Existence of Equilibria in Discontinuous Bayesian Games

Wei He* and Nicholas C. Yannelis†

May 7, 2015

Abstract

We provide easily-verifiable sufficient conditions on the primitives of a Bayesian game to guarantee the existence of a behavioral-strategy Bayes-Nash equilibrium. We allow players' payoff functions to be discontinuous in actions, and illustrate the usefulness of our results via an example of an all-pay auction with general tie-breaking rules which cannot be handled by extant results.

JEL classification: C62; C72; D82

Keywords: Discontinuous Bayesian Game; Behavioral Strategy; Random Disjoint Payoff Matching; Equilibrium Existence; All-pay Auction

*Department of Economics, Henry B. Tippie College of Business, The University of Iowa, 108 John Pappajohn Business Building, Iowa City, IA 52242-1994. Email: he.wei2126@gmail.com.

†Department of Economics, Henry B. Tippie College of Business, The University of Iowa, 108 John Pappajohn Business Building, Iowa City, IA 52242-1994. Email: nicholasyannelis@gmail.com.

Contents

1 Introduction

Bayesian games, where each player observes his own private information and then all players choose actions simultaneously, have been extensively studied and found wide applications in many fields of economics. The notion of Bayesian equilibrium is a fundamental game-theoretic concept for analyzing such games. In many applied work, Bayesian games with discontinuous payoffs arise naturally. For example, in auctions or price competitions, players' payoffs may not be continuous when a tie occurs. However, many previous works focus on the case of continuous payoffs,[1] while little is known about equilibrium existence results in Bayesian games with payoff discontinuities.

In a complete information environment, Reny (1999) showed that a better-reply secure game possesses a pure-strategy Nash equilibrium, and proposed the payoff security condition which is sufficient for a game to be better-reply secure together with some other conditions.[2] Recently, several authors have generalized the work of Reny (1999) to an incomplete information setting. Specifically, Carbonell-Nicolau and McLean (2014) extended the "uniform payoff security" condition of Monteiro and Page (2007) and the "uniform diagonal security" condition of Prokopovych and Yannelis (2014) to the setting of Bayesian games, and showed the existence of behavioral/distributional-strategy equilibria. He and Yannelis (2015) proposed the "finite payoff security" condition and proved the existence of pure-strategy equilibria.

The purpose of this paper is to provide a new equilibrium existence result for Bayesian games with discontinuous payoffs. Our result is based on a Bayesian generalization of the clever condition called "disjoint payoff matching", which was introduced by Allison and Lepore (2014) for a normal form game. The advantage of this condition is that one only needs to check the payoff at each strategy profile itself. The standard payoff security-type condition forces one to check the payoffs in the neighborhood of each strategy profile, which is more demanding. Thus, our condition is relatively straightforward, and the equilibrium existence result can be easily verified for a large class of Bayesian games. Our result widens the applications in economics as we can cover situations that previous results in the literature are not readily applicable. As an illustrative example, we provide an application to an all-pay auction with general tie-breaking rules.

The rest of the paper is organized as follows. The model and our main results

[1]See, for example, Milgrom and Weber (1985) and Balder (1988).

[2]A number of recent papers have generalized the work of Reny (1999) in several directions; see Bagh and Jofre (2006), Carmona (2009), Bagh (2010), Carbonell-Nicolau and McLean (2013), Prokopovych (2013), Reny (2013) and Carmona and Podczeck (2014) among others. See also the recent paper of Carmona and Podczeck (2015) for additional references.

are presented in Section 2. Some preparatory results and the proof of the main theorem are collected in Section 3. An illustrative application to an all-pay auction with general tie-breaking rules is given in Section 4. Section 5 provides a purification result. Section 6 concludes the paper.

2 Model

2.1 Bayesian game and behavioral-strategy equilibrium

We consider a **Bayesian game** as follows:

$$G = \{u_i, X_i, (T_i, \mathcal{T}_i), \lambda\}_{i\in I}.$$

- There is **a finite set of players**, $I = \{1, 2, \ldots, n\}$.
- Player i's **action space** X_i is a nonempty compact metric space, which is endowed with the Borel σ-algebra $\mathcal{B}(X_i)$. Denote $X = \prod_{i\in I} X_i$ and $\mathcal{B}(X) = \otimes_{i\in I}\mathcal{B}(X_i)$; that is, $\mathcal{B}(X)$ is the product Borel σ-algebra.
- The measurable space $(T_i, \mathcal{T}_i)$ represents the **private information space** of player i. Let $T = \prod_{i\in I} T_i$ and $\mathcal{T} = \otimes_{i\in I}\mathcal{T}_i$.
- The **common prior** λ is a probability measure on the measurable space $(T, \mathcal{T})$.
- For every player $i \in I$, $u_i : X \times T \to \mathbb{R}_+$ is a $\mathcal{B}(X) \otimes \mathcal{T}$-measurable function representing the **payoff** of player i, which is bounded by some $\gamma > 0$.[3]

As usual, we write t_{-i} for an information profile of all players other than i, and T_{-i} as the space of all such information profiles. We adopt similar notation for action profiles, strategy profiles and payoff profiles.

For every player $i \in I$, a **pure strategy** is a $\mathcal{T}_i$-measurable function from T_i to X_i. Let $\mathcal{L}_i$ be **the set of all possible pure strategies** of player i, and $\mathcal{L} = \prod_{i\in I}\mathcal{L}_i$.

A **behavioral strategy** of player i is a $\mathcal{T}_i$-measurable function from T_i to $\triangle(X_i)$, where $\triangle(X_i)$ denotes the space of all Borel probability measures on X_i under the topology of weak convergence.[4] A pure strategy can be viewed as a special case of a

[3]Since u_i is bounded, we can assume that u_i takes values in $\mathbb{R}_+$ without loss of generality.

