have a joint normal distribution, and, in fact, need not have joint (or any) density. And second, the specific quadratic form of payoffs can be generalized to a *semi-quadratic* one, which allows components that are non-linear (but concave) functions of own actions, without affecting BNE existence.[12]

The paper is organized as follows. In Section 2 we describe the general set-up, recall the notions of a Bayesian game and BNE, and define w-Bayesian potential games. Section 3 contains our BNE existence result, and applications are discussed in Section 4. All proofs are given in the Appendix.

2 Bayesian potential games

2.1 Bayesian games

Let $N = \{1, ..., n\}$ be a finite set of players. Games are played in an uncertain environment. The underlying uncertainty is described by a probability space $(\Omega, \mathcal{F}, \mu)$, where Ω is a set of states of nature, $\mathcal{F}$ is a σ-field of measurable events, or subsets of Ω, and μ is a countably additive probability measure on $(\Omega, \mathcal{F})$, representing the common prior belief of the players about the actual state of nature. Private information of player $i \in N$ is given by a σ-subfield $\mathcal{F}_i$ of $\mathcal{F}$, consisting of events that are

[11] That class includes linear oligopolies as well as variants of games considered in Crémer (1990) and Morris and Shin (2002).

[12] That is because a Bayesian potential that is concave at each state would still be easily constructible for such games.

discernible by i.

For each $i \in N$ there is $d_i \in \mathbb{N}$ such that, at any $\omega \in \Omega$, player i has a set $A_i(\omega)$ of actions that is a non-empty, convex and compact subset of $\mathbb{R}^{d_i}$; that is, i's action set may depend on the state of nature. Let us denote by A_i the corresponding set-valued function $A_i : \Omega \to 2^{\mathbb{R}^{d_i}}$, and define $A : \Omega \to 2^{\mathbb{R}^d}$ (where $d := \sum_{i=1}^n d_i$) as $A(\omega) = A_1(\omega) \times ... \times A_n(\omega)$ for every $\omega \in \Omega$; the values of A are also non-empty, convex and compact. We will assume that the graph of A (and, respectively, of each A_i) is $\mathcal{F} \otimes \mathcal{B}(\mathbb{R}^d)$-measurable[13] (respectively, $\mathcal{F}_i \otimes \mathcal{B}(\mathbb{R}^{d_i})$-measurable) and that $\sup_{a \in A(\cdot)} \|a\|_{\mathbb{R}^d}$ is μ-integrable (i.e., A and each A_i are *integrably bounded*).

Each $i \in N$ has a payoff function $u_i : \Omega \times \mathbb{R}^d \to \mathbb{R}$. We will assume that u_i is $\mathcal{F} \otimes \mathcal{B}(\mathbb{R}^d)$-measurable and that $\sup_{a \in A(\cdot)} |u_i(\cdot, a)|$ is μ-integrable. A Bayesian game will be identified with the collection of its above-described attributes, $G = (N, (\Omega, \mathcal{F}, \mu), (\mathcal{F}_i, A_i, u_i)_{i=1}^n)$.

A (pure Bayesian) strategy of player $i \in N$ in the game G is an $\mathcal{F}_i$-measurable function $x_i : \Omega \to \mathbb{R}^{d_i}$, with $x_i(\cdot) \in A_i(\cdot)$ μ-a.e. The set of all strategies of player i will be denoted by X_i, which is non-empty by the measurable selection theorem. Each player i evaluates his *ex-ante* prospect in the game via the expected payoff function U_i on the product set $X = X_1 \times ... \times X_n$ of strategy profiles, given by

$$U_i(x) = \int_\Omega u_i(\omega, x(\omega)) d\mu(\omega) \tag{1}$$

for any $x = (x_1, ..., x_n) \in X$.[14] As usual, $x \in X$ is a (pure-strategy) Bayesian Nash equilibrium of the game G, or BNE for short, if it is a Nash equilibrium of the normal form of G, namely, if the inequality

$$U_i(x) \geq U_i(y_i, x_{-i})$$

holds for every $i \in N$ and $y_i \in X_i$, where $(y_i, x_{-i}) \in X$ denotes the strategy profile obtained from x by substituting y_i for x_i.

[13] Here and henceforth, $\mathcal{B}(K)$ will denote the Borel σ-field on a Borel set K in some Euclidean space.

[14] Note that the function $u_i(\cdot, x(\cdot))$ is $\mathcal{F}$-measurable as a composition of an $\mathcal{F} \otimes \mathcal{B}(\mathbb{R}^d)$-measurable u_i with an $\mathcal{F}$-measurable $\Omega \times \mathbb{R}^d$-valued function $\omega \mapsto (\omega, x_1(\omega), ..., x_n(\omega))$. The function $u_i(\cdot, x(\cdot))$ is therefore also μ-integrable, being bounded in absolute value from above by $\sup_{a \in A(\cdot)} |u_i(\cdot, a)|$ (which is μ-integrable by assumption).

