Pure-strategy Equilibrium in Bayesian Potential Games with Absolutely Continuous Information

Ezra Einy*and Ori Haimanko†

September 2022

Abstract

We prove the existence of a pure-strategy Bayesian Nash equilibrium in Bayesian games with absolutely continuous information and a Bayesian potential that is upper semi-continuous in actions for any realization of the players' types. In particular, all finite Bayesian potential games with absolutely continuous information possess a pure-strategy Bayesian Nash equilibrium.

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

Key words: Bayesian games, Bayesian potential, pure-strategy equilibrium, continuous payoffs, absolute continuity of information, purification.

1 Introduction

Potential games, formally introduced by Monderer and Shapley (1996) but brought to attention already by Rosenthal (1973), are used in applications in very diverse fields. From oligopoly theory to team decisions to congestion and network problems, important insights have been gained due to the fact that potential games naturally arise in the corresponding contexts (to mention just a sample, see, e.g., Radner (1962), Raith (1996), Fabrikant et al. (2004), Ui (2009)). Having a potential is, in particular, a significant facilitating factor in establishing the most basic and desirable feature of

*Department of Economics, Ben-Gurion University of the Negev, Beer Sheva, Israel. E-mail: einy@bgu.ac.il.

†Department of Economics, Ben-Gurion University of the Negev, Beer Sheva, Israel. E-mail: orih@bgu.ac.il.

a game, which is the existence of a pure-strategy equilibrium: finding an equilibrium may be accomplished by maximizing the potential, as if the latter were the common payoff of all players, because any potential maximizer is an equilibrium of the game.

The existence of a pure-strategy equilibrium in all finite[1] potential games is therefore self-evident. However, when incomplete information is allowed the situation is somewhat less clear. For a finite Bayesian game, its Bayesian potential was defined by van Heumen et al. (1996) (with a subsequent extension by Ui (2009) and Einy and Haimanko (2020) for infinite information structures) as a common surrogate payoff function that mimics all marginal changes in payoffs produced by the players unilateral deviations for any realization of uncertainty. The ex-ante evaluation of the Bayesian potential, namely its expectation, acts as a potential of the game's normal form: any pure Bayesian strategy profile that maximizes the expected potential is a (Bayesian Nash) equilibrium of the game. When the game's information structure is discrete, the expected potential clearly has a maximum in pure Bayesian strategies, and hence the game has a pure-strategy equilibrium.[2] But when the players have uncountable type sets, the existence of an expected potential maximizer is in question because the topologies in which the set of pure Bayesian strategy profiles may be compact tend to be too weak to ensure that the expected potential is continuous on that set.[3]

The problem underlined above can sometimes be avoided if the action sets of players are infinite but *convex* and *compact*. That is for instance the case in many Bayesian potential games with semi-quadratic payoffs that appear in works summarized and extended in Raith (1996), Ui (2009) and Einy and Haimanko (2020).[4] These authors mainly focus on such games under the assumption that their Bayesian potential is *concave* in actions, and there is a strong underlying reason for that. Indeed, as Einy and Haimanko (2020) have shown, the concavity of a Bayesian potential (coupled with the latter's ex-post continuity in actions) is a sufficient condition for a pure-strategy equilibrium existence in general games, precisely because that con-

[1] By *finite* games we refer to games with a finite action set.

[2] This was observed by Heumen et al. (1996).

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

[4] Manifesting the extent of a Bayesian potential applicability, these games arise *inter alia* in the context of oligopolies with linear demand, team decision problems, coordination problems and networks.

dition makes the expected potential upper semi-continuous in the *weak* topology on the set of pure Bayesian strategy profiles, which is *weakly* compact, and hence the expected potential attains its maximum in pure strategies.

Here we consider the question of pure-strategy equilibrium existence in Bayesian potential games where the action sets are not necessarily convex (and the information structure not necessarily discrete). In particular, the question applies to *finite* Bayesian potential games, examples of which abound. Standard finite congestion games with incomplete information naturally come to mind because (complete-information) congestion games are the best-known instances of a potential game.[5] But many other finite Bayesian potential games have come under specific scrutiny, such as, e.g., the public good provision game of Palfrey and Rosenthal,[6] and the investment and regime change games of Morris et al. (2022). Yet another source of finite Bayesian potential games can be obtained by restricting the games with convex action space mentioned in the preceding paragraph to have finitely many feasible actions. Examples of such games would be, e.g., linear-demand oligopolies considered in Einy and Haimanko (2020)[7] where quantities/prices are restricted to discrete units/scale.