[4]That is, a behavioral strategy f_i is a transition probability from $(T_i, \mathcal{T}_i)$ to $(X_i, \mathcal{B}(X_i))$ such that $f_i(\cdot|t_i)$ is a probability measure on $(X_i, \mathcal{B}(X_i))$ for all $t_i \in T_i$, and $f_i(B|\cdot)$ is a $\mathcal{T}_i$-measurable function for every $B \in \mathcal{B}(X_i)$. If λ_i is a probability measure on $(T_i, \mathcal{T}_i)$, then $\lambda_i \diamond f_i$ denotes a probability measure on $T_i \times X_i$ such that $\lambda_i \diamond f_i(A \times B) = \int_A f_i(B|t_i)\lambda_i(\mathrm{d}t_i)$ for any measurable subsets $A \subseteq T_i$ and $B \subseteq X_i$.

behavioral strategy by considering it as a Dirac measure for every t_i. **The set of all behavioral strategies** of player i is denoted by $\mathcal{M}_i$, and $\mathcal{M} = \prod_{i \in I} \mathcal{M}_i$.

Given a behavioral strategy profile $f = (f_1, f_2, \dots, f_n) \in \mathcal{M}$, the **expected payoff** of player i is

$$U_i(f) = \int_T \int_{X_1} \dots \int_{X_n} u_i(x_1, \dots, x_n, t_1, \dots, t_n) f_n(\mathrm{d}x_n|t_n) \dots f_1(\mathrm{d}x_1|t_1) \lambda(\mathrm{d}t).$$

Definition 1. *A* ***behavioral-strategy equilibrium*** *is a behavioral strategy profile $f^* = (f_1^*, f_2^*, \dots, f_n^*) \in \mathcal{M}$ such that f_i^* maximizes $U_i(f_i, f_{-i}^*)$ for any $f_i \in \mathcal{M}_i$ and each player $i \in I$.*[5]

We impose the following assumption on the information structure. Let λ_i be the marginal probability of λ on $(T_i, \mathcal{T}_i)$ for each $i \in I$. Suppose that $(T, \mathcal{T}, \lambda)$ and $(T_i, \mathcal{T}_i, \lambda_i)$ are complete probability measure spaces.

Assumption (Absolue Continuity (AC))**.** *The probability measure λ is absolutely continuous with respect to $\otimes_{i \in I} \lambda_i$ with the corresponding Radon-Nikodym derivative $\psi \colon T \to \mathbb{R}_+$.*

This assumption is widely adopted in the setting of Bayesian games; see, for example, Milgrom and Weber (1985), Balder (1988), Jackson et al. (2002) and Carbonell-Nicolau and McLean (2014). Notice that the (AC) assumption is imposed in Milgrom and Weber (1985) and Balder (1988) even when the payoff function is continuous in the action variables. If players have independent priors in the sense that $\lambda = \otimes_{i \in I} \lambda_i$, then the (AC) assumption holds trivially.

2.2 Normal form game

Below, we convert a Bayesian game G to an (ex ante) normal form game G_0. If one can show the existence of a Nash equilibrium in the game G_0, then this equilibrium corresponds to a behavioral-strategy equilibrium in the original Bayesian game G.

A normal form game G_d is a collection $(X_i, u_i)_{i \in I}$, where X_i and u_i are the action space and payoff function of player i, respectively. We view a Bayesian game G as a normal norm game and denote it by $G_0 = (\mathcal{M}_i, U_i)_{i \in I}$, where $\mathcal{M}_i$ is the set of all possible behavioral strategies, and U_i is the expected payoff function of player i.

[5] Milgrom and Weber (1985) considered distributional strategies and Balder (1988) extended their results to behavioral strategies. As remarked in Milgrom and Weber (1985), every behavioral strategy gives rise to a natural distributional strategy, and every distributional strategy corresponds to an equivalent class of behavioral strategies defined as the induced regular conditional probabilities. We consider behavioral strategies in this paper for simplicity, but all the results can be easily extended to distributional strategies.

A **Nash equilibrium** in the game G_0 is a strategy profile $f^* = (f_1^*, f_2^*, \dots, f_n^*) \in \mathcal{M}$ such that f_i^* maximizes $U_i(f_i, f_{-i}^*)$ for any $f_i \in \mathcal{M}_i$ and each player $i \in I$. Thus, if f^* is a Nash equilibrium in the game G_0, then it is also a behavioral-strategy equilibrium in the original Bayesian game G.

2.3 Main result

Reny (1999) proved that under some regularity conditions, a payoff secure game has a pure-strategy equilibrium.[6] To prove that the mixed extension of a normal form game is payoff secure, Allison and Lepore (2014) introduced the interesting notion of "disjoint payoff matching" in games with complete information. Below, we extend this notion to the setting of Bayesian games, and show that the ex ante game G_0 is payoff secure.

First, we describe the notion of "payoff security", which is due to Reny (1999).

Definition 2. *In a normal form game G_d, player i can secure a payoff $\alpha \in \mathbb{R}$ at $(x_i, x_{-i}) \in X_i \times X_{-i}$ if there is some $\overline{x_i} \in X_i$ such that $u_i(\overline{x_i}, y_{-i}) \geq \alpha$ for all y_{-i} in some open neighborhood of x_{-i}.*

The game G_d is called "payoff secure" if for every $i \in I$, $(x_i, x_{-i}) \in X_i \times X_{-i}$ and $\epsilon > 0$, player i can secure a payoff $u_i(x_i, x_{-i}) - \epsilon$ at (x_i, x_{-i}).

Consider the points at which a player's payoff function is discontinuous in other players' strategies. In particular, let $D_i \colon T_i \times X_i \to T_{-i} \times X_{-i}$ be defined by

$$D_i(t_i, x_i) = \{(t_{-i}, x_{-i}) \in T_{-i} \times X_{-i} \colon u_i(x_i, \cdot, t_i, t_{-i}) \text{ is discontinuous in } x_{-i}\}.$$

Suppose that D_i has a $\mathcal{B}(X) \otimes \mathcal{T}$-measurable graph for each $i \in I$. Given a pure strategy f_i of player i, denote $D_i^{f_i}(t_i) = D_i(t_i, f_i(t_i))$.

Remark 1. *In many applications such as auctions and price competition, the discontinuity arises due to the action variables, and independently of the state variables. That is, the correspondence D_i does not depend on T in the sense that if $(t, x) \in Gr(D_i)$, then $(t', x) \in Gr(D_i)$ for any $t' \in T$. It is usually easy to check that D_i has a measurable graph in such cases.*[7]

Definition 3. *A Bayesian game G is said to satisfy the condition of "random disjoint payoff matching" if for any $f_i \in \mathcal{L}_i$, there exists a sequence of deviations $\{g_i^k\}_{k=1}^{\infty} \subseteq \mathcal{L}_i$ such that the following conditions hold:*

[6] See Prokopovych (2011) for an alternative proof for metric games.

