EQUILIBRIUM EXISTENCE IN GAMES WITH A CONCAVE BAYESIAN POTENTIAL

Ezra Einy and Ori Haimanko

Discussion Paper No. 19-11

November 2019

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†

November 2019

Abstract

We establish existence of a pure-strategy Bayesian Nash equilibrium in Bayesian games with convex and compact action sets that have a continuous and concave potential at any state of nature. No assumptions are made on the information structure in these game; 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, Bayesian potential, equilibrium existence, concave payoffs, absolute continuity, information structures.

1 Introduction

The extensive use of Bayesian games in economic theory, particularly in the subfields of auctions and industrial organization, has been made possible by the fact that quite general categories of games with incomplete information possess a Bayesian Nash equilibrium (henceforth, BNE). For finite games with a discrete information

*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.

structure, the existence of BNE has been known since that concept's introduction in the fundamental work of Harsanyi (1967), but it was the work of Milgrom and Weber (1985) that proved BNE existence with remarkable generality: players' action and type sets were allowed to be (possibly uncountable) metric spaces, compact in the case of actions and separable and complete in the case of types.[1]

There are two conditions that are jointly required for Milgrom and Weber's BNE existence result. One is the continuity[2] of the players' payoff functions on the set of action profiles for any realization of the players' types; this condition has since been generalized in multiple works to allow for some modes of discontinuity (see, e.g., Carbonell-Nicolau and McLean (2018) and the literature survey in Section 5.1 therein). The other condition requires the joint distribution of the players' types to 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 of the 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"[3] or its variants,[4] while focusing on a relaxation of the payoff continuity assumption. 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

[1] All topological assumptions on the type sets were subsequently removed in the BNE existence result of Balder (1988).

[2] While Milgrom and Weber (1985) originally required equicontinuity of payoffs, this was reduced to continuity in Balder (1988).

[3] That is how we will refer to the second condition of Milgrom and Weber (1985) from now on. An alternative term for absolutely continuous information is "diffuse information," as used by Stinchcombe (2011).

[4] The reader is again referred to the literature survey in Carbonell-Nicolau and McLean (2018). The majority of relevant works are in the Harsanyi types setting, for which the absolute continuity condition was stated originally. In those papers that adopt the state-space setting (such as Yannelis and Rustichini (1991) and He and Yannelis (2016)), the absolute continuity condition needs to be replaced because it only applies to a distribution of types. To remain within the absolutely continuous information paradigm, Yannelis and Rustichini (1991) consider countable partitions of the space of states of nature, while He and Yannelis (2016) assume that the space is countable.

(2003) and Hellman (2014).

An interesting case of non-absolutely continuous information is obtained, for instance, when the types are a non-atomic continuum, but the information revealed by each type is purely atomic, that is, given his type, each player knows with certainty that the others' types belong to a finite or countable set.[5] Hellman and Levy (2017) characterize the exact circumstances in which a BNE exists for such "purely atomic" Bayesian games with finitely many actions. Thus, a BNE may exist, but not necessarily so, in a general Bayesian game with a payoff function that is continuous at any state of nature. A natural question that arises is whether there are interesting classes of games which, if played at every state of nature, would guarantee existence of a BNE in the corresponding Bayesian game for *all* varieties of information structures (and, in particular, for non-absolutely continuous ones).

It has long been known that two-player zero-sum games are one such class. Mamer and Schilling (1986) and Einy et al. (2008) have shown that when a game is zero-sum at every state of nature[6] and the payoff function is continuous at each action separately at every state, then a BNE exists for general information structures, without any need for absolute continuity of information. In this work, we will present another class of games for which the fact that they are played at each state will imply BNE existence for any information structure.