2.2 w-Bayesian Potential games

Consider a Bayesian game $G = \left(N, (\Omega, \mathcal{F}, \mu), (\mathcal{F}_i, A_i, u_i)_{i=1}^{n}\right)$ and an n-tuple $w = (w_i)_{i\in N}$ of *weight functions* such that $w_i : \Omega \to (0, \infty)$ is $\mathcal{F}_i$-measurable for each $i \in N$. We say that G is a *w-Bayesian potential game* (or w-BP game for short) if there exists $p : \Omega \times \mathbb{R}^d \to \mathbb{R}$ (called a *w-Bayesian potential*, or w-BP, for G) that satisfies the following:

(a) p is $\mathcal{F} \otimes \mathcal{B}\left(\mathbb{R}^d\right)$-measurable;

(b) there exist a μ-integrable $\psi : \Omega \to [0, \infty)$ and a constant $M > 0$ such that

$$|p(\omega, a)| \leq \psi(\omega) + M \|a\|_{\mathbb{R}^d} \tag{2}$$

for every $\omega \in \Omega$ and $a \in \mathbb{R}^d$;

and

(c) for μ-almost every $\omega \in \Omega$, every $i \in N$, and every $a \in A(\omega)$, $b_i \in A_i(\omega)$,

$$u_i(\omega, (b_i, a_{-i})) - u_i(\omega, a) = w_i(\omega)\left[p(\omega, (b_i, a_{-i})) - p(\omega, a)\right] \tag{3}$$

(where $(b_i, a_{-i}) \in A(\omega)$ is the action profile obtained from a by substituting b_i for a_i).

Note that a w-BP game G consists of playing a weighted potential game (in the sense of Monderer and Shapley (1996)) at any state of nature ω, but the weights $(w(\omega))_{i\in N}$ may vary with the state of nature. Our notion of a w-BP game has two well-known simpler versions that it generalizes.[15] When the weight functions w are constant, that is, invariant of the state of nature, G is known as a *weighted Bayesian potential game*; these games were introduced by Facchini et al. (1997) for a finite Ω, and by Ui (2009) for general Ω. In turn, the weighted notion extends the basic concept of an (exact) *Bayesian potential game* (BP game, for short), originally due to van Heumen et al. (1996), for which all weights are identical and equal to 1. If G is a BP game, condition (3) on the corresponding p (which in the weight-free context is referred to simply as *Bayesian potential*, or BP) becomes particularly straightforward: it is required that

$$u_i(\omega, (b_i, a_{-i})) - u_i(\omega, a) = p(\omega, (b_i, a_{-i})) - p(\omega, a) \tag{4}$$

[15] Our notion is also related to the one that was described in the statement of Theorem 5 in Ui (2009), in the context of differentiable payoff functions.

for μ-almost every $\omega \in \Omega$, every $i \in N$, and every $a \in A(\omega)$, $b_i \in A_i(\omega)$. Thus, changes in the state-payoff resulting from unilateral deviations by any player are precisely accounted for by changes in the BP, at every state of nature.

Given a w-BP game G with a w-BP p, consider its *expected potential* (or EP for short), $E(p) : X \to \mathbb{R}$, given by

$$E(p)(x) = \int_{\Omega} p(\omega, x(\omega)) d\mu(\omega) \tag{5}$$

for any $x = (x_1, ..., x_n) \in X$. If G is a BP game and p is its BP, then the corresponding EP obviously retains the property expressed in (4), which is now given in terms of the players' expected payoffs:

$$U_i(y_i, x_{-i}) - U_i(x) = E(p)(y_i, x_{-i}) - E(p)(x) \tag{6}$$

for every $i \in N$ and every $x \in X$, $y_i \in X_i$. Thus, if G is a BP game then its normal form is a potential game in the usual sense (of Monderer and Shapley (1996)), and, clearly, any maximizer $x \in X$ of its normal-form potential $E(p)$ is a BNE of G.

It is important to note that (4) would not typically hold if G is a w-BP game for weight functions w that are not identically 1; in general, there need not be any simple relation between the marginal changes of the expected payoffs and those of the EP.[16] However, one obvious implication of (4), that a unilateral maximization of the player's expected payoff leads him to the same strategies as a unilateral maximization of the EP, turns out to be a property common to all w-BP games. According to the following lemma, unilateral maximization of the EP preserves all best responses in such games (or, using the terminology of Ui (2009), the EP is a "Bayesian best-response potential" for the game).

Lemma. Let G be a game with a w-BP p for some collection w of weight functions. Then, for any $x \in X$ and any $i \in N$, $\arg\max_{y_i \in X_i} U_i(y_i, x_{-i}) = \arg\max_{y_i \in X_i} E(p)(y_i, x_{-i})$.

The proof is given in the Appendix.

Remark. The assumption of $\mathcal{F}_i$-measurability of each player i's weight function w_i in the definition of a w-BP is necessary in order for the Lemma to hold. Indeed,

[16] One such relation exists when G is a weighted Bayesian potential game: if all weight functions w_i have constant (but possibly distinct) values, (6) holds in a modified form, with the right-hand side expression being multiplied by the constant value of w_i.

consider a one-player game (i.e., a decision problem) with two equiprobable states of nature in $\Omega = \{0, 1\}$, $\mathcal{F}_1 = \{\varnothing, \Omega\}$, $A_1 = [0, 1]$, and the payoff function $u_1(\omega, a_1) = -(\omega + 1)(a_1 - \omega)^2$. Clearly, strategies of player 1 can be identified with his actions, and $\arg\max_{y_1 \in X_1} U_1(y_1) = \{\frac{2}{3}\}$. Now notice that $p(\omega, a_1) = -(a_1 - \omega)^2$ satisfies (3) for $i = 1$ and $w_1(\omega) = \omega + 1$, and that the function w_1 is not $\mathcal{F}_1$-measurable. Also, $\arg\max_{y_1 \in X_1} E(p)(y_1) = \{\frac{1}{2}\}$. Thus, p cannot be just any weighted potential (or, more generally, ordinal potential) in the state-games.