[7] If A is a correspondence from a space Y to Z, then $\mathrm{Gr}(A) \subseteq Y \times Z$ denotes the graph of A.

1. *for λ-almost all $t = (t_i, t_{-i}) \in T$ and all $x_{-i} \in X_{-i}$,*

$$\liminf_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) \geq u_i(f_i(t_i), x_{-i}, t_i, t_{-i});$$

2. $\limsup_{k\to\infty} D_i(t_i, g_i^k(t_i)) = \emptyset$ *for any $i \in I$ and λ_i-almost all $t_i \in T_i$.*

When T_i is a singletons set for any player $i \in I$, the above definition reduces to be the notion of disjoint payoff matching introduced by Allison and Lepore (2014) in a complete information environment.

In a Bayesian game G, if the above deviations of a player depend only on the action variables, but not on the state variables, then it is typically much easier to check the random disjoint payoff matching condition. In particular, such properties are satisfied in many applications where the discontinuity arises due to the presence of the ties. The following lemma provides a sufficient condition which is easy to check.

Lemma 1. *For each player i, suppose that there exists a sequence of measurable function $h_i^k \colon X_i \to X_i$ for $k \geq 1$ such that the following conditions hold:*

1. *for any $x_i \in X_i$,*

$$\liminf_{k\to\infty} u_i(h_i^k(x_i), x_{-i}, t_i, t_{-i}) \geq u_i(x_i, x_{-i}, t_i, t_{-i})$$

for λ-almost all $t \in T$ and all $x_{-i} \in X_{-i}$;

2. $\limsup_{k\to\infty} D_i(t_i, h_i^k(x_i)) = \emptyset$ *for any $i \in I$, $x_i \in X_i$ and λ_i-almost all $t_i \in T_i$.*

Then the condition of random disjoint payoff matching is satisfied.

Proof. For any $i \in I$ and $f_i \in \mathcal{L}_i$, let $g_i^k(t_i) = h_i^k(f_i(t_i))$ for each $k \geq 1$. Then the sequence $\{g_i^k\}$ satisfies the conditions in Defintion 3. □

The following theorem is our main result. It shows that the random disjoint payoff matching condition of a Bayesian game G could guarantee the payoff security of the game G_0.

Theorem 1. *Under Assumption (AC), if a Bayesian game G satisfies the random disjoint payoff matching condition, then the game G_0 is payoff secure.*

2.4 Existence of behavioral-strategy equilibria

Theorem 1 above shows that the random disjoint payoff matching condition of a Bayesian game G guarantees the payoff security of the ex ante game G_0. Reny (1999) showed that a payoff secure game has a pure-strategy Nash equilibrium provided that the game has a compact action spaces, and each player's payoff function is quasiconcave in his own actions and satisfies a certain upper semicontinuity condition. In particular, the condition of aggregate upper semicontinuity of Dasgupta and Maskin (1986) suffices for this aim.

Definition 4. *A normal form game G_d is said to be "aggregate upper semicontinuous" if the summation of the utility functions of all players is upper semicontinuous.*

This notion can be extended to the setting of Bayesian games: a Bayesian game G is called aggregate upper semicontinuous if $\sum_{i \in I} u_i(\cdot, t) \colon X \to \mathbb{R}$ is upper semicontinuous for any $t \in T$.

It is easy to see that the aggregate upper semicontinuity of a Bayesian game G implies the aggregate upper semicontinuity of G_0. We provide a proof below for the sake of completeness (see also Lemma 3 in Carbonell-Nicolau and McLean (2014)).

Lemma 2. *Under Assumption (AC), if a Bayesian game G is aggregate upper semicontinuous, then the game G_0 is aggregate upper semicontinuous.*

Proof. Recall that $\psi \colon T \to \mathbb{R}_+$ is the Radon-Nikodym derivative of the probability measure λ with respect to the product probability $\otimes_{i \in I} \lambda_i$. Let $\phi(x, t) = \sum_{i \in I} u_i(x, t)\psi(t)$. Then ϕ is jointly measurable and upper semicontinuous in x. Define a mapping $H^u \colon \mathcal{M} \to \mathbb{R}$ as follows: for any $f = (f_1, \ldots, f_n) \in \mathcal{M}$

$$
\begin{aligned}
H^u(f_1, \ldots, f_n) &= \int_T \int_X \phi(x, t) f(\mathrm{d}x|t) \otimes_{i \in I} \lambda_i(\mathrm{d}t) \\
&= \int_T \int_X \sum_{i \in I} u_i(x, t)\psi(t) f(\mathrm{d}x|t) \otimes_{i \in I} \lambda_i(\mathrm{d}t) \\
&= \sum_{i \in I} \int_T \int_X u_i(x, t) f(\mathrm{d}x|t) \lambda(\mathrm{d}t) \\
&= \sum_{i \in I} U_i(f).
\end{aligned}
$$

By Lemma 4, H^u is upper semicontinuous. Thus, G_0 is aggregate upper semicontinuous. □

The existence of a behavioral-strategy equilibrium follows as an immediate corollary.

Corollary 1. *Under Assumption (AC), if a Bayesian game G satisfies the random disjoint payoff matching condition and is aggregate upper semicontinuous, then the game G_0 has a Nash equilibrium, which is a behavioral-strategy equilibrium for G.*

Proof. By Theorem 1, the game G_0 is payoff secure. As proved in Lemma 2, G_0 is aggregate upper semicontinuous. By Proposition 3.2 and Theorem 3.1 of Reny (1999), it has a Nash equilibrium, which implies that G has a behavioral-strategy equilibrium. □

Remark 2. *By extending the uniform payoff security condition of Monteiro and Page (2007) and adopting the (AC) assumption, Carbonell-Nicolau and McLean (2014) proved the existence of behavioral/distributional-strategy equilibria in Bayesian games with discontinuous payoffs. In particular, they showed that the ex ante game G_0 is payoff secure when the Bayesian game G satisfies their uniform payoff security condition. Our result does not cover the result of Carbonell-Nicolau and McLean (2014) and vice versa. Notice that our condition only needs to check the payoffs at each strategy profile itself, but not for those payoffs in the neighborhood of the strategy profile.*

3 Proof of Theorem 1

3.1 Preparatory results

The proof of Theorem 1 is based on a clever argument of Allison and Lepore (2014). However, our incomplete information framework introduces several subtle difficulties and necessitates new arguments that are far from trivial. Below, we present some technical results needed for the proof of Theorem 1.