The popular concept of a potential game, introduced in Monderer and Shapley (1996), quickly found its way into the incomplete information paradigm: Heumen et al. (1996) defined a Bayesian potential game as one in which, at every state of nature, a potential game is played;[7] Ui (2009) extended that definition to more general (infinite) information structures and action sets. A Bayesian potential is then

[5]Stinchcombe (2011) generically characterized the non-absolutely continuous information structures as those in which there exists a non-null event, "perhaps not in any player's information set, conditional on which two or more players can infer the value of some continuously distributed random variable" (see Stinchcombe 2011, p. 657).

[6]In what follows, we will mostly use the language of the "states of nature" model of incomplete information that supersedes a simpler model of Harsanyi types. In the context of the latter model, "state" will refer to any realization of the players' types.

[7]That is, at every state of nature, the payoff differences of a player that are brought about by his unilateral deviations are precisely mimicked by a fictitious payoff function (the potential) that is common to all players.

defined as a state-dependent potential function for the state games. It has been well understood that any maximizer of the expectation of a Bayesian potential for such a game over the set of all *pure* Bayesian strategy profiles is a BNE of the Bayesian potential game. However, outside a rather specific context that will be discussed later, there has been no attempt to establish existence of a BNE in Bayesian potential games with a general (possibly non-absolutely continuous) information structure by first proving the existence of a maximizer for the expected Bayesian potential.

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 that make the set of pure Bayesian strategies of each player compact are in general too weak to guarantee continuity of the expected Bayesian potential.[8] However, it turns out that, with action sets being convex subsets of a Euclidean space, the topological tension between continuity and compactness does not arise when a Bayesian potential is not only continuous but also *concave* in all states of nature. The proof is based on the result of Balder and Yannelis (1993), who have shown – in the context of expected utilities defined for a variety of sets of contingent consumption plans – that for concave and continuous state-utilities the expected utility is *upper semi-continuous* in the *weak topology* on the set of contingent plans. Here we will show that, with some measurability-related fixes, the latter result can be applied to Bayesian games whose Bayesian potential is continuous and concave at each state, thereby establishing[9] (weak) upper semi-continuity of the expectation of the Bayesian potential on the set of pure Bayesian strategy profiles, which is (weakly) compact. This implies the existence of a Bayesian potential maximizer, and hence of a BNE in pure strategies. 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

[8] With actions in a Euclidean space $\mathbb{R}^m$, (pure) Bayesian strategies can be viewed as bounded or, more generally, integrable $\mathbb{R}^m$-valued functions of the state of nature, and hence as elements of a corresponding L^1 space. In the *strong*, or norm, topology induced by L^1 on the Bayesian strategy sets, the expected payoffs are continuous (assuming, e.g., that the payoffs are continuous at each state, and also integrably bounded), but that is not the case in the *weak* topology on L^1 (see, e.g., Example 2 in Milgrom and Weber (1985)). However, it is the weak, and not the strong, topology in which the Bayesian strategy sets tend to be compact (see, e.g., Einy et al. (2008), and, more generally, Corollary 2.5 of Balder and Yannelis (1993)).

[9] This requires an additional, mild, integrability assumption on the Bayesian potential.

in that respect.[10]

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 (belonging to the real line). 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;[11] the strategy of each player in the unique maximizer was shown to be linear in the player's signal. The games of Radner (1962) were found to be very 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 potentials for linear oligopoly games 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, that includes linear oligopolies alongside other well-known payoff specifications.[12]

Our result on BNE existence applies in the above-mentioned contexts when players' actions are restricted to be compact intervals,[13] because, for those Bayesian

[10] There is a considerable strand of literature on "equilibrium purification," studying existence of *pure strategy* BNE in Bayesian games, that began with Radner and Rosenthal's (1982) work on finite-action games with independent non-atomic types and private values. Although Radner and Rosenthal's framework and results have been significantly extended in several directions (see, e.g., He and Sun (2019), Khan and Zhang (2014) and the references therein), the assumptions on the information structure that are needed for equilibrium purification remain quite restrictive, and in particular leave out non-absolutely continuous information structures and type sets with atoms.