The following is an obvious but important corollary of the Lemma:

Corollary. If G is a game with a w-BP p for some collection w of weight functions, then any maximizer $x \in X$ of $E(p)$ is a BNE of G.

3 BNE Existence

The existence of a w-BP that is concave and upper semi-continuous at almost every state implies existence of a BNE in the game *without* any assumption on the information structure. In particular, the set of states of nature Ω may be uncountable, and players' private information may be given by σ-fields that are not generated by partitions of Ω.

Theorem. Let $G = (N, (\Omega, \mathcal{F}, \mu), (\mathcal{F}_i, A_i, u_i)_{i=1}^n)$ be a game with a w-BP p for some collection w of weight functions. Assume that, for μ-almost every $\omega \in \Omega$: (i) $p(\omega, \cdot)$ is concave on $A(\omega)$; and (ii) $p(\omega, \cdot)$ is upper semi-continuous[17] on $A(\omega)$. Then G possesses a (pure-strategy) BNE.

We will now comment on the method of the theorem's proof (which appears in the Appendix), and describe its structure. It follows from the Corollary that establishing BNE existence is reducible to showing that the expected potential $E(p)$ attains a maximum over the set X of strategy profiles. The existence of a maximum is explored in the following way. The definition of $E(p)$ in (5) allows to extends its

[17] It is well-known that any concave function on a convex polytope is *lower* semi-continuous (see, e.g., Gale et al. (1968)). Hence, if $A(\omega)$ is a polytope then we, in effect, assume that $p(\omega, \cdot)$ is continuous.

domain from X to $X' =$ the set of all A-valued and $\mathcal{F}$-measurable functions on Ω. (In defining X', we drop the the condition of $\mathcal{F}_i$-measurability that characterizes each i's component of a strategy profile in X, and so X' can be regarded as consisting of profiles of "pseudo-strategies," in which all players have the same information field $\mathcal{F}$).[18] We then apply some of the results of Balder and Yannelis (1993) (who in turn obtain them using strong tools[19] from functional analysis) to establish the following two facts: given our assumptions on A and p, the set X' is compact in the weak topology and $E(p)$ is (weakly) upper semi-continuous on X'.

Although the above implies existence of $\max_{x' \in X'} E(p)(x')$, it remains to show that a maximum exists on a smaller domain of true strategy-profiles, X. We will prove that X is, in fact, a weakly closed subset of X' by first showing that it is strongly closed, and then applying Mazur's theorem on the equivalence of strong and weak closedness of convex sets in a Banach space. The implied weak compactness of X then guarantees the existence of $\max_{x \in X} E(p)(x)$, and hence of a BNE in the game G.

4 Applications

Our existence result can be applied in a number of well-recognized contexts, which are presented in the following subsections. The player set N, the space $(\Omega, \mathcal{F}, \mu)$ and the private information fields $(\mathcal{F}_i)_{i=1}^n$ will be fixed throughout. Although not by any means necessary, it will be assumed for simplicity of presentation that all players' action sets are state-independent, and so A and each A_i will henceforth be treated as *subsets* of the corresponding Euclidean spaces, not as set-valued functions. The payoff functions and the BPs will only be defined on $\Omega \times A$; they can be extended arbitrarily in a measurable way onto $\Omega \times \mathbb{R}^d$ in order to fit into the framework of Section 2.

[18]The notation X' is used, for brevity, only in this summary; the actual set of "pseudo-strategies" is $L^1((\Omega, \mathcal{F}, \mu); A)$, precisely defined in the proof.

[19]These are Mazur's theorem (see, e.g., Corollary 23 in Royden (1988)) and Diestel's theorem (see Diestel (1977)).

4.1 Motivating model: oligopoly with linear demand

Cournot oligopoly is a showcase of the usefulness of a Bayesian potential approach. We consider the following description of the model, partially based on Raith (1996). The members of N are firms; each $i \in N$ produces a separate good (also denoted by i), and its action set $A_i \subset \mathbb{R}_+$ is a compact interval of possible output levels of good i. In choosing output level a_i, firm i incurs a state-dependent production cost of $c_i(\omega, a_i)$, where $c_i : \Omega \times A_i \to \mathbb{R}_+$ is an $\mathcal{F} \otimes \mathcal{B}(A_i)$-measurable function that is continuous and convex in its second variable a_i, and integrably bounded. The state-dependent linear *inverse demand* (i.e., price function) of the firms' output is given by

$$\mathbf{P}_i(\omega, a) = \mathbf{A}_i(\omega) - \sum_{j \neq i} \varepsilon(\omega) a_j - \delta(\omega) a_i \tag{7}$$

for every $\omega \in \Omega$, $a \in A$, where $(\mathbf{A}_i)_{i=1}^n$, ε and δ are $\mathcal{F}$-measurable and μ-integrable functions, with $(\mathbf{A}_i)_{i=1}^n$ and δ being strictly positive and $\varepsilon(\omega) \in (-\frac{\delta(\omega)}{n-1}, \delta(\omega)]$ for every $\omega \in \Omega$. The state-dependent net-profit function of firm i is therefore

$$u_i(\omega, a) = \left(\mathbf{A}_i(\omega) - \sum_{j \neq i} \varepsilon(\omega) a_j - \delta(\omega) a_i\right) a_i - c_i(\omega, a_i), \tag{8}$$

for every $\omega \in \Omega$ and $a \in A$. It is easy to see that the following function $p : \Omega \times A \to \mathbb{R}$ is an (exact) BP for our incomplete information oligopoly:

$$p(\omega, a) = \sum_{i=1}^n \mathbf{A}_i(\omega) a_i - \left(\delta(\omega) \sum_{i=1}^n a_i^2 + \varepsilon(\omega) \sum_{1 \le i < j \le n} a_i a_j\right) - \sum_{i=1}^n c_i(\omega, a) \tag{9}$$

for every $\omega \in \Omega$ and $a \in A$. By our assumptions on $(\mathbf{A}_i)_{i=1}^n$, ε, δ and $(c_i)_{i=1}^n$, p is $\mathcal{F} \otimes \mathcal{B}(A)$-measurable and integrably bounded, and it can be readily seen that the function $p(\omega, \cdot)$ is concave and continuous for any fixed $\omega \in \Omega$. Hence, the oligopoly falls within the purview of our theorem – it has a BNE, and BNE existence is obtained without any direct restriction on the information structure. In contrast, the BNE existence result in Raith (1996) is predicated upon ε, δ being state-independent (i.e., known), costs being linear, and all uncertain parameters having a joint normal distribution with the players' private signals.[20]

[20] Since prices are rarely negative in reality, the functional form in (8) – even if believed to be a good approximation of the real-world consumer demand – may need to be truncated when reaching

As a particular case, when $\mathbf{A}_i = \mathbf{A}$ for all i and $\varepsilon = \delta$, we obtain a Cournot oligopoly with a single homogeneous good, which (in the complete information setting) served as the first example of a potential game in Monderer and Shapley (1996). On the other hand, when all $(c_i)_{i=1}^n$ are taken to be zero, and the actions of firms are the *prices* they charge for their goods rather than the quantities that they produce, equation (7) can be viewed as a description of a state-dependent linear *demand* for good i given the vector a of prices, and hence (8) can be viewed as a payoff function in a *Bertrand oligopoly* with price competition. Thus, such Bertrand oligopoly is also a BP game, with the ensuing claim of BNE existence.[21]

4.2 Games with semi-quadratic payoffs

The first, quadratic, term of the firm's utility function (8) in the oligopoly model of Section 4.1 points towards some natural generalizations. Common concave payoffs of quadratic form have been considered by Radner (1962) in the context of "team games," for which he established the existence of a BNE under an implicit integrability-related condition linking the game parameters and its information structure. We will follow Ui's (2009) account[22] that views those common payoffs as (exact) BPs for a sizable category of payoff functions. Ui's payoffs will be generalized in the following respect: the term that depends on the player's own action will not necessarily be linear.

Assume that A_i is a compact interval for each $i \in N$, and that each i's payoff function has the following, *semi-quadratic*, form:

$$u_i(\omega, a) = -\frac{1}{2} q_{ii}(\omega) a_i^2 - a_i \sum_{j \neq i} q_{ij}(\omega) a_j + f_i(\omega, a_i) + h_i(\omega, a_{-i}), \tag{10}$$

for every $\omega \in \Omega$ and $a \in A$, where $Q(\omega) = [q_{ij}(\omega)]_{n \times n}$ is an $\mathcal{F}$-measurable, μ-integrable and symmetric matrix, $f_i : \Omega \times A_i \to \mathbb{R}$ is $\mathcal{F} \otimes \mathcal{B}(A_i)$-measurable and

zero. This truncation may cause non-existence of BNE even in two-states-of-nature settings, as shown by Einy et al. (2010). It is, therefore, advisable to check that any positive-price BNE found in the current model remains such when the inverse demand is truncated to keep prices non-negative.

[21] Notice also that linear costs of output can be added to payoff functions, and accommodated by the potential.

[22] Following Radner (1962), Ui (2009) found closed-form expressions for the unique BNE equilibria in certain contexts when the game's linear parameters and the players' signals have a joint normal distribution.

integrably bounded, and $h_i : \Omega \times A_{-i} \to \mathbb{R}$ is $\mathcal{F} \otimes \mathcal{B}(A_{-i})$-measurable and integrably bounded.[23] It is easy to see that the game has an (exact) BP, p, that is given by

$$p(\omega, a) = -\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij}(\omega) a_i a_j + \sum_{i=1}^{n} f_i(\omega, a_i) \tag{11}$$

for every $\omega \in \Omega$ and $a \in A$.

We will henceforth assume that, at every $\omega \in \Omega$, the matrix $Q(\omega)$ is positive semi-definite and each $f_i(\omega, \cdot)$ is continuous and concave, which obviously implies that the BP p is concave (and continuous) in a. BNE existence is, therefore, guaranteed by our theorem, regardless of what information structure is imposed on the game. To compare, the sufficient condition in the general existence result of Radner (1962) (namely, his Theorems 2 and 3) links together the information structure and the parameters of the game,[24] requires Q to be (strictly) positive definite, and, most importantly, the functions $(f_i)_{i=1}^{n}$ in (10) need to be *linear* in the second variable. What Radner's result affords, however, is the possibility to work with an unrestricted action set $\mathbb{R}$, instead of *a priori* confining actions to compact intervals as we do.

Notice that when $q_{ii}(\omega) = 2\delta(\omega)$, $q_{ij}(\omega) = \varepsilon(\omega)$ if $i \neq j$, $f_i(\omega, a_i) = \mathbf{A}_i(\omega) - c_i(\omega, a_i)$ and $h_i \equiv 0$, (10) and (11) correspond to (8) and (9) in the case of Cournot oligopoly with linear demand that was analyzed in Section 4.1. In the following examples we will briefly describe some other specific classes of incomplete information games that the semi-quadratic functional form in (10) can accommodate.