We will show that any Bayesian game with a Bayesian potential that is ex-post upper semi-continuous in actions possesses a pure-strategy equilibrium, and hence, in particular, any finite Bayesian potential game has such an equilibrium. Unlike in the existence result of Einy and Haimanko (2020) for concave Bayesian potential games, however, here we impose the condition of absolute continuity of information, common to nearly all results on Bayesian Nash equilibrium existence.[8] Introduced by Milgrom and Weber (1986), absolute continuity of information requires the joint distribution of the players' types to be absolutely continuous with respect to the product of its marginal distributions. This condition is useful in applications because it holds, e.g., when the players' types are independent, or have a joint density.

Our method of proof deviates from the standard one because, as mentioned previously, compactifying the set of pure Bayesian strategy profiles by endowing it with the

[5] As shown by Monderer and Shapley (1996), any finite potential game is in fact isomorphic to a congestion game.

[6] See the rendering of that game on pp. 211-213 in Fudenberg and Tirole (1991).

[7] See Example 1 in Section 4 therein.

[8] See, e.g., Carbonell-Nicolau and McLean (2018) for a survey of results on equilibrium existence in Bayesian games.

weak topology may leave the expected potential discontinuous. Instead, the game is extended into behavioral Bayesian strategies, on whose product set the narrow topology of Balder (1988) leads both to that set's compactness and to upper-continuity of the expected potential. That ensures the existence of an expected potential's maximizer in behavioral strategies, which can then be purified to produce a pure-strategy maximizer. The latter is a desired pure-strategy equilibrium.

In the literature on pure-strategy equilibrium existence in general Bayesian games,[9] the recent result of He and Sun (2019) invites comparison to ours because those authors also allow uncountable type spaces, do not use structural assumptions[10] on types, and purify a behavioral-strategy equilibrium in order to obtain a pure-strategy one. They rely, however, on the type spaces being atomless, and on a further "coarser inter-player information" condition that requires a modicum of conditional freedom in each player's type. It is due to these assumptions that a sophisticated purification procedure,[11] whereby a given behavioral strategy is replaced by a pure one without affecting any of the players' expected payoffs in all associated strategy profiles, can be performed. In contrast, in our set-up there is no assumption of atomlessness.[12] Also, our purification claim concerning an expected potential's maximizer is more straightforward, requiring no appeal to classical purification results. Specifically, on account of its multi-linearity in behavioral strategies, the expected potential is first shown to have a maximizer whose components are extreme points of convex sets of strategies; these extreme points are then shown to be pure strategies.

The paper is organized as follows. The framework is presented in sections 2 and 3. Section 4 states the main result, which is proved in Section 5. Section 6 concludes.

[9] See, e.g., He and Sun (2019, pp. 14–15) for a brief summary.

[10] Such as, e.g., the assumption that the type sets are lattices, as in Van Zandt and Vives (2007) or Reiny (2011).

[11] This procedure in particular extends the purification results in Khan et al. (2006), which are in turn based on the classical purification introduced in Dvoretsky et al. (1951).

[12] Atomlessness of type spaces is well known to be *necessary* for the standard purification procedures in Bayesian games, so it is of particular note that it is not required for pure-strategy equilibrium existence in Bayesian potential games. Furthermore, our result allows infinite action sets, unlike most of purification literature.

2 Bayesian games

A Bayesian game G considered in this work has a finite player set $N = \{1, ..., n\}$. The information endowment of each player $i \in N$ is given by a standard Borel (measurable) type space $(T_i, \mathcal{T}_i)$. The players are assumed to have a common prior μ on the product space $(T, \mathcal{T}) := (\times_{i \in N} T_i, \otimes_{i \in N} \mathcal{T}_i)$ of all type profiles. The Milgrom and Weber (1986) condition of *absolute continuity of information* holds: μ is postulated to be absolutely continuous w.r.t. the product of its marginals, $\otimes_{i \in N} \mu_i$.

Each player $i \in N$ has a set A_i of actions that is a metrizable compact space. The compact product set $A = A_1 \times ... \times A_n$ is comprised of all action profiles. Each $i \in N$ has an *ex-post* payoff function $u_i : T \times A \to \mathbb{R}$. We assume that u_i is $\mathcal{T} \otimes \mathcal{B}(A)$-measurable[13] and that $\sup_{a \in A} |u_i(\cdot, a)|$ is μ-integrable.