[11] When all parameters of the quadratic payoff function are uncertain and have a general distribution, Radner offered a sufficient condition for the maximum existence, that implicitly links together the information structure in the game and the distribution of its parameters.

[12] These specifications include variants of games considered in Crémer (1990) and Morris and Shin (2002).

[13] When players' actions belong to $\mathbb{R}$, Theorem 5 in Radner (1962) has already established the existence of a maximum for the expectation of Radner's function (with the only uncertainty being

games, Radner's payoff function constitutes a Bayesian potential that is continuous and concave at each state. But our result also extends the scope of what has already been shown for those Bayesian games with quadratic payoffs 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 at all. 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, because a concave Bayesian potential would still be easily constructible for such games.[14]

The paper is organized as follows. In Section 2 we describe the general set-up and recall the notions of a Bayesian game, Bayesian potential and BNE. Section 3 contains our BNE existence result and remarks on its possible extensions. Applications are discussed in Section 4.

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 discernible by i.

Each player $i \in N$ has a set A_i of actions, which is a convex and compact subset of

on the linear term's parameters, and with a joint normal distribution of those parameters and the players' signals). Notice, however, that this does not imply existence of a maximum when action sets are taken to be some compact intervals, unless the pure Bayesian strategy profile that is the Radner's maximizer happens to have values in the restricted action sets at every state of nature.

[14]Such a potential would not, however, necessarily have a quadratic form, and hence won't be amenable to Radner's (1962) analysis.

$\mathbb{R}^{d_i}$ for some $d_i \in \mathbb{N}$.[15] The product set $A = A_1 \times ... \times A_n \subset \mathbb{R}^d$ (where $d := \sum_{i=1}^n d_i$) is thus also convex and compact. Each $i \in N$ has a payoff function $u_i : \Omega \times A \to \mathbb{R}$. We will assume that u_i is $\mathcal{F} \otimes \mathcal{B}(A)$-measurable[16] and integrably bounded (i.e., $\sup_{a \in A} |u_i(\cdot, a)|$ is μ-integrable). A 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 $x_i : \Omega \to A^i$. The set of all strategies of player i will be denoted by X_i. 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 $x = (x_1, ..., x_n) \in X$. 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.

2.2 Potential games

In our definition of a Bayesian potential we extend the original notion of van Heumen et al. (1996) that was stated for finite information structures (and appropriately adapt the definition of Ui (2009) for general structures). We say that $G = (N, (\Omega, \mathcal{F}, \mu), (\mathcal{F}_i, A_i, u_i)_{i=1}^n)$ a *Bayesian potential game* if there exists an $\mathcal{F} \otimes \mathcal{B}(A)$-measurable and integrably bounded function $p : \Omega \times A \to \mathbb{R}$ that satisfies the following: for μ-almost every $\omega \in \Omega$, every $i \in N$, and every $a \in A$, $b_i \in A_i$,

$$u_i(\omega, (b_i, a_{-i})) - u_i(\omega, a) = p(\omega, (b_i, a_{-i})) - p(\omega, a) \tag{2}$$

(where $(b_i, a_{-i}) \in A$ is the action profile obtained from a by substituting b_i for a_i). Such p will be called a *Bayesian potential* for G.[17]

[15] See Remark 1 for a generalization that allows the action sets to depend on the state of nature.

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

[17] See Remark 2 for more general notions of Bayesian potentials.