We consider the topology on the behavioral strategy space $\mathcal{M}_i$ for each player $i \in I$.

Definition 5. *A sequence $\{f_i^k\}$ in $\mathcal{M}_i$ is said to weakly converge to some $f_i^0 \in \mathcal{M}_i$ ($f_i^k \Longrightarrow f_i^0$) if for every integrably bounded Carathéodory function*[8] $c : T_i \times X_i \to \mathbb{R}$

$$\lim_{k\to\infty} \int_{T_i} \int_{X_i} c(t_i, x_i) f_i^k(\mathrm{d}x_i|t_i)\lambda_i(\mathrm{d}t_i) = \int_{T_i} \int_{X_i} c(t_i, x_i) f_i^0(\mathrm{d}x_i|t_i)\lambda_i(\mathrm{d}t_i).$$

The weak topology on $\mathcal{M}_i$ is defined as the weakest topology for which the functional $f_i \to \int_{T_i} \int_{X_i} c(t_i, x_i) f_i(\mathrm{d}x_i|t_i)\lambda_i(\mathrm{d}t_i)$ is continuous for every integrably bounded Carathéodory function c.

Let $\mathcal{M}_i$ be endowed with the above topology, and $\mathcal{M} = \Pi_{i\in I}\mathcal{M}_i$ be endowed with the corresponding product topology. The following lemma shows that each player i in the game G_0 is endowed with a nonempty, convex and compact strategy space $\mathcal{M}_i$

[8] The function c is said to be a Carathéodory function if $c(\cdot, x_i)$ is $\mathcal{T}_i$-measurable for each $x_i \in X_i$ and $c(t_i, \cdot)$ is continuous on X_i for each $t_i \in T_i$. In addition, c is called integrably bounded if there exists a λ_i-integrable function $\chi\colon T_i \to \mathbb{R}_+$ such that $|c(t_i, x_i)| \le \chi(t_i)$ for all $(t_i, x_i) \in T_i \times X_i$.

Lemma 3. *$\mathcal{M}_i$ is a convex and compact subset of a topological vector space.*

Proof. See Theorem 2.3 of Balder (1988). □

Lemma 4. *If a sequence $\{f_i^k\}$ in $\mathcal{M}_i$ weakly converges to some $f_i^0 \in \mathcal{M}_i$, then for every integrably bounded measurable function $c\colon T_i \times X_i \to \mathbb{R}$ such that $c(t_i, \cdot)$ is lower semicontinuous in x_i, we have*

$$\liminf_{k\to\infty} \int_{T_i}\int_{X_i} c(t_i, x_i) f_i^k(\mathrm{d}x_i|t_i)\lambda_i(\mathrm{d}t_i) \geq \int_{T_i}\int_{X_i} c(t_i, x_i) f_i^0(\mathrm{d}x_i|t_i)\lambda_i(\mathrm{d}t_i).$$

Proof. See Theorem 2.2 in Balder (1988). □

In the proof of our Theorem 1, we need to deal with some subtle measurability issues based on the projection theorem and Aumann's measurable selection theorem. These theorems are stated below for the convenience of the reader.

Projection Theorem: Let X be a Polish space and $(S, \mathcal{S}, \mu)$ a complete finite measure space. If a set E belongs to $\mathcal{S} \otimes \mathcal{B}(X)$, then the projection of E on S belongs to $\mathcal{S}$.

Aumann's measurable selection theorem: Let X be a Polish space and $(S, \mathcal{S}, \mu)$ a complete finite measure space. Suppose that F is a nonempty valued correspondence from S to X having an $\mathcal{S} \otimes \mathcal{B}(X)$-measurable graph. Then F admits a measurable selection; that is, there is a measurable function f from S to X such that $f(s) \in F(s)$ for μ-almost all $s \in S$.

3.2 Proof

We now proceed with the proof of Theorem 1.

Fix a behavioral strategy profile $(f_1, \ldots, f_n) \in \mathcal{M}$, a player $i \in I$ and $\epsilon > 0$.

Let $S_i\colon T_i \to X_i$ be a correspondence defined by

$$\begin{aligned} S_i(t_i) = \{x_i \in X_i\colon &\int_{T_{-i}}\int_{X_{-i}} u_i(x_i, x_{-i}, t_i, t_{-i})\psi(t_i, t_{-i}) f_{-i}(\mathrm{d}x_{-i}|t_{-i}) \otimes_{j\neq i} \lambda_i(\mathrm{d}t_{-i}) \\ &\geq \int_{T_{-i}}\int_{X} u_i(x_i, x_{-i}, t_i, t_{-i})\psi(t_i, t_{-i}) f(\mathrm{d}x|t_i, t_{-i}) \otimes_{j\neq i} \lambda_i(\mathrm{d}t_{-i})\}. \end{aligned}$$

It is obvious that for each fixed t_i, $S_i(t_i)$ is nonempty. Since u_i is jointly measurable, and f and ψ are measurable, the correspondence S_i has a $\mathcal{B}(X_i) \otimes \mathcal{T}_i$-measurable graph. By the Aumann measurable selection theorem, S_i has a $\mathcal{T}_i$-measurable selection f_i'.

Therefore, we have

$$\int_T \int_{X_{-i}} u_i(f_i'(t_i), x_{-i}, t_i, t_{-i}) f_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t) \geq \int_T \int_X u_i(x_i, x_{-i}, t_i, t_{-i}) f(\mathrm{d}x|t)\lambda(\mathrm{d}t).$$

By the random disjoint payoff matching condition, there exists a sequence of pure strategies $\{g_i^k\} \subseteq \mathcal{L}_i$ such that for λ-almost all $t = (t_i, t_{-i}) \in T$ and all $x_{-i} \in X_{-i}$,

$$\liminf_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) \geq u_i(f_i'(t_i), x_{-i}, t_i, t_{-i}),$$

and $\limsup_{k\to\infty} D_i(t_i, g_i^k(t_i)) = \emptyset$ for λ_i-almost all $t_i \in T_i$.