Example 1 (Network games). In a network game, players' payoffs depend on the realized action profile $a \in \mathbb{R}_{+}^{N}$ and on the network (i.e., a graph) that links different players to one another. We consider a semi-quadratic generalization of one of the network game analyzed in Bramoullé et al. (2014) (based, in turn, on the model in Ballester et al (2006)), in which player i's payoff is

$$u_i(a_i, a_{-i}) = f_i(a_i) - \tfrac{1}{2} a_i^2 - \delta \sum_{j=1}^{n} g_{ij} a_i a_j,$$

[23] Here, as usual, A_{-i} stands for $\times_{j \neq i} A_j$, and $a_{-i} \in A_{-i}$ is obtained by omitting the i^{th} coordinate of a.

[24] If stated in the present set-up, the condition requires an $\mathcal{F}$-measurable state-by-state maximizer z of the potential p to have a finite "distance" from at least one strategy profile $x \in X$, in the sense that $\int_\Omega \sum_{i=1}^{n} \sum_{j=1}^{n} q_{ij}(\omega)(x_i(\omega) - z_i(\omega))(x_j(\omega) - z_j(\omega))\, d\mu(\omega) < \infty$.

where f_i is an increasing, continuous and concave function that vanishes at 0, $\delta > 0$, the values $g_{ij} \in \{0,1\}$ indicate whether players i and j are linked or not, $g_{ii} \equiv 0$ and $g_{ij} = g_{ji}$ for every $i \neq j$; w.l.o.g., each player i can be constrained to use actions in some compact interval $A_i = [0, M]$. Thus, each player's activity has decreasing returns to scale, and he is subject to negative externality from being linked to other players. By making f_i, the externality parameter δ and the link matrix $[g_{ij}]_{n\times n}$ state-dependent (in a measurable, integrable fashion), this game turns into a BP game, with a BP p that is given by

$$p(\omega, a) = \sum_{i=1}^{n} f_i(\omega, a_i) - \frac{1}{2}\sum_{i=1}^{n} a_i^2 - \delta(\omega) \sum_{1\leq i<j\leq n}^{n} g_{ij}(\omega) a_i a_j$$

for every $\omega \in \Omega$ and $a \in [0, M]^n$. The function $p(\omega, \cdot)$ is obviously continuous. It is also concave if the matrix[25] $[\delta_{ij} + \delta(\omega) g_{ij}(\omega)]_{n\times n}$ is positive semi-definite at each state of nature, and a BNE then exists by our theorem.

Example 2 (Coordination games). In Ui's (2009) two-player version of the game of Morris and Shin (2002), each player needs to take an action serving two possibly conflicting objectives: being close to (what is required by) the fundamental state $\theta(\omega)$,[26] and being close to the action of the other player (in the spirit of Keynes's "beauty contest" example). His utility function additively combines two loss terms representing the two objectives: for each $i = 1, 2$,

$$u_i(\omega, a) = -\lambda(a_i - \theta(\omega))^2 - (1-\lambda)(a_i - a_j)^2$$

for some $0 < \lambda < 1$, and for every $\omega \in \Omega$ and $a \in \mathbb{R}_+^2$. As a BP, one may use the function given by

$$p(\omega, a) = -\lambda(a_1 - \theta(\omega))^2 - \lambda(a_2 - \theta(\omega))^2 - (1-\lambda)(a_1 - a_2)^2$$

that is obviously continuous and concave in a. As long as θ^2 is μ-integrable and action sets are truncated from above (with a weak inequality), our theorem assures BNE existence under any information structure.

[25] Here δ_{ij} is the Kronecker delta.

[26] In Ui's (2009) specification, θ has a joint normal distribution with signals that the two player obtain (and that constitute their private information).

Example 3 (**Team-theoretical model of a firm**). Crémer (1990) considered a model in which two agents with a common interest have uncertainly about a single integrable parameter[27] $\theta(\omega)$ that affects as follows their (identical) utilities:

$$u_1(\omega, a) = u_2(\omega, a) = \theta(\omega)(a_1 + a_2) - \frac{B(a_1 + a_2)^2 - C(a_1 - a_2)^2}{2} \tag{12}$$

for some $B, C > 0$, and for every $\omega \in \Omega$ and $a \in \mathbb{R}^2_+$. (The case of $B > C$ corresponds to strategic substitutability of actions, while the case of $C > B$ to strategic complementarity.) Our result guarantees BNE existence in general when the action sets are weakly truncated from above, since the common utility – which is also a BP – is clearly continuous and concave in a. Moreover, the influence of the parameter θ on the actions' direct impact need not be linear: the first term in (12) can be replaced by any integrably bounded function of ω that is continuous and concave in a_1 and a_2 without affecting BNE existence.

Our last example retains the quadratic form of utility functions but has multi-dimensional strategy sets.

