games

Article

Invariant Equilibrium in Discontinuous Bayesian Games

Blake A. Allison [1] and Jason J. Lepore [2,*]

[1] Department of Economics, Emory University, Atlanta, GA 30322, USA; baallison@emory.edu
[2] Department of Economics, Orfalea College of Business, California Polytechnic State University, San Luis Obispo, CA 93407, USA
* Correspondence: jlepore@calpoly.edu

Abstract: We provide sufficient conditions on the primitives of a class of discontinuous Bayesian games such that all games in the class share equilibria. If a Bayesian game in the class also satisfies a weak efficiency condition, then we show its normal form is better-reply secure. The invariance property then provides an existence result for all Bayesian games in the class. Results are shown for both pure strategy and behavioral strategy equilibrium. We illustrate the application of the results with an example of a class of contests with bid caps.

Keywords: discontinuous Bayesian game; invariance; equilibrium existence; random superior payoff matching; random weak efficiency

check for updates

Citation: Allison, B.A.; Lepore, J.J. Invariant Equilibrium in Discontinuous Bayesian Games. *Games* **2024**, *15*, 17. https://doi.org/10.3390/g15030017

Academic Editors: Klaus Ritzbergera and Ulrich Berger

Received: 29 February 2024
Revised: 22 April 2024
Accepted: 16 May 2024
Published: 20 May 2024

1. Introduction

In many applications of Bayesian games, such as auctions, contests, or in oligopoly pricing games, payoff discontinuities naturally occur. Recent literature has developed that provides sufficient conditions for the existence of equilibrium in discontinuous[1]. This literature leverages the complete information environment results of [3], who shows that a "better-reply secure" game possesses a Nash equilibrium, by establishing conditions on the primitive of a Bayesian game that are sufficient for its normal form to be better-reply secure[2]. Ref. [4] introduce "finite payoff security" and use it to show the existence of pure strategy equilibrium in a Bayesian game. Ref. [5] extend the "uniform payoff security" condition of [6], and the "uniform diagonal security" condition of [7] to show the existence of behavioral strategy equilibrium in a Bayesian game. Ref. [8] extend the "disjoint payoff matching" condition of [9] to show the existence of behavioral strategy equilibrium in a Bayesian game. These contributions provide valuable new results of the domain of discontinuous Bayesian games.

In this paper, we provide new sufficient conditions, both for the existence of pure strategy equilibrium and behavioral strategy equilibrium, for a class of Bayesian games with discontinuous payoffs. Our approach is based on a Bayesian generalization of the complete information conditions of [10]. The results for both pure and behavioral strategies are based on two types of conditions. The first type of condition is a Bayesian game generalization of "superior payoff matching", which requires that at any given strategy profile, each player can match the highest payoff that they would receive near that strategy profile across all games within that class. There are separate conditions that apply to pure strategies and behavioral strategies. These matching conditions are used to show that a class of games possesses the same equilibrium, what we call an invariant class of games. The second type of condition is a Bayesian game generalization of normal form "weak efficiency" from [11], which requires that in almost all fixed-type sections of the game, all players receive their highest possible payoff at any strategy profile for which this payoff selection is simultaneously feasible. The same second type of condition is used for both pure strategies and behavioral strategies. We show that if a Bayesian game satisfies these two types of conditions in pure strategies (behavioral strategies), then its normal form (mixed extension) is better-reply secure. Based on the invariance results, the existence of equilibrium is shown for all Bayesian games in the same class, which can include games that are not payoff secure.

Our results cover situations in which the extant literature is not directly applicable. In particular, our results avoid the need to verify reciprocal upper semicontinuity[3]. Such verification is particularly challenging, both in the abstract and in application, as there is no clear connection between reciprocal upper semicontinuity of the game with fixed types and the Bayesian game, nor between the normal form and mixed extension of a game. In the existing literature, this property has been guaranteed by the far more restrictive assumption that the sum of the payoffs is upper semicontinuous. Our adaptation of weak efficiency along with invariance is a novel approach that allows verification of better-reply security in a class of games for which the sum of the payoffs is not upper semicontinuous. Since satisfaction of our matching conditions provides equilibrium invariance across a class of Bayesian games, only one game in the class must satisfy our weak efficiency condition to show existence for the entire class of Bayesian games. We apply our results to a contest with bid caps as an illustrative example in which the contest type-section games violate reciprocal upper semicontinuity.