A pure (Bayesian) strategy of player $i \in N$ in the game G is a $\mathcal{T}_i$-measurable function $s_i : T_i \to A_i$; that is, upon learning his type t_i, player i chooses an action $s_i(t_i) \in A_i$. The set of all strategies of player i is denoted by S_i, with the product set $S = \times_{i=1}^n S_i$ comprising all strategy profiles. Each player i evaluates his *ex-ante* prospect in the game via the expected payoff function U_i on S, given by

$$U_i(s) = \int_T u_i \left(t, (s_1(t_1), ..., s_n(t_n))\right) \mu(dt),$$

for any $s = (s_i)_{i \in N} \in S$. As usual, $s \in S$ is a pure-strategy Bayesian Nash equilibrium of the game G, or PS-BNE for short, if it is a Nash equilibrium of the normal form of G, namely, if the inequality

$$U_i(s) \geq U_i(r_i, s_{-i})$$

holds for every $i \in N$ and $r_i \in S_i$, where $(r_i, s_{-i}) \in S$ denotes the strategy profile obtained from s by substituting r_i for s_i.

3 Bayesian potential and expected potential

As in Einy and Haimanko (2020), who extend the corresponding notion of van Heumen et al. (1996) and follow Ui (2009), G is said to be a *Bayesian potential game* (or BP game, for short) if there exists $\mathbf{p} : T \times A \to \mathbb{R}$ (called a *Bayesian potential*, or BP, for G) that satisfies the following:

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

(a) $\mathbf{p}$ is $\mathcal{T} \otimes \mathcal{B}(A)$-measurable;

(b) $\sup_{a \in A} |\mathbf{p}(\cdot, a)|$ is μ-integrable;

and

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

$$u_i(t, (b_i, a_{-i})) - u_i(t, a) = \mathbf{p}(t, (b_i, a_{-i})) - \mathbf{p}(t, a) \tag{1}$$

(where $(b_i, a_{-i}) \in A$ is the action profile obtained from a by substituting b_i for a_i). Thus, the marginal change in the ex-post payoff resulting from a unilateral deviation by any player is precisely reflected by marginal change in the BP.

If G has a BP $\mathbf{p}$, consider the corresponding *expected potential* (or EP for short), $E(\mathbf{p}) : S \to \mathbb{R}$, given by

$$E(\mathbf{p})(s) = \int_T \mathbf{p}\left(t, (s_1(t_1), ..., s_n(t_n))\right) \mu\left(dt\right) \tag{2}$$

for any $s = (s_i)_{i \in N} \in S$. The EP obviously retains the property expressed in (1), now given in terms of the expected payoffs:

$$U_i(r_i, s_{-i}) - U_i(s) = E(\mathbf{p})(r_i, s_{-i}) - E(\mathbf{p})(s) \tag{3}$$

for every $i \in N$ and every $s \in S$, $r_i \in S_i$. Thus, if G is a BP game then it is a normal-form potential game (in the sense of Monderer and Shapley (1996)), and, clearly, any maximizer $s \in S$ of its normal-form potential $E(\mathbf{p})$ is a PS-BNE of G.

4 PS-BNE existence

The straightforward way of establishing PS-BNE existence in a BP game G is by showing that the EP attains a maximum on S. That was done, e.g., by van Heumen et al. (1996) for Bayesian games with finite information structures and by Einy and Haimanko (2020) for games possessing a BP that is concave and (ex-post) upper semi-continuous on A. In general, however, the natural topology in which the strategy profile set S is *compact* – with each strategy set S_i having the weak topology – is *too weak to obtain continuity* of the EP, and hence to imply existence of an EP maximizer even if the BP is ex-post continuous.

It turns out that the existence of a PS-BNE in any Bayesian game with an ex-post upper semi-continuous BP can still be established:

Theorem 1 *If G has a BP $\mathbf{p}$ such that $\mathbf{p}(t,\cdot)$ is upper semi-continuous on A for every $t \in T$, then G possesses a PS-BNE.*

The proof of Theorem 1, given in the next section, proceeds in an indirect manner. First, G needs to be extended into *behavioral strategies* because the behavioral strategy sets can then be equipped with a topology that resolves the previously mentioned tension between compactness of the strategy sets and continuity of the expected payoff functions. Indeed, Balder (1988) showed that in the *narrow topology*[14] the behavioral strategy sets are compact and the expected payoffs are continuous (provided the ex-post payoffs are continuous in actions and the information is absolutely continuous); furthermore, the continuity in that claim can be replaced by upper semi-continuity. The existence of an EP maximizer in *behavioral strategies* is then a corollary. But, with the EP being affine in each strategy separately, the components of its behavioral-strategy maximizer can be sequentially "purified" (i.e., replaced by pure strategies in such a way that the resulting strategy profiles remain EP maximizers), thus producing a genuine PS-BNE.