If p is a Bayesian potential for G then the *expected potential (function)* $P : X \to \mathbb{R}$, given by

$$P(x) = \int_\Omega p(\omega, x(\omega)) d\mu(\omega) \tag{3}$$

for any $x = (x_1, ..., x_n) \in X$, obviously retains the property expressed in (2) in that it precisely mimics all unilateral deviations of each player in terms of his expected payoff. That is,

$$U_i(y_i, x_{-i}) - U_i(x) = P(y_i, x_{-i}) - P(x) \tag{4}$$

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

3 BNE Existence

The existence of a Bayesian potential 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. If $G = \left(N, (\Omega, \mathcal{F}, \mu), (\mathcal{F}_i, A_i, u_i)_{i=1}^n\right)$ has a Bayesian potential p, and $p(\omega, \cdot)$ is a concave and upper semi-continuous[18] function on A for μ-almost every $\omega \in \Omega$, then G possesses a (pure-strategy) BNE.

Proof. 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 $B \subset \mathbb{R}^m$ we will mean either $A \subset \mathbb{R}^d$ or $A_i \subset \mathbb{R}^{d_i}$ for some $i \in N$.

The Banach space $L^1\left((\Omega, \Sigma, \mu); \mathbb{R}^m\right)$ consists of all (equivalence classes[19] of) $\mathbb{R}^m$-valued, Σ-measurable and μ-integrable functions on Ω, with the L^1-norm given by

$$\|x\|_{1,\mathbb{R}^m} = \int_\Omega \|x(\omega)\|_{\mathbb{R}^m} d\mu(\omega) \tag{5}$$

[18]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 is a polytope then we, in effect, assume that $p(\omega, \cdot)$ is continuous.

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

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 $\mathbb{R}^m$).

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 $L^1((\Omega, \Sigma, \mu); B)$ that consists of Σ-measurable and μ-almost everywhere B-valued functions, where $B \subset \mathbb{R}^m$ is a given convex and compact set. Corollary 2.5 of Balder and Yannelis (1993) implies that $L^1((\Omega, \Sigma, \mu); B)$ is weakly compact. In particular, it is also weakly (and hence strongly) closed.

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) $L^1((\Omega, \mathcal{F}, \mu); A)$.[20] 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.[21] 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 $y \in L^1\left((\Omega, \mathcal{F}, \mu); \mathbb{R}^d\right)$ 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, $L^1\left((\Omega, \mathcal{F}_i, \mu), \mathbb{R}^{d_i}\right)$ 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 $x_i \in L^1((\Omega, \mathcal{F}_i, \mu); A_i)$. Moreover, since x_i is A_i-valued modulo an $\mathcal{F}_i$-measurable function that vanishes μ-almost everywhere, it can be assumed that x_i is, in fact, A_i-valued (and $\mathcal{F}_i$-measurable). In other words, $x_i \in X_i$ for each $i \in N$, and thus $x = (x_1, ..., x_n) \in X$. But, clearly, $\{x^k\}_{k=1}^\infty$ converges to $x \in X$ in $\| \|_{1,\mathbb{R}^d}$.[22] This

[20] This is done by identifying any $x = (x_1, ..., x_n) \in X$ with the A-valued function $\omega \longmapsto (x_1(\omega), ..., x_n(\omega))$, modulo the set of $\mathcal{F}$-measurable functions that differ from it on a null set of the measure μ.

[21] Now and henceforth, concrete functions will be used to represent the corresponding equivalence classes.

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

shows that X is a strongly closed subset of $L^1((\Omega, \mathcal{F}, \mu); A)$.

Due to a well-known 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 expected potential function P can be (well-)defined by (3) on the entire $L^1((\Omega, \mathcal{F}, \mu); A)$.[23] Since the Bayesian potential p is integrably bounded by assumption, its concavity and continuity properties stated in the premise of the theorem allow an appeal to Theorem 2.8 of Balder and Yannelis (1993), which asserts weak upper semi-continuity of the expectation P on $L^1((\Omega, \mathcal{F}, \mu); A)$. In particular, P is weakly upper semi-continuous on the weakly compact subset X of $L^1((\Omega, \mathcal{F}, \mu); A)$. As such, P attains its supremum on X at some $x \in X$. Clearly, x remains a maximizer of P also when genuine strategy profiles in X are considered (instead of their equivalence classes). A standard argument, based on (4), now establishes that x is a BNE of G. ■