Example 4 (Routing problems). In a class of routing problems described in Altman et al. (2007), a transportation network is modelled as a directed graph. Each player i decides how to split his traffic of size $\Lambda_i > 0$ (that needs to pass from an i-specific "source" node to a "destination" node on the graph) between the links in the graph. The action set A_i of player i is thus a subset of $[0, \Lambda_i]^L$ (where L denotes the set of links) of traffic volume assignments that satisfy flow-conservation constraints,[28] which is convex and compact. It is assumed that a per-unit common congestion (dis)utility at a link l has the form $c_l(v) = b_l + d_l v$ for a total traffic volume v passing through l (where $b_l, d_l < 0$). Player i's utility is then the total of his (dis)utility experienced at all links, namely,

$$u_i(a) = \sum_{l \in L} \left[b_l + d_l \sum_{j=1}^{n} a_j(l) \right] a_i(l)$$

for each $a = \left((a_i(l))_{l \in L} \right)_{i=1}^{n}$, where $a_j(l)$ denotes the volume of traffic put by player

[27] See the previous footnote.

[28] For a full desription, see p. 2 in Altman et al. (2007).

j through link l. Clearly, the function p that is given by

$$p(a) = \sum_{l \in L} \left[b_l \sum_{i=1}^{n} a_i(l) + d_l \left(\sum_{i=1}^{n} a_i^2(l) + \sum_{1 \le i < j \le n} a_i(l)\, a_j(l) \right) \right],$$

for any $a = \left((a_i(l))_{l \in L}\right)_{i=1}^{n}$, is a potential for the game, and it is strictly concave and continuous in a. The extension to the incomplete information case, with a concomitant claim of BNE existence, can be performed effortlessly (similarly to what has been done, e.g., in Example 1), by adding uncertainty on the parameters $(b_l)_{l \in L}$ and $(d_l)_{l \in L}$.

5 Appendix

Proof of the Lemma. Fix $i \in N$, and notice that $z_i \in \arg\max_{y_i \in X_i} U_i(y_i, x_{-i})$ if and only if the inequality

$$\int_C [u_i(\omega, (z_i(\omega), x_{-i}(\omega))) - u_i(\omega, (y_i(\omega), x_{-i}(\omega)))] d\mu(\omega) \ge 0 \qquad (13)$$

holds for every $y_i \in X_i$ and $C \in \mathcal{F}_i$ (the "if" direction is trivial, and to see the "only if" direction consider only those deviations y_i that coincide with z_i on $\Omega \backslash C$).

Next, for any μ-integrable random variable Z on $(\Omega, \mathcal{F}, \mu)$, denote by $\mathbb{E}[Z(\cdot) \mid \mathcal{F}_i]$ the conditional expectation[29] of Z given the field $\mathcal{F}_i$. By taking $Z(\cdot) = u_i(\cdot, (z_i(\cdot), x_{-i}(\cdot))) - u_i(\cdot, (y_i(\cdot), x_{-i}(\cdot)))$ for some $y_i, z_i \in X_i$, the defining property of the conditional expectation guarantees that, for every $C \in \mathcal{F}_i$,

$$\begin{aligned} & \int_C [u_i(\omega, (z_i(\omega), x_{-i}(\omega))) - u_i(\omega, (y_i(\omega), x_{-i}(\omega)))] d\mu(\omega) \\ = & \int_C \mathbb{E}[u_i(\cdot, (z_i(\cdot), x_{-i}(\cdot))) - u_i(\cdot, (y_i(\cdot), x_{-i}(\cdot)))] \mid \mathcal{F}_i](\omega)\, d\mu(\omega). \end{aligned}$$

Thus, (13) holds for every $y_i \in X_i$ and $C \in \mathcal{F}_i$ if and only if

$$\mathbb{E}[u_i(\cdot, (z_i(\cdot), x_{-i}(\cdot))) - u_i(\cdot, (y_i(\cdot), x_{-i}(\cdot)))] \mid \mathcal{F}_i](\omega) \ge 0 \qquad (14)$$

for every $y_i \in X_i$ and μ-a.e. $\omega \in \Omega$. It follows that the last condition on z_i characterizes the set $\arg\max_{y_i \in X_i} U_i(y_i, x_{-i})$. Similarly, $z_i \in \arg\max_{y_i \in X_i} E(p)(y_i, x_{-i})$ if and only if

$$\mathbb{E}[E(p)(\cdot, (z_i(\cdot), x_{-i}(\cdot))) - E(p)(\cdot, (y_i(\cdot), x_{-i}(\cdot)))] \mid \mathcal{F}_i](\omega) \ge 0 \qquad (15)$$

[29] For a definition see, e.g., Section 34 in Billingsley (1995).

for every $y_i \in X_i$ and μ-a.e. $\omega \in \Omega$.

Now observe that

$$\begin{aligned} & \mathbb{E}\left[u_i(\cdot,(z_i(\cdot),x_{-i}(\cdot))) - u_i(\cdot,(y_i(\cdot),x_{-i}(\cdot)))\right] \mid \mathcal{F}_i] \\ \text{(by (3))} \quad = \; & \mathbb{E}\left[w_i(\cdot)\left[E(p)(\cdot,(z_i(\cdot),x_{-i}(\cdot))) - E(p)(\cdot,(y_i(\cdot),x_{-i}(\cdot)))\right]\right] \mid \mathcal{F}_i] \\ = \; & w_i(\cdot)\,\mathbb{E}\left[p(\cdot,(y_i,x_{-i})(\cdot)) - p(\cdot,x(\cdot))\right] \mid \mathcal{F}_i], \end{aligned}$$

where the last equality is a consequence of $\mathcal{F}_i$-measurability of w_i (see, e.g., Theorem 34.3 in Billingsley (1995)). But $w_i(\cdot) > 0$, and so (14) holds if and only if (15) holds. This shows that $z_i \in \arg\max_{y_i \in X_i} U_i(y_i, x_{-i})$ if and only if $z_i \in \arg\max_{y_i \in X_i} E(p)(y_i, x_{-i})$, establishing the claimed equality. ■

Proof of the Theorem. We begin by recalling the notion of an L^1 space. In what follows, Σ will denote a σ-field on Ω that is equal to either $\mathcal{F}$ or $\mathcal{F}_i$ for some $i \in N$, and by a set-valued map $B : \Omega \to 2^{\mathbb{R}^m}$ we will mean either $A : \Omega \to 2^{\mathbb{R}^d}$ or $A_i : \Omega \to 2^{\mathbb{R}^{d_i}}$ for some $i \in N$.