Also notice that if G is *finite* (i.e., all action sets A_i are finite) and has a BP, then its BP is trivially ex-post continuous. Thus, we have the following immediate corollary:

Corollary 2 *If G is a finite BP game then it possesses a PS-BNE.*

5 Proof of Theorem 1

5.1 Part 1: Existence of an EP maximizer in behavioral strategies

We begin by extending the game into behavioral strategies. Formally, a *behavioral strategy* of $i \in N$ in G is a mapping $\sigma_i : T_i \times \mathcal{B}(A_i) \to [0,1]$, such that $\sigma_i(t_i,\cdot)$ is a probability measure on A_i for every $t_i \in T_i$ and $\sigma_i(\cdot,B)$ is $\mathcal{T}_i$-measurable for every $B \in \mathcal{B}(A_i)$. Let Σ_i denote the set of i's behavioral strategies; the product set $\Sigma = \times_{i=1}^{n}\Sigma_i$ consists of strategy profiles. Any pure strategy $s_i \in S_i$ is clearly identifiable with a behavioral strategy in Σ_i in which, for every $t_i \in T_i$, $\sigma_i(t_i,\cdot)$ is the

[14] We borrow this term from Carbonell-Nicolau and McLean (2018).

Dirac measure concentrated on $s_i(t_i)$; henceforth, such an identification will be made whenever convenient, and the symbol s_i will be used both for a pure strategy in S_i and for its behavioral form in Σ_i.

For any strategy profile $\sigma = (\sigma_i)_{i\in N} \in \Sigma$, the expected payoff of player $i \in N$ is given by

$$U_i(\sigma) = \int_T \int_A u_i\,(t, a)\,\sigma_1(t_1, da_1)...\sigma_n(t_n, da_n)\mu(dt).$$

In the same fashion, the expected potential $E(\mathbf{p})$ extends into a function $E(\mathbf{p}) : \Sigma \to \mathbb{R}$ given by

$$E(\mathbf{p})(\sigma) = \int_T \int_A \mathbf{p}\,(t, a)\,\sigma_1(t_1, da_1)...\sigma_n(t_n, da_n)\mu(dt). \tag{4}$$

We endow the behavioral strategy set Σ_i of each player i with the *narrow topology* of Balder (1988). Using one of its equivalent definitions (see Theorem 2.2(b) in Balder (1988)), this is the coarsest topology in which, for every Carathéodory integrand[15] $g : T_i \times A_i \to \mathbb{R}$, the functional $I_g : \Sigma_i \to \mathbb{R}$ that is given for any $\sigma_i \in \Sigma_i$ by

$$I_g(\sigma_i) = \int_{T_i} \int_{A_i} g(t_i, a_i)\sigma_i(t_i, da_i)\mu_i(dt_i) \tag{5}$$

is continuous. By Theorem 2.3(a) of Balder (1988), each Σ_i is *compact* in the narrow topology, and hence Σ is compact in the product topology.

According to Lemma 3 of Carbonell-Nicolau and McLean (2018), upper semi-continuity of ex-post payoffs on A implies upper semi-continuity of the expected payoff functions on Σ. We will apply this lemma[16] to $E(\mathbf{p})$: since, by assumption, $\mathbf{p}(t, \cdot)$ is upper semi-continuous on A for every $t \in T$, it follows that $E(\mathbf{p})$ is *upper semi-continuous* on Σ. Also, as remarked earlier, Σ is compact, and it follows that there *exists* $\overline{\sigma} \in \Sigma$ which maximizes $E(\mathbf{p})$.