The remainder of the paper proceeds as follows: The game environment and all preliminary definitions are presented in Section 2. The primary results are presented in Section 3. The example of a class of contests with bid caps is presented in Section 4.

2. Preliminaries

2.1. A Class of Bayesian Games

Consider a class of Bayesian games $\mathcal{G}$. Each Bayesian game $G = (u, X, (T, \mathcal{T}), \lambda)$ in the class $\mathcal{G}$ is as follows:

There is a finite set of players $I = \{1, 2, \ldots, n\}$, which is identical for all games within the class $\mathcal{G}$. Each player i's action space X_i is a nonempty compact metric space endowed with a Borel σ-algebra $\mathcal{B}(X_i)$. As is standard, we denote the action space by $X = \prod_{i \in I} X_i$ and the product Borel σ-algebra by $\mathcal{B}(X) = \otimes_{i \in I} \mathcal{B}(X_i)$.

The measurable space $(T_i, \mathcal{T}_i)$ represents the private information space of each player i. We denote the products as $T = \prod_{i \in I} T_i$ and $\mathcal{T} = \otimes_{i \in I} \mathcal{T}_i$. The common prior λ is a probability measure on $(T, \mathcal{T})$. Denote by λ_i the marginal probability of λ on $(T_i, \mathcal{T}_i)$ for each $i \in I$. The measure spaces $(\lambda_i, T_i, \mathcal{T}_i)$ and $(\lambda, T, \mathcal{T})$ are assumed to be complete probability measure spaces. The common prior λ is absolutely continuous with respect to $\otimes_{i \in I} \lambda_i$ with the corresponding Radon-Nikodym derivative $\psi : T \mapsto \mathbb{R}_+$.

Each game in the class has a particular payoff selection from a countable set of payoff functions $\mathcal{U}$. The set of payoffs is such that for every $u = (u_1, u_2, \ldots, u_n) \in \mathcal{U}$, each player i's payoff $u_i : X \times T \mapsto \mathbb{R}_+$ is a bounded function that is $\mathcal{B}(X) \otimes \mathcal{T}$-measurable. A class of games only varies based on the payoff selection $u \in \mathcal{U}$. The notation $G(u)$ is used when it is necessary to be explicit about a particular payoff selection u.

We denote a selection of information profiles by $t \in T$. As is standard, we refer to the information profile of all players other than i by t_{-i} and the set of all such information profiles T_{-i}. A similar notation is used for action profiles, strategy profiles, and payoff profiles. We refer to an information selection $t_i \in T_i$ as a type for player i.

2.2. Strategies and Expected Payoffs

A *pure strategy* of player i is a $\mathcal{T}_i$-measurable function $s_i : T_i \mapsto X_i$. Denote by S_i, the set of all possible pure strategies for player i. Let the set of all possible pure strategy profiles for the game be denoted by $S = \prod_{i=1}^{n} S_i$. A *behavioral strategy* for player i is a $\mathcal{T}_i$-measurable function $\delta_i : T_i \mapsto \Delta(X_i)$, where $\Delta(X_i)$ is the set of all Borel probability measures on X_i endowed with the topology of weak convergence. Denote the set of all behavioral strategies for player i by M_i, with $M = \prod_{i=1}^{n} M_i$.

The expected utility of player i given the pure strategy profile $s \in S$ is

$$U_i(s) = \int_T u_i(s(t), t)\lambda(dt). \tag{1}$$

Given a behavioral strategy profile $\delta \in M$, the expected utility of player i

$$\boldsymbol{U}_i(\delta) = \int_T \int_X u_i(x, t)\delta(dx|t)\lambda(dt). \tag{2}$$

With some abuse of notation, the expected utility of player i given the pure strategy $s_i \in S_i$ and behavioral strategies $\delta_{-i} \in M_{-i}$ is written as

$$\boldsymbol{U}_i(s_i, \delta_{-i}) = \int_T \int_{X_{-i}} u_i(s_i(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt). \tag{3}$$

Note that a pure strategy $s_i \in S_i$ or pure strategy profile $s \in S$ has an associated behavioral strategy or strategy profile that will be denoted by $f_{s_i} \in M_i$ or $f_s \in M$, respectively, where f_s is the Dirac measure for which $f_s(E|t) = 1$ if and only if $s(t) \in E$.