The Banach space $L^1((\Omega,\Sigma,\mu);\mathbb{R}^m)$ consists of all (equivalence classes[30] of) $\mathbb{R}^m$-valued, Σ-measurable and μ-integrable functions on Ω, with the L^1-norm given by

$$\|x\|_{1,\mathbb{R}^m} = \int_\Omega \|x(\omega)\|_{\mathbb{R}^m}\, d\mu(\omega) \tag{16}$$

for every $x \in L^1((\Omega,\mathcal{F},\mu);\mathbb{R}^m)$, where $\| \; \|_{\mathbb{R}^m}$ denotes the Euclidean norm on $\mathbb{R}^m$. The topology that the L^1-norm induces on $L^1((\Omega,\Sigma,\mu);\mathbb{R}^m)$ is called *strong*. The *weak* topology on $L^1((\Omega,\Sigma,\mu);\mathbb{R}^m)$ is the minimal one in which, for every $y \in L^\infty((\Omega,\Sigma,\mu);\mathbb{R}^m)$($\equiv$the space of equivalence classes of all $\mathbb{R}^m$-valued, bounded and Σ-measurable functions on Ω), the linear functional $x \longmapsto \int_\Omega \langle x(\omega), y(\omega)\rangle\, d\mu(\omega)$ is continuous (where $\langle , \rangle$ denotes the scalar product on $\mathbb{R}^m$).

Given any $B : \Omega \to 2^{\mathbb{R}^m}$ as above, denote by $L^1((\Omega,\Sigma,\mu);B)$ the set of all $x \in L^1((\Omega,\Sigma,\mu);\mathbb{R}^m)$ with $x(\omega) \in B(\omega)$ for μ-almost every $\omega \in \Omega$. The strong (respectively, the weak) topology on $L^1((\Omega,\Sigma,\mu);\mathbb{R}^m)$ induces the strong (respectively, the weak) topology and on its subset $L^1((\Omega,\Sigma,\mu);B)$. Since B ($= A$ or A_i) has convex and compact values and is integrably bounded, $L^1((\Omega,\Sigma,\mu);B)$ is weakly compact by Corollary 2.5 of Balder and Yannelis (1993). In particular, $L^1((\Omega,\Sigma,\mu);B)$ is also weakly (and hence strongly) closed.

[30] The underlying equivalence relation identifies any two Σ-measurable functions that coincide μ-almost everywhere on Ω.

We will now apply the above to the issue at hand. Notice that X, the set of strategy profiles in G, can be naturally viewed as a convex subset of the weakly compact (and also strongly closed) $L^1((\Omega, \mathcal{F}, \mu); A)$.[31] We will first show that X is a strongly closed subset of $L^1((\Omega, \mathcal{F}, \mu); A)$. To this end, let $\{x^k\}_{k=1}^{\infty} \subset X$ be a strongly ($\| \|_{1,\mathbb{R}^d}$-)convergent sequence.[32] In particular, $\{x^k\}_{k=1}^{\infty}$ is a Cauchy sequence w.r.t. $\| \|_{1,\mathbb{R}^d}$.

For each $i \in N$ and $k \geq 1$, $x_i^k \in X_i$ represents an equivalence class in $L^1((\Omega, \mathcal{F}_i, \mu); A_i)$. Since $\|y_i\|_{1,\mathbb{R}^{d_i}} \leq \|y\|_{1,\mathbb{R}^d}$ for any $y \in L^1((\Omega, \mathcal{F}, \mu); \mathbb{R}^d)$ and its restriction y_i to (any) d_i coordinates, $\{x_i^k\}_{k=1}^{\infty} \subset X_i$ is a Cauchy sequence in $L^1((\Omega, \mathcal{F}_i, \mu); A_i)$ w.r.t. $\| \|_{1,\mathbb{R}^{d_i}}$. Being a Banach space, $L^1((\Omega, \mathcal{F}_i, \mu), \mathbb{R}^{d_i})$ is complete, and so is its strongly closed subset $L^1((\Omega, \mathcal{F}_i, \mu); A_i)$. Therefore, $\{x_i^k\}_{k=1}^{\infty}$ $\| \|_{1,\mathbb{R}^{d_i}}$-converges to a limit $x_i \in L^1((\Omega, \mathcal{F}_i, \mu); A_i)$ (and, being in $L^1((\Omega, \mathcal{F}_i, \mu); A_i)$, x_i belongs to X_i). It follows that $\{x^k\}_{k=1}^{\infty}$ converges to $x = (x_1, ..., x_n)$ in $\| \|_{1,\mathbb{R}^d}$,[33] and $x \in X$. This shows that X is a strongly closed subset of $L^1((\Omega, \mathcal{F}, \mu); A)$.

Due to Mazur's theorem on the equivalence between the strong and weak closedness of convex sets in a Banach space (see, e.g., Corollary 23 in Royden (1988)), a convex and strongly closed X is also weakly closed. It is, moreover, a subset of the weakly compact $L^1((\Omega, \mathcal{F}, \mu); A)$. Therefore, X is weakly compact.