Let

$$E_i^k(t_i) = \{(t_{-i}, x_{-i})\colon u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) > u_i(f_i'(t_i), x_{-i}, t_i, t_{-i}) - \epsilon\}.$$

Since the functions u_i, g_i^k and f_i' are measurable, the correspondence E_i^k has a $\mathcal{B}(X_{-i}) \otimes \mathcal{T}$-measurable graph. Notice that since $\lambda \diamond f_{-i}$ is a probability measure on $T \times X_{-i}$ and

$$\lambda \diamond f_{-i}\left(\liminf_{k\to\infty} \mathrm{Gr}(E_i^k)\right) = 1,$$

we have that $\lim_{k\to\infty} \lambda \diamond f_{-i}\left(\mathrm{Gr}(E_i^k)\right) = 1$.

Since $\limsup_{k\to\infty} D_i(t_i, g_i^k(t_i)) = \emptyset$ for λ_i-almost all $t_i \in T_i$, it follows that

$$\limsup_{k\to\infty} \lambda \diamond f_{-i}\left(\mathrm{Gr}(D_i^{g_i^k})\right) \leq \lambda \diamond f_{-i}\left(\limsup_{k\to\infty} \mathrm{Gr}(D_i^{g_i^k})\right) = 0.$$

Thus, $\lim_{k\to\infty} \lambda \diamond f_{-i}\left(\mathrm{Gr}(E_i^k) \setminus \mathrm{Gr}(D_i^{g_i^k})\right) = 1$. As a result, there exists some positive integer $K > 0$ such that for any $k \geq K$,

$$\lambda \diamond f_{-i}\left(\mathrm{Gr}(E_i^k) \setminus \mathrm{Gr}(D_i^{g_i^k})\right) > 1 - \epsilon.$$

Let $g_i = g_i^K$ and $F = \mathrm{Gr}(E_i^K) \setminus \mathrm{Gr}(D_i^{g_i^K})$. Then we have

$$\begin{aligned}&\int_F u_i(g_i(t_i), x_{-i}, t_i, t_{-i})\lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i}))\\ &\geq \int_F u_i(f_i'(t_i), x_{-i}, t_i, t_{-i})\lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i})) - \epsilon,\end{aligned}$$

which implies that

$$\int_{T\times X_{-i}} u_i(g_i(t_i), x_{-i}, t_i, t_{-i})\lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i}))$$

$$
\begin{aligned}
&= \int_F u_i(g_i(t_i), x_{-i}, t_i, t_{-i}) \lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i})) \\
&+ \int_{F^c} u_i(g_i(t_i), x_{-i}, t_i, t_{-i}) \lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i})) \\
&\geq \int_F u_i(f_i'(t_i), x_{-i}, t_i, t_{-i}) \lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i})) - \epsilon \\
&+ \int_{F^c} u_i(f_i'(t_i), x_{-i}, t_i, t_{-i}) \lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i})) - \gamma \cdot \epsilon \\
&= \int_{T \times X_{-i}} u_i(f_i'(t_i), x_{-i}, t_i, t_{-i}) \lambda \diamond f_{-i}(\mathrm{d}(t_i, t_{-i}, x_{-i})) - (\gamma + 1)\epsilon.
\end{aligned}
$$

Since X_{-i} is a compact metric space, it is second countable (see Royden and Fitzpatrick (2010, Proposition 25, p.204)). Thus, we can find a countable base $\{V_m\}_{m \geq 1}$ for X_{-i}. Let

$$
h_i^m(x_{-i}, t) = \begin{cases} \inf_{x'_{-i} \in V_m} u_i(g_i(t_i), x'_{-i}, t_i, t_{-i}), & \text{if } x_{-i} \in V_m; \\ -2\gamma, & \text{otherwise.} \end{cases}
$$

It is easy to see that $h_i^m(\cdot, t)$ is lower semicontinuous on X_{-i} for each fixed $t \in T$ and $m \geq 1$. It can be easily checked that h_i^m is a jointly measurable function. Indeed, it suffices to show that for any $c \geq 0$, the set $\{(x_{-i}, t) \in X_{-i} \times T \colon h_i^m(x_{-i}, t) < c\}$ is $\mathcal{B}(X_{-i}) \otimes \mathcal{T}$-measurable. Since u_i is jointly measurable and g_i is $\mathcal{T}_i$-measurable, the set

$$
\{(x_{-i}, t) \in V_m \times T \colon u_i(g_i(t_i), x_{-i}, t_i, t_{-i}) < c\}
$$

is $\mathcal{B}(X_{-i}) \otimes \mathcal{T}$-measurable. By the Projection Theorem, the projection of the above set on T, denoted as T_m, is a $\mathcal{T}$-measurable subset. Notice that

$$
\{(x_{-i}, t) \in X_{-i} \times T \colon h_i^m(x_{-i}, t) < c\} = (V_m \times T_m) \cup (V_m^c \times T), \text{[9]}
$$

which is $\mathcal{B}(X_{-i}) \otimes \mathcal{T}$-measurable. Thus, h_i^m is a jointly measurable function.

Let $\underline{u}_i(x_{-i}, t) = \sup_{m \geq 1} h_i^m(x_{-i}, t)$. For each fixed $t \in T$, as in the proof of Reny (1999, Theorem 3.1), $\underline{u}_i(\cdot, t)$ is the pointwise supremum of a sequence of lower semicontinuous function, which is also lower semicontinuous on X_{-i}. In addition, $\underline{u}_i$ is the supremum of a sequence of $\mathcal{B}(X_{-i}) \otimes \mathcal{T}$-measurable functions, which is also $\mathcal{B}(X_{-i}) \otimes \mathcal{T}$-measurable. Define a function $H_i^l \colon \prod_{j \neq i} \mathcal{M}_j \to \mathbb{R}$ as follows: for $g_{-i} = (g_1, \ldots, g_{i-1}, g_{i+1}, \ldots, g_n)$,

$$
H_i^l(g_{-i}) = \int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t) \psi(t) g_{-i}(x_{-i} | t_{-i}) \otimes_{i \in I} \lambda_i(\mathrm{d}t).
$$

[9]For any subset E, E^c denotes the complement of the set E.