A *pure strategy equilibrium* is a strategy profile $s^* \in S$ such that $U_i(s^*) \geq U_i(s_i, s^*_{-i})$ for all $s_i \in S_i$ and each $i \in I$.

A *behavioral strategy equilibrium* is a strategy profile $\delta^* \in M$ such that $\boldsymbol{U}_i(\delta^*) \geq \boldsymbol{U}_i(\delta_i, \delta^*_{-i})$ for all $\delta_i \in M_i$ and each $i \in I$.

Denote by $EQ(u)$ the set of pure strategy equilibria of the Bayesian game $G(u)$ and denote by $\widetilde{EQ}(u)$ the set of behavioral strategy equilibria of the Bayesian game $G(u)$.

Before continuing, we must establish the topology on M_i. Let $\mathcal{H}_i$ be the space of uniformly finite transition measures from $(T_i, \mathcal{T}_i, \lambda_i)$ to $(X_i, \mathcal{B}_i(X_i))$. The weak topology on $\mathcal{H}_i$ is the weakest topology for which the functional $\nu \to \int_{T_i} \int_{X_i} c(t_i, x_i)\nu(dx_i|t_i)\lambda_i(dt_i)$ is continuous on $\mathcal{H}_i$ for every integrably bounded Caratheodory function c, i.e., for every function c for which $c(\cdot, x_i)$ is $\mathcal{T}_i$-measurable and $c(t_i, \cdot)$ is continuous. The space M_i is a subspace of $\mathcal{H}_i$ endowed with the relative topology, which we denote by Υ_i. The space M is thus endowed with the product topology $\Upsilon = \otimes_{i \in N}\Upsilon_i$.

2.3. Normal Form

We now express the class of *ex ante* normal form games $\mathbb{G}$ that corresponds to the class of Bayesian games $\mathcal{G}$. First, let us denote by $G_d = (X, u)$ a normal form game with the set of players I, the action space X, and a payoff selection $u \in \mathcal{U}$. The Bayesian game G can be expressed as a normal form game $G_0 = (S, U)$ with pure strategies S and expected payoffs defined by (1). Further, the mixed extension of the normal form version of G is denoted by $\widetilde{G}_0 = (M, \boldsymbol{U})$.

A *Nash equilibrium* of the game G_0 is a strategy profile $s^* \in S$ such that $U_i(s^*) \geq U_i(s_i, s^*_{-i})$ for all $s_i \in S_i$ and each $i \in I$. A *mixed strategy Nash equilibrium* of the game G_0 is a Nash equilibrium of $\widetilde{G_0}$, that is, a strategy profile $\delta^* \in M$ such that $\boldsymbol{U}_i(\delta^*) \geq \boldsymbol{U}_i(\delta_i, \delta^*_{-i})$ for all $\delta_i \in M_i$ and each $i \in I$. Clearly, the set of pure strategy equilibria of the Bayesian game G is the same as the set of Nash equilibria in its normal form and the set of behavioral strategy equilibria of the Bayesian game G is the same as the set of mixed strategy Nash equilibria in its normal form.

2.4. Extreme Payoffs

We define the upper and lower envelopes of a player's payoff for each fixed information profile within a given game as well as for the entire class of games. First, given a countable base $\{\mathcal{V}_m\}_{m \geq 1}$ for X, we define for each payoff selection $u \in \mathcal{U}$ the functions $\underline{u}^m_i(x, t)$ and $\overline{u}^m_i(x, t)$ as follows, where γ is an arbitrary upper bound on all u_i[4]:

$$\overline{u}^m_i(x, t) = \begin{cases} \sup_{x \in \mathcal{V