The EP function $E(p)$ can be (well-)defined by (5) on the entire $L^1((\Omega, \mathcal{F}, \mu); A)$. Since the w-BP p satisfies conditions (a), (b) stated in Section 2.2, and (i), (ii) in the premise of the theorem, by Theorem 2.8 of Balder and Yannelis (1993) p's expectation $E(p)$ is weakly upper semi-continuous on $L^1((\Omega, \mathcal{F}, \mu); A)$. In particular, $E(p)$ is weakly upper semi-continuous on the weakly compact subset X of $L^1((\Omega, \mathcal{F}, \mu); A)$. As such, $E(p)$ attains its maximum on X at some $x \in X$. By the Corollary, that x is a BNE of G. ■

[31] This is done by identifying any $x = (x_1, ..., x_n) \in X$ with the function $\omega \longmapsto (x_1(\omega), ..., x_n(\omega))$, whose values belong to $A(\omega)$ for μ-a.e. ω.

[32] Now and henceforth, concrete functions will be used to represent the corresponding equivalence classes in L^1.

[33] This is because $\|x - x^k\|_{1,\mathbb{R}^d} \leq \sum_{i=1}^{n} \|x_i - x_i^k\|_{1,\mathbb{R}^{d_i}}$.

References

[1] Altman E., Hayel Y., Kameda H., 2007. Evolutionary dynamics and potential games in noncooperative routing. In: Proceedings of IEEE WiOpt, 1–5.

[2] Balder E.J., 1988. Generalized equilibrium results for games with incomplete information. Math. Oper. Res. 13, 265–276.

[3] Balder E.J., Yannelis N., 1993. On the continuity of expected utility. Econ. Theory 3, 625-643.

[4] Ballester C., Calvo-Armengol A., Zenou Y., 2002. Who's who in networks. Wanted: the key player. Econometrica 74, 1403–1417.

[5] Billingsley P., 1995. Probability and measure. John Willey&Sons, New York, NY.

[6] Bramoullé Y., Kranton R., D'Amours M., 2014. Strategic interaction and networks. Am. Econ. Rev. 104, 898–930.

[7] Carbonell-Nicolau O., McLean R.P., 2018. On the existence of Nash equilibrium in Bayesian games. Math. Oper. Res. 43, 100-129.

[8] Crémer J., 1990. Common knowledge and coordination of economic activities, in M. Aoki, B. Gustafsson, and O.E. Williams (Eds.), The Firm as a Nexus of Treaties, Sage Publishers, London, pp. 53-76.

[9] Diestel J., 1977. Remarks on weak compactness in $L_1(\mu, X)$. Glasgow Math. J. 18, 87-91.

[10] Einy E., Haimanko O., Moreno D., Shitovitz B., 2008. Uniform continuity of the value in zero-sum games with differential information. Math. Oper. Res. 33, 552-560.

[11] Einy E., Haimanko O., Moreno D., Shitovitz B., 2010. On the existence of Bayesian Cournot equilibrium. Game. Econ. Behav. 68, 77-94.

[12] Facchini G., van Megen, F., Borm, P., Tijs S., 1997. Congestion models and weighted Bayesian potential games. Theor. Decis. 42, 193–206.

[13] Gale D., Klee V., Rockafellar R.T., 1968. Convex functions on convex polytopes. P. Am. Math. Soc. 19, 867–873.

[14] Harsanyi J.C., 1967. Games of incomplete information played by Bayesian players, part I: The basic model. Manage. Sci. 14, 159–182.

[15] He W., Sun Y., 2019. Pure-strategy equilibria in Bayesian games. J. Econ. Theory 180, 11–49.

[16] He W., Yannelis N.C., 2016. Existence of equilibria in discontinuous Bayesian games. J. Econ. Theory 162, 181–194.

[17] Hellman Z., Levy Y.J., 2017. Bayesian games with a continuum of states. Theor. Econ. 12, 1089–1120.

[18] Mamer J, Schilling K., 1986. A zero-sum game with incomplete information and compact action spaces. Math. Oper. Res. 11, 627–631.

[19] Milgrom P., Weber R.J., 1986. Distributional strategies for games with incomplete information. Math. Oper. Res. 10, 619-632.

[20] Monderer D., Shapley S.H., 1996. Potential games. Game. Econ. Behav. 14, 124-143.

[21] Morris S., Shin S.H., 2002. Social value of private information. Am. Econ. Rev. 92, 1521-1534.

[22] Radner R., 1962. Team decision problems. Ann. Math. Stat. 33, 857-881.

[23] Radner R., Rosenthal, R.W., 1982. Private information and pure strategy equilibria. Math. Oper. Res. 7, 401–409.

[24] Raith, M., 1996. A general model of information sharing in oligopoly. J. Econ. Theory 71, 260-288.

[25] Royden, H.L., 1988. Real Analysis. Macmillan Publishing Company, New York, NY.

[26] Simon R.S., 2003. Games of incomplete information, ergodic theory, and the measurability of equilibria. Israel J. Math. 138, 73-92.

[27] Ui T., 2009. Bayesian potentials and information structures: team decision problems revisited. Int. J. Econ. Theory 5, 271–291.

[28] van Heumen R., Peleg B., Tijs S., Borm P., 1996. Axiomatic characterizations of solutions for Bayesian games. Theor. Decis. 40, 103–129.