By Lemma 4, H_i^l is lower semicontinuous. Thus, there is an open neighborhood $\mathcal{N}_{-i}(f_{-i}) \subseteq \prod_{j \neq i} \mathcal{M}_j$ of f_{-i} such that for any $g_{-i} \in \mathcal{N}_{-i}(f_{-i})$,

$$\int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t)\psi(t) g_{-i}(x_{-i}|t_{-i}) \otimes_{i\in I} \lambda_i(\mathrm{d}t)$$
$$- \int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t)\psi(t) f_{-i}(x_{-i}|t_{-i}) \otimes_{i\in I} \lambda_i(\mathrm{d}t) - \epsilon.$$

That is,

$$\int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t) g_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t)$$
$$\geq \int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t) f_{-i}(x_{-i}|t_{-i})\lambda(\mathrm{d}t) - \epsilon.$$

Recall that $F = \mathrm{Gr}(E_i^K) \setminus \mathrm{Gr}(D_i^{g_i^K})$. Since $u_i(t, g_i(t_i), \cdot)$ is continuous on the t-section $\{x_{-i} \in X_{-i} \colon (x_{-i}, t) \in F\}$ of F, we have $\underline{u}_i(x_{-i}, t) = u_i(g_i(t_i), x_{-i}, t)$ for any $(x_{-i}, t) \in F$. As a result,

$$\int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t) f_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t)$$
$$= \int_F \underline{u}_i(x_{-i}, t)\lambda \diamond f_{-i}(\mathrm{d}(t, x_{-i})) + \int_{F^c} \underline{u}_i(x_{-i}, t)\lambda \diamond f_{-i}(\mathrm{d}(t, x_{-i}))$$
$$\geq \int_F u_i(g_i(t_i), x_{-i}, t)\lambda \diamond f_{-i}(\mathrm{d}(t, x_{-i}))$$
$$> \int_F u_i(g_i(t_i), x_{-i}, t)\lambda \diamond f_{-i}(\mathrm{d}(t, x_{-i})) + \int_{F^c} u_i(g_i(t_i), x_{-i}, t)\lambda \diamond f_{-i}(\mathrm{d}(t, x_{-i})) - \gamma \cdot \epsilon$$
$$= \int_T \int_{X_{-i}} u_i(g_i(t_i), x_{-i}, t) f_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t) - \gamma \cdot \epsilon.$$

Therefore, for any $g_{-i} \in \mathcal{N}_{-i}(f_{-i})$, we have

$$\int_T \int_{X_{-i}} u_i(g_i(t_i), x_{-i}, t) g_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t)$$
$$\geq \int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t) g_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t)$$
$$\geq \int_T \int_{X_{-i}} \underline{u}_i(x_{-i}, t) f_{-i}(x_{-i}|t_{-i})\lambda(\mathrm{d}t) - \epsilon$$
$$\geq \int_T \int_{X_{-i}} u_i(g_i(t_i), x_{-i}, t) f_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t) - (\gamma + 1) \cdot \epsilon$$
$$\geq \int_T \int_{X_{-i}} u_i(f_i'(t_i), x_{-i}, t) f_{-i}(\mathrm{d}x_{-i}|t_{-i})\lambda(\mathrm{d}t) - 2(\gamma + 1) \cdot \epsilon$$

$$\geq \int_T \int_X u_i(x_i, x_{-i}, t) f(\mathrm{d}x|t)\lambda(\mathrm{d}t) - 2(\gamma + 1) \cdot \epsilon,$$

and consequently, the game G_0 is payoff secure.

4 An Application

Below, we provide an example of an all-pay auction with general tie-breaking rules to demonstrate the usefulness of our result.[10]

All-pay auction with general tie-breaking rules

Suppose that N bidders compete for an object. Each bidder's valuation of the object is given by a measurable function $v\colon \Pi_{i\in I} T_i \to [0,1]$, where T_i is the state space, $i = 1, \ldots, N$. The common prior is λ, and λ is absolutely continuous with respect to $\otimes_{i\in I}\lambda_i$. Bidder i observes his own state t_i and submits a bid $x_i \in X_i = [0,1]$. The bidder who submits the highest bid wins the object and all bidders need to pay their bids. If multiple bidders submit the highest bid simultaneously, then the tie is broken as follows:

$$u_i(x_1, \ldots, x_N, t_1, \ldots, t_N) =$$

$$\begin{cases} -x_i, & x_i < \max_{j\in I} x_j, \\ \frac{\xi_i(x_1,\ldots,x_N)}{\sum_{k\in I: x_k = \max_{j\in I} x_j} \xi_k(x_1,\ldots,x_N)} \cdot v(t_1, \ldots, t_N) - x_i, & x_i = \max_{j\in I} x_j; \end{cases}$$

where $\xi = (\xi_1, \ldots, \xi_N)\colon [0,1]^N \to (0,1]^N$ is a continuous function which assesses the relative importance of each bidder's position when breaking the tie. In particular, if $\xi_i \equiv 1$ for any i, then the tie is broken via the equal probability rule. However, this is not necessary.

Proposition 1. *An all-pay auction with general tie-breaking rules satisfies the random disjoint payoff matching condition.*

[10] Jackson et al. (2002) showed the existence of a distributional-strategy equilibrium for discontinuous games with incomplete information by proposing a solution concept where the payoff is "endogenously defined" at the discontinuities. Araujo, De Castro and Moreira (2008) first considered non-monotonic functions in auctions and showed that an all-pay auction tie-breaking rule is sufficient for the existence of pure-strategy equilibrium for a class of auctions. Carbonell-Nicolau and McLean (2014) considered an all-pay auction with the standard tie-breaking rule that the winning players share the object with equal probability. For other variations, see, for example, Klose and Kovenock (2015) for an all-pay auction with identity-dependent externalities. The results of this section are not covered by any of the above papers.