Remark 1 (Extension of the theorem to games with state-dependent action sets). Given a space of states of nature $(\Omega, \mathcal{F}, \mu)$ and information fields $(\mathcal{F}_i)_{i=1}^n$ for players in the set N, consider a generalized concept of a Bayesian game, in which the action set of each player i at any state $\omega \in \Omega$ is a convex and compact $A_i(\omega) \subset \mathbb{R}^{d_i}$; that is, the action set may depend on the state of nature. Let us now denote by A_i and A the corresponding set-valued functions $A_i : \Omega \to 2^{\mathbb{R}^{d_i}}$ and $A : \Omega \to 2^{\mathbb{R}^d}$ (where the latter is given by $A(\omega) = A_1(\omega) \times \ldots \times A_n(\omega)$ for every $\omega \in \Omega$), further assuming that the graph of A (respectively, A_i) is $\mathcal{F} \otimes \mathcal{B}(\mathbb{R}^d)$-measurable (respectively, $\mathcal{F}_i \otimes \mathcal{B}(\mathbb{R}^{d_i})$-measurable) and that $\sup_{a \in A(\omega)} \|a\|_{\mathbb{R}^d}$ is μ-integrable. A Bayesian strategy of player i is then an $\mathcal{F}_i$-measurable function $x_i : \Omega \to \mathbb{R}^{d_i}$ with the property that $x_i(\omega) \in A_i(\omega)$ for μ-almost every $\omega \in \Omega$; in order for the expected payoffs $(U_i)_{i=1}^n$ to be well-defined by (1), we also assume that each payoff function $u_i : \Omega \times \mathbb{R}^d \to \mathbb{R}$ is $\mathcal{F} \otimes \mathcal{B}(\mathbb{R}^d)$-measurable and satisfies $|u_i(\omega, a)| \leq \psi_i(\omega) + M_i \|a\|_{\mathbb{R}^d}$ for every $\omega \in \Omega$ and $a \in \mathbb{R}^d$, where $\psi_i : \Omega \to \mathbb{R}_+$ is some μ-integrable function and

[23] Indeed, extend p arbitrarily from $\Omega \times A$ to $\Omega \times \mathbb{R}^d$ (preserving measurability). Then, for any $x, y \in L^1((\Omega, \mathcal{F}, \mu); A)$ that are identical μ-almost everywhere, defining $P(x)$ and $P(y)$ by (3) produces equal expressions.

$M_i > 0$ is a constant. Assumptions of an identical nature will be made on a Bayesian potential function $p : \Omega \times \mathbb{R}^d \to \mathbb{R}$.

The statement of our theorem and its proof will remain in force in this generalized setting. The proof remains valid because, for any $B : \Omega \to 2^{\mathbb{R}^m}$ with a $\Sigma \otimes \mathcal{B}(\mathbb{R}^m)$-measurable graph, convex and compact values, and μ-integrable $\sup_{b \in B(\omega)} \|b\|_{\mathbb{R}^m}$, the set $L^1((\Omega, \Sigma, \mu); B)$ [of $x \in L^1((\Omega, \Sigma, \mu); \mathbb{R}^m)$ with $x(\omega) \in B(\omega)$ for μ-almost every $\omega \in \Omega$] is still weakly compact, by Corollary 2.5 of Balder and Yannelis (1993).[24] Similarly, by Theorem 2.8 of Balder and Yannelis (1993), the expected potential function P will still be weakly upper semi-continuous on $L^1((\Omega, \mathcal{F}, \mu); A)$.