5.2 Part 2: Purification of the EP maximizer

In this part of the proof we will show that, given a maximizer $\overline{\sigma}$ of $E(\mathbf{p})$ over Σ, the components of $\overline{\sigma}$ may be sequentially replaced by pure strategies with the resulting

[15] Carathéodory integrand is a $\mathcal{T}_i \otimes \mathcal{B}\,(A_i)$-measurable function $g : T_i \times A_i \to \mathbb{R}$ such that $g(t_i, \cdot)$ is continuous for every $t_i \in T_i$, and there exists a $\mathcal{T}_i$-measurable and μ_i-integrable $\varphi : T_i \to \mathbb{R}_+$ satisfying $|g\,(t_i, a_i)| \leq \varphi\,(t_i)$ for every $(t_i, a_i) \in T_i \times A_i$.

[16] To use that lemma, take all players' ex-post payoffs to be equal to $\mathbf{p}$.

strategy profiles remaining $E(\mathbf{p})$'s maximizers. In such a way, a pure-strategy profile (in S) that maximizes $E(\mathbf{p})$ will be obtained.[17]

We first define a modification of the set Σ_1 that will be a Hausdorff space,[18] proceeding as in Section 2 of Balder (1988). Denote by $\widehat{\Sigma}_1$ the space of *uniformly finite transition measures*, defined as mappings $\widehat{\sigma}_1 : T_1 \times \mathcal{B}(A_1) \to \mathbb{R}$ for which: (i) $\widehat{\sigma}_1(t_1, \cdot)$ is a signed bounded measure on A_1 for every $t_1 \in T_1$; (ii) $\sup_{t_1 \in T_1} |\widehat{\sigma}_1|(t_1, A_1) < \infty$; and (iii) $\widehat{\sigma}_1(\cdot, B)$ is $\mathcal{T}_1$-measurable for every $B \in \mathcal{B}(A_1)$. The narrow topology on $\widehat{\Sigma}_1$ is defined in the same way as on Σ_1 (with $I_g(\widehat{\sigma}_1)$ being defined by (5) using $\widehat{\sigma}_1 \in \widehat{\Sigma}_1$ instead of σ_1). Notice that, clearly, Σ_1 is a subset of $\widehat{\Sigma}_1$, and its narrow topology is the same as the relative narrow topology induced from $\widehat{\Sigma}_1$.

Consider the set

$$\widehat{N}_1 := \left\{ \widehat{\sigma}_1 \in \widehat{\Sigma}_1 \mid I_g(\widehat{\sigma}_1) = 0 \text{ for every Carathéodory integrand on } T_1 \times A_1 \right\},$$

which is closed in the narrow topology. As observed in Balder (1988, p. 267), $\widehat{N}_1$ consists precisely of $\widehat{\sigma}_1 \in \widehat{\Sigma}_1$ such that $\widehat{\sigma}_1(t_1, \cdot)$ is the zero measure for μ_1-almost every t_1. Under the usual quotient mapping $\pi_1 : \widehat{\Sigma}_1 \to \widehat{\Sigma}_1 / \widehat{N}_1$, given by $\pi_1(\widehat{\sigma}_1) = \widehat{\sigma}_1 + \widehat{N}_1 := \left\{ \widehat{\sigma}_1 + \widehat{\sigma}_1' \mid \widehat{\sigma}_1' \in \widehat{N}_1 \right\}$, the quotient space $\widehat{\Sigma}_1 / \widehat{N}_1 = \pi_1\left(\widehat{\Sigma}_1\right)$ is a Hausdorff locally convex topological vector space when equipped with the narrow quotient topology.

Now fix a maximizer $\overline{\sigma} \in \Sigma$ of $E(\mathbf{p})$ whose existence was established in Part 1 of the proof, and consider the compact set $E_1 \subset \Sigma_1$ of the maximizers of $E(\mathbf{p})(\cdot, \overline{\sigma}_{-1})$, which is non-empty because $\overline{\sigma}_1 \in E_1$. The quotient set $\pi_1(E_1)$ is then a compact subset of $\pi_1\left(\widehat{\Sigma}_1\right)$ because π_1 is trivially continuous. Being a compact subset of a Hausdorff locally convex topological vector space, $\pi_1(E_1)$ has an extreme point (see, e.g., Lemma in Holmes (1975, p. 74)). Consider one such extreme point, which must have the form $\pi_1(\tau_1) = \tau_1 + \widehat{N}_1$ for some $\tau_1 \in E_1$. We will show that the strategy τ_1 makes a pure choice of action with μ_1-probability 1; the following lemma will be