Proof. Given any bidder i and $f_i \in \mathcal{L}_i$, let

$$g_i^k(t_i) = \begin{cases} \min\{f_i(t_i) + \frac{1}{k}, 1\}, & f_i(t_i) < 1; \\ \frac{1}{k}, & f_i(t_i) = 1. \end{cases}$$

It is obvious that $g_i^k \in \mathcal{L}_i$ for any $k \geq 1$.

Fix any $t \in T$ and $x_{-i} \in X_{-i}$. If $f_i(t_i) = 1$, then $u_i(f_i(t_i), x_{-i}, t_i, t_{-i}) \leq 0$ and $\liminf_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) \geq 0$. If $f_i(t_i) < 1$, we need to consider three possible cases.

1. If bidder i is the unique winner, then he is still the unique winner by adopting the strategy $g_i^k(t_i)$ since $g_i^k(t_i) > f_i(t_i)$. Since $g_i^k(t_i) \to f_i(t_i)$ and ξ is a continuous function, we have $\lim_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) = u_i(f_i(t_i), x_{-i}, t_i, t_{-i})$.
2. If bidder i is one of the multiple winners, then he becomes the unique winner by adopting the strategy $g_i^k(t_i)$. Then

$$\lim_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) = v_i(t_i, t_{-i}) - f_i(t_i) \geq u_i(f_i(t_i), x_{-i}, t_i, t_{-i}).$$

3. If bidder i does not get the object, then he still loses the game by adopting $g_i^k(t_i)$ for sufficiently large k. As a result, $\lim_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) = u_i(f_i(t_i), x_{-i}, t_i, t_{-i})$.

Thus, we have

$$\liminf_{k\to\infty} u_i(g_i^k(t_i), x_{-i}, t_i, t_{-i}) \geq u_i(f_i(t_i), x_{-i}, t_i, t_{-i}),$$

which implies that condition (1) of Definition 3 is satisfied. In addition, for all $t_i \in T_i$, $D_i(t_i, g_i^k(t_i)) = \left\{[0, g_i^k(t_i)]^{N-1} \setminus [0, g_i^k(t_i))^{N-1}\right\} \times T_{-i}$. Since $g_i^k(t_i) \neq g_i^{k'}(t_i)$ for sufficiently large k and k', we have

$$\limsup_{k\to\infty} D_i(t_i, g_i^k(t_i)) = \emptyset$$

for any $t_i \in T_i$. Thus, condition (2) of Definition 3 also holds.

Therefore, an all-pay auction with general tie-breaking rules satisfies the random disjoint payoff matching condition. □

Since $\sum_{i\in I} u_i(t, x) = v(t) - \sum_{i\in I} x_i$, the aggregate upper semicontinuity condition is satisfied. Thus, the existence of a behavioral-strategy equilibrium follows immediately by combining Corollary 1 and Proposition 1.

Corollary 2. *A behavioral-strategy equilibrium exists in an all-pay auction with general tie-breaking rules.*

Remark 3. *Allison and Lepore (2014) presented a Bertrand-Edgeworth oligopoly model which has general specifications of costs, residual demand rationing, and tie-breaking rules. They showed that this price competition problem satisfies the disjoint payoff matching condition and a mixed-strategy equilibrium exists. One can easily extend their model to an incomplete information environment and formulate the problem as a Bayesian game. Then by referring to our Theorem 1 and Corollary 1, one can prove the existence of a behavioral-strategy equilibrium. For further applications on Bayesian games with discontinuous payoffs including the war of attrition, Cournot competition and rent seeking, see Carbonell-Nicolau and McLean (2014).*

5 Purification

By adopting the "random disjoint payoff matching" condition and the "relative diffuseness" condition in He and Sun (2014), we will present a purification result for behavioral-strategy equilibrium in Bayesian games with private values and independent priors, and hence obtain the existence of pure-strategy equilibrium for such games.

For each $i \in I$, let $(T_i, \mathcal{T}_i, \lambda_i)$ be the private information space, and $\mathcal{F}_i \subseteq \mathcal{T}_i$ be the smallest σ-algebra relative to which u_i is measurable. The σ-algebras $\mathcal{T}_i$ and $\mathcal{F}_i$ will represent the diffuseness of information from the aspect of strategies and from the aspect of payoffs, respectively. The probability spaces $(T_i, \mathcal{T}_i, \lambda_i)$ and $(T_i, \mathcal{F}_i, \lambda_i)$ will be used to model the information space and the payoff-relevant information space.

For any non-negligible subset $D \in \mathcal{T}_i$, the restricted probability space $(D, \mathcal{F}_i^D, \lambda_i^D)$ is defined as follows: $\mathcal{F}_i^D$ is the σ-algebra $\{D \cap D' \colon D' \in \mathcal{F}_i\}$ and λ_i^D the probability measure re-scaled from the restriction of λ_i to $\mathcal{F}_i^D$. Furthermore, $(D, \mathcal{T}_i^D, \lambda_i^D)$ can be defined similarly.

Definition 6. *Following the notations above, $\mathcal{F}_i$ is said to be **setwise coarser** than $\mathcal{T}_i$ if for every $D \in \mathcal{T}_i$ with $\lambda_i(D) > 0$, there exists a $\mathcal{T}_i$-measurable subset D_0 of D such that $\lambda_i(D_0 \triangle D_1) > 0$ for any $D_1 \in \mathcal{F}_i^D$.*

The following assumption due to He and Sun (2014) states that on any non-negligible subset $D \subseteq T_i$, $\mathcal{T}_i^D$ is always larger than $\mathcal{F}_i^D$. That is, the strategy-relevant diffuseness of information is essentially richer than the payoff-relevant diffuseness of information.

Assumption (RD). *For each $i \in I$, $(T_i, \mathcal{T}_i, \lambda_i)$ is atomless and $\mathcal{F}_i$ is setwise coarser than $\mathcal{T}_i$.*

Given a behavioral strategy profile $f = (f_1, \dots, f_n) \in \mathcal{M}$, a purification is a pure strategy profile $g = (g_1, \dots, g) \in \mathcal{L}$ such that the expected payoff $U_i(g) = U_i(f)$ for each player $i \in I$. Below, we show the existence of a purification for any behavioral-strategy equilibrium.

Corollary 3. *Suppose that*

1. *Assumption (RD) holds, u_i is measurable with respect to $\mathcal{F}_i$ for each $i \in I$, and $\lambda = \otimes_{i \in I} \lambda_i$;*
2. *every player has private values in the sense that u_i is a measurable function from $X \times T_i$ to $\mathbb{R}_+$;*
3. *the Bayesian game satisfies the random disjoint payoff matching condition, and is aggregate upper semicontinuous.*

Then there exists a purification for any behavioral-strategy equilibrium, and hence a pure-strategy equilibrium exists.

Proof. By Corollary 1, there exists a behavioral-strategy equilibrium f. Then due to Theorem 2 of He and Sun (2014), f has a purification g, which is a pure-strategy equilibrium. □

Remark 4. *By adopting the "relative diffuseness" condition of He and Sun (2014) and the "uniform payoff security" condition of Carbonell-Nicolau and McLean (2014), He and Yannelis (2015) presented a purification result for behavioral-strategy equilibrium in Bayesian games with discontinuous payoffs. As discussed in Remark 2 above, our result here has the advantage that we only need to check the payoffs at each strategy profile itself, but not for those payoffs in the neighborhood of each strategy profile.*

Remark 5. *It was pointed out in He and Yannelis (2015), an existence result of a mixed-strategy equilibrium in a normal form game can be understood as an existence result of a pure-strategy equilibrium in a Bayesian game with state-irrelevant payoffs. In particular, suppose that each player can only observe his own private signal from the unit interval $[0, 1]$, which is endowed with the uniform distribution η. Let $T = [0, 1]^n$ be the state space. The payoff of each player only depends on the action profile, but not on the state profile. Then the normal form game is reformulated as a Bayesian game with state-irrelevant payoffs. The mixed strategy m_i of player i in the normal form game can be realized by his private signal (like a randomization device) to be a pure strategy f_i in the sense that $m_i = \eta \circ f_i^{-1}$. It is easy to check that $f = (f_1, \dots, f_n)$ is a pure-strategy equilibrium in this Bayesian game.*

If we view a normal form game as a Bayesian game, then $\mathcal{F}_i = \{\emptyset, [0,1]\}$ for each player $i \in I$, since players' payoffs do not depend on the states, and $\mathcal{T}_i$ is the Borel σ-algebra on $[0,1]$. Thus, Assumption (RD) trivially holds, and our Corollary 3 extends Allison and Lepore (2014) by allowing for state-dependent payoffs.

6 Concluding Remarks

The purpose of this paper was to prove a new theorem on the existence of behavioral-strategy equilibria for Bayesian games with discontinuous payoffs. Our result is different from the recent ones in Carbonell-Nicolau and McLean (2014) and He and Yannelis (2015). We applied our equilibrium existence theorem to an all-pay auction with general tie-breaking rules, and also indicated further applications to oligopoly theory. It remains an open question whether the existence result of this paper can be extended to a setting of a continuum of players. Such an extension will further widen the economic applications.

References

B. A. Allison and J. J. Lepore, Verifying payoff security in the mixed extension of discontinuous games, *Journal of Economic Theory* **152** (2014), 291–303.

A. Araujo, L. I. De Castro and H. Moreira, Non-monotoniticies and the all-pay auction tie-breaking rule, *Economic Theory* **35** (2008), 407–440.

A. Bagh, Variational convergence: approximation and existence of equilibria in discontinuous games, *Journal of Economic Theory* **145** (2010), 1244–1268.

A. Bagh and A. Jofre, Reciprocal upper semicontinuity and better reply secure games: a comment, *Econometrica* **74** (2006), 1715–1721.

E. J. Balder, Generalized equilibrium results for games with incomplete information, *Mathematics of Operations Research* **13** (1988), 265–276.

O. Carbonell-Nicolau and R. P. McLean, Approximation results for discontinuous games with an application to equilibrium refinement, *Economic Theory* **54** (2013), 1–26.

O. Carbonell-Nicolau and R. P. McLean, On the existence of Nash equilibrium in Bayesian games, working paper, Rutgers University, 2014.

G. Carmona, An existence result for discontinuous games, *Journal of Economic Theory* **144** (2009), 1333–1340.

G. Carmona and K. Podczeck, Existence of Nash equilibrium in games with a measure space of players and discontinuous payoff functions, *Journal of Economic Theory* **152** (2014), 130–178.

G. Carmona and K. Podczeck, Existence of Nash equilibrium in ordinal games with discontinuous preferences, working paper, 2015.

P. Dasgupta and E. Maskin, The existence of equilibrium in discontinuous economic games. Part I: theory, *Review of Economic Studies* **53** (1986), 1–26.

W. He and X. Sun, On the diffuseness of incomplete information game, *Journal of Mathematical Economics* **54** (2014), 131–137.

W. He and N. C. Yannelis, Discontinuous games with asymmetric information: An extension of Reny's existence theorem, *Games and Economic Behavior* **91** (2015), 26–35.

M. O. Jackson, L. K. Simon, J. M. Swinkels and W. R. Zame, Communication and equilibrium in discontinuous games of incomplete information, *Econometrica* **70** (2002), 1711–1740.

B. Klose and D. Kovenock, The all-pay auction with complete information and identity-dependent externalities, *Economic Theory* **59** (2015), 1–19.

P. R. Milgrom and R. J. Weber, Distributional strategies for games with incomplete information, *Mathematics of Operations Research* **10** (1985), 619–632.

P. K. Monteiro and F. H. Page Jr., Uniform payoff security and Nash equilibrium in compact games, *Journal of Economic Theory* **134** (2007), 566–575.

P. Prokopovych, On equilibrium existence in payoff secure games, *Economic Theory* **48** (2011), 5–16.

P. Prokopovych, The single deviation property in games with discontinuous payoffs, *Economic Theory* **53** (2013), 383–402.

P. Prokopovych and N. C. Yannelis, On the existence of mixed strategy Nash equilibria, *Journal of Mathematical Economics* **52** (2014), 87–97.

P. J. Reny, On the existence of pure and mixed strategy Nash equilibria in discontinuous games, *Econometrica* **67** (1999), 1029–1056.

P. J. Reny, Nash equilibrium in discontinuous games, working paper, University of Chicago, 2013.

H. L. Royden and P. M. Fitzpatrick, Real Analysis, 4th edn., Prentice Hall, Boston, 2010.