Remark 2 (Extension of the theorem to generalized notions of potential). The notion of Bayesian potential does not depend on the private information in the game, given by the fields $(\mathcal{F}_i)_{i=1}^n$. This fact makes our theorem that much stronger: the existence of a Bayesian potential with the attributes required by the theorem guarantees existence of a BNE in a way that is robust in regard to specific details of the private information endowments. A somewhat weaker notion of a *weighted* Bayesian potential (also due to van Heumen et al. (1996)) would have preserved the existence claim and its robustness: $p : \Omega \times A \to \mathbb{R}$ is such a potential if there exist a vector of positive weights $(w_i)_{i=1}^n$ such that, for μ-almost every $\omega \in \Omega$, every $i \in N$, and every $a \in A$, $b_i \in A_i$,

$$u_i(\omega, (b_i, a_{-i})) - u_i(\omega, a) = w_i \left[p(\omega, (b_i, a_{-i})) - p(\omega, a) \right]. \tag{6}$$

If one is ready to dispense with the requirement that BNE existence be independent of the specific $(\mathcal{F}_i)_{i=1}^n$, then more general notions of potential may be used in our theorem. For instance, one may allow each w_i in (6) to be a strictly positive and bounded $\mathcal{F}_i$-measurable function. Even more generally, one may consider *Bayesian best-response potentials* defined in Ui (2009), whose expectation precisely mimics the best responses of players to all Bayesian strategy profiles of others.[25]

[24] The sets $L^1((\Omega, \mathcal{F}, \mu); A)$ and $L^1((\Omega, \mathcal{F}_i, \mu); A_i)$ that are used in the original proof will need to be redefined following this principle, by treating A and A_i as the set-valued functions on Ω and not as fixed sets.

[25] See Theorem 5 in Ui (2009) for an example of a condition generalizing (6) that makes p a Bayesian best-response potential.

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.

4.1 Motivating model: oligopoly with linear demand

Cournot oligopoly is a showcase of the Bayesian potential usefulness. 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 a Bayesian potential for our incomplete information oligopoly:

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

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.

As a particular case, when $\mathbf{A}_i = \mathbf{A}$ for all i and $\varepsilon = \delta$, we obtain Cournot oligopoly with a single homogeneous good, which (in a 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 Bayesian potential game, with the ensuing claim of BNE existence.[26]

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[27] that views those common payoffs as Bayesian potentials 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:

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

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

[27] 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.

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 integrably bounded, and $h_i : \Omega \times A_{-i} \to \mathbb{R}$ is $\mathcal{F} \otimes \mathcal{B}(A_{-i})$-measurable and integrably bounded.[28] It is easy to see that the game has a Bayesian potential, 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 Bayesian potential 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,[29] 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

[28] 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.

[29] 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 $\int_\Omega \sum_{i=1}^{n}\sum_{j=1}^{n} q_{ij}(\omega)(x_i(\omega) - z_i(\omega))(x_j(\omega) - z_j(\omega))\, d\mu(\omega) < \infty$.

model in Ballester et al (2006)), in which player i's payoff is

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

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 Bayesian potential game, with a Bayesian potential p that is given by

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

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[30] $[\delta_{ij} + \delta(\omega)\, g_{ij}(\omega)]_{n\times n}$ 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 $\theta(\omega)$,[31] 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 Bayesian potential, one may use the function given by

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

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

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

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.

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[32] $\theta(\omega)$ 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 Bayesian potential – 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,[33] 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,

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

[32] See the previous footnote.

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

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

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

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 $(d_l)_{l\in L}$.

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

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

[7] 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.

[8] 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.

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

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

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

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

[13] Hellman Z., 2014. A game with no Bayesian approximate equilibrium. J. Econ. Theory 153, 138–151.

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

[15] Khan M.A., Zhang Y., 2014. On the existence of pure-strategy equilibria in games with private information: a complete characterization. J. Math. Econ. 50, 197–202.

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

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

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

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

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

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

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

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

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

[25] Stinchcombe M.B., 2011. Balance and discontinuities in infinite games with type-dependent strategies. J. Econ. Theory 146, 656–671.

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

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