[17] Had the set S_i of each player i's pure strategies been a closed subset of Σ_i, the existence of a maximizer of $E(p)$ over S would have been guaranteed from the beginning, but S_i is not necessarily closed in the narrow topology. For instance, if $A_i = [-1, 1], T_i = [0, 1]$ (with $\mathcal{T}_i = \mathcal{B}([0, 1])$ and μ_i being the Lebesgue measure), then the sequence $\{s_i^m\}_{m=1}^{\infty} \subset S_i$ of Rademacher functions – namely, $s_i^m(t_i) = sgn[\sin(2^m \pi t_i)]$ for every $m \geq 1$ – converges in the narrow topology to $\sigma_i \in \Sigma_i \diagdown S_i$ for which $\sigma_i(t_i, \cdot)$ is the uniform distribution on $\{-1, 1\}$ for every $t_i \in T_i$.

[18] Σ_1 itself is not a Hausdorff space because changing any strategy $\sigma_1 \in \Sigma_1$ on a μ_1-null set of types in T_1 produces a distinct element $\sigma_1' \in \Sigma_1$, but σ_1, σ_1' cannot be separated by disjoint open sets in Σ_1.

crucial.

Lemma 3 *For any* $B \in \mathcal{B}(A_1)$, $\mu_1(\{t_1 \mid \tau_1(t_1, B) \in \{0,1\}\}) = 1$.

Proof of Lemma 3. Suppose to the contrary that, for some $C \in \mathcal{B}(A_1)$,

$$\mu_1(\{t_1 \mid \tau_1(t_1, C) \notin \{0,1\}\}) > 0. \tag{6}$$

It follows that there exists $0 < \varepsilon < \frac{1}{2}$ and a set $T' \in \mathcal{T}_1$ such that $\mu_1(T') > 0$ and $\varepsilon < \tau_1(t_1, C) < 1 - \varepsilon$ for every $t_1 \in T'$.

Fix some $\alpha \in \left(2 - \frac{1}{1-\varepsilon}, 1\right)$ (which must be positive since $\varepsilon < \frac{1}{2}$), and define two transition measures $\tau_1^-, \tau_1^+ \in \widehat{\Sigma}_1$ as follows: for any $B \in \mathcal{B}(A_1)$, if $t_1 \in T'$ then

$$\tau_1^-(t_1, B) := \alpha \cdot \tau_1(t_1, B \cap C) + \frac{1 - \alpha\tau_1(t_1, C)}{1 - \tau_1(t_1, C)} \cdot \tau_1(t_1, B \diagdown C)$$

and

$$\tau_1^+(t_1, B) := (2-\alpha) \cdot \tau_1(t_1, B \cap C) + \frac{1 - (2-\alpha)\tau_1(t_1, C)}{1 - \tau_1(t_1, C)} \cdot \tau_1(t_1, B \diagdown C);$$

and if $t_1 \notin T'$ then

$$\tau_1^-(t_1, B) = \tau_1^+(t_1, B) := \tau_1(t_1, B).$$

It is easy to see that $\tau_1^-, \tau_1^+ \in \Sigma_1$, i.e., that τ_1^-, τ_1^+ are behavioral strategies. Furthermore, clearly,

$$\frac{1}{2}\tau_1^- + \frac{1}{2}\tau_1^+ = \tau_1, \tag{7}$$

and hence

$$\frac{1}{2}\pi_1(\tau_1^-) + \frac{1}{2}\pi_1(\tau_1^+) = \pi_1(\tau_1). \tag{8}$$

Also,

$$\pi_1(\tau_1^-) \neq \pi_1(\tau_1^+) \tag{9}$$

because $\tau_1^-(\cdot, C) = \alpha \cdot \tau_1(\cdot, C) \neq (2-\alpha) \cdot \tau_1(\cdot, C) = \tau_1^+(\cdot, C)$ on a positive-probability set T' of types.

Since $E(\mathbf{p})(\sigma)$, by (4), is affine in σ_1, it follows from (7) that

$$E(\mathbf{p})(\tau_1, \overline{\sigma}_{-i}) = \frac{1}{2}E(\mathbf{p})(\tau_1^-, \overline{\sigma}_{-i}) + \frac{1}{2}E(\mathbf{p})(\tau_1^+, \overline{\sigma}_{-i}). \tag{10}$$

But because $\tau_1 \in E_1$, the two summands in the right-hand side of (10) cannot exceed $E(\mathbf{p})(\tau_1, \overline{\sigma}_{-i})$, and hence $E(\mathbf{p})(\tau_1, \overline{\sigma}_{-i}) = E(\mathbf{p})(\tau_1^-, \overline{\sigma}_{-i}) = E(\mathbf{p})(\tau_1^+, \overline{\sigma}_{-i})$, implying that $\tau_1^-, \tau_1^+ \in E_1$. Therefore $\pi_1(\tau_1^-), \pi_1(\tau_1^+) \in \pi_1(E_1)$. However, in light of (8) and

(9), $\pi_1(\tau_1)$ is not an extreme point of $\pi_1(E_1)$. This is a contradiction to the choice of $\pi_1(\tau_1)$. We conclude that there exists no $C \in \mathcal{B}(A_1)$ satisfying (6), which establishes the lemma. ■

Since A_1 is a metrizable compact space, we can fix some metric that induces the topology on A_1, and find a sequence of closed balls $\{B_m\}_{m=1}^{\infty} \subset A_1$ whose diameters converge to 0 as $m \to \infty$, and which satisfy $\cup_{m=k}^{\infty} B_m = A_1$ for every $k \geq 1$. By applying Lemma 3 to every set $B = B_m$ for $m \geq 1$, we have

$$\mu_1(\{t_1 \mid \forall m \geq 1 : \tau_1(t_1, B_m) \in \{0,1\}\}) = 1,$$

i.e., there exists $T_1' \in \mathcal{T}_1$ with $\mu_1(T_1') = 1$ such that $\tau_1(t_1, B_m) \in \{0,1\}$ for every $t_1 \in T_1'$ and every $m \geq 1$. For any fixed $t_1 \in T_1'$, consider a subsequence $\{B_{m_k}\}_{k=1}^{\infty}$ such that $\tau_1(t_1, B_{m_k}) = 1$ for every k; since the diameter of the (closed) sets in $\{B_{m_k}\}_{k=1}^{\infty}$ converges to 0 and A_1 is compact, there exists $s_1(t_1) \in A_1$ such that $\cap_{k=1}^{\infty} B_{m_k} = \{s_1(t_1)\}$.[19] By letting $s_1(t_1)$ to have some constant value $a_1' \in A_1$ for every $t_1 \in T_1 \diagdown T_1'$, a function $s_1 : T_1 \to A_1$ is fully defined.

Observe now that s_1 is $\mathcal{T}_1$-measurable. Indeed, for every set B_m, $s_1^{-1}(B_m) \cap T_1'$ is precisely the set $\tau_1(\cdot, B_m)^{-1}(\{1\}) \cap T_1' \in \mathcal{T}_1$. Because $s_1^{-1}(B_m) \cap (T_1 \diagdown T_1') = T_1 \diagdown T_1'$ if $a_1' \in B_m$ and $s_1^{-1}(B_m) \cap (T_1 \diagdown T_1') = \varnothing$ if $a_1' \notin B_m$, it follows that $s_1^{-1}(B_m) \in \mathcal{T}_1$ for every B_m. But the sets $\{B_m\}_{m=1}^{\infty}$ generate $\mathcal{B}(A_1)$ by Theorem 3.3 of Mackey (1957) since they obviously separate points, and hence s_1 is $\mathcal{T}_1$-measurable by, e.g., Proposition 2.3 of Çinlar (2010). Therefore, $s_1 : T_1 \to A_1$ is in fact a *pure strategy* of player 1.

We conclude that τ_1 makes deterministic choices (given by s_1) for all types in T_1', while the difference between τ_1 and s_1 may only occur for types in $T_1 \diagdown T_1'$; but $\mu_1(T_1 \diagdown T_1') = 0$, and therefore that difference is (expected) payoff-irrelevant for player 1. The implication is that $E(\mathbf{p})(\tau_1, \overline{\sigma}_{-i}) = E(\mathbf{p})(s_1, \overline{\sigma}_{-i})$, and hence that $s_1 \in E_1$.

We have thus shown that the set $E_1 \subset \Sigma_1$ of the maximizers of $E(\mathbf{p})(\cdot, \overline{\sigma}_{-1})$ contains a pure strategy s_1. One may, therefore, replace the first component $\overline{\sigma}_1$ of the EP maximizer $\overline{\sigma} \in \Sigma$ by a pure strategy (if $\overline{\sigma}_1$ is not pure to begin with). Identical

[19] Indeed, the sets in $\{B_{m_k}\}_{k=1}^{\infty}$ have the finite intersection property because $\cap_{k=1}^{l} B_{m_k}$ is of $\tau_1(t_1, \cdot)$-probability 1 for any $l \geq 1$, and hence non-empty. Thus $\cap_{k=1}^{\infty} B_{m_k} \neq \varnothing$ by the compactness of A_1. Furthermore, because the diameters of $\{B_{m_k}\}_{k=1}^{\infty}$ become vanishingly small as $m \to \infty$, $\cap_{k=1}^{\infty} B_{m_k}$ cannot contain two distinct elements.

treatment can now be sequentially applied to the other components of $\overline{\sigma}$, leading to the conclusion that $E(\mathbf{p})$ attains its maximum at a pure strategy profile. From (3), any pure-strategy maximizer of $E(\mathbf{p})$ is a PS-BNE of the game G.

6 Concluding remark

It is easy to see why a treatment akin to the one used in the proof of Theorem 1 cannot be applied to show PS-BNE existence in a general, non-BP, Bayesian game. Indeed, assume that the existence of a behavioral-strategy Bayesian Nash equilibrium (henceforth, BS-BNE) is guaranteed; that is the case, e.g., if all ex-post payoffs $\{u_i(t,\cdot)\}_{i\in N}$ are continuous on A for every $t \in T$, by Theorem 3.1 of Balder (1988). Fixing any BS-BNE $\overline{\sigma}$, one would be able (arguing as in Part 2 of the proof of Theorem 1) to change the behavioral strategy $\overline{\sigma}_1$ of player 1 into a pure strategy s_1 that gives player 1 the same expected payoff, i.e., $U_1(s_1, \overline{\sigma}_{-1}) = U_1(\overline{\sigma})$, and hence remains his best response. However, without advanced purification techniques that would have been enabled by assuming atomlessness and other conditions on the types' distribution (as in, e.g., He and Sun (2019)), the replacement of $\overline{\sigma}_1$ by s_1 would typically alter the expected payoffs of players other than 1, and in particular some player $i \neq 1$ could now have a profitable unilateral deviation given the strategy profile $(s_1, \overline{\sigma}_{-1})$. Thus, there is no guarantee that a sequential purification of BS-BNE strategies results in a BS-BNE, even after just one step.

Only in BP games, with the EP being a surrogate expected payoff that is *common* to all players, any chain of unilateral strategy purifications that starts from an EP maximizer $\overline{\sigma}$ creates a new EP maximizer at each stage; thus, sequentially purifying the strategies of all players leads to a pure-strategy outcome that is an EP maximizer and a PS-BNE.

References

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

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

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

[4] Çinlar E., 2010. Probability and stochastics. Graduate texts in mathematics 261, Springer, New York.

[5] Dvoretsky A., Wald J., Wolfowitz J., 1951. Elimination of randomization in certain statistical decision procedures and zero-sum two-person games. Ann. Math. Stat. 22, 1–21.

[6] Einy E., Haimanko O., 2020. Equilibrium existence in games with a concave Bayesian potential. Game. Econ. Behav. 123, 288-294.

[7] Fabrikant A., Papadimitriou C., Talwar K., 2004. The complexity of pure Nash equilibria. STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing, 604–612.

[8] Fudenberg D., Tirole J., 1991. Game theory. The MIT Press, Cambridge, Massachusetts.

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

[10] Holmes R., 1975. Geometric functional analysis and its applications. Graduate texts in mathematics 24, Springer, New York.

[11] Khan M.A., Rath K.P., Sun Y., 2006. The Dvoretzky-Wald-Wolfowitz theorem and purification in atomless finite-action games. Int. J. Game Theory 34, 91–104.

[12] Mackey G.W., 1957. Borel structures in groups and their duals. Trans. Amer. Math. Soc. 85, 134-135.

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

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

[15] Morris S., Oyama D., Takahashi S., 2022. Implementation via information design in binary-action supermodular games. Manuscript.

[16] Reny P., 2011. On the existence of monotone pure-strategy equilibria in Bayesian games. Econometrica 79(2), 499–553.

[17] Rosenthal R.W., 1973. A class of games possessing pure-strategy Nash equilibria. Int. J. Game Theory 2, 65–67.

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

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

[20] Van Zandt T., Vives X., 2007. Monotone equilibria in Bayesian games of strategic complementarities. J. Econ. Theory 134(1), 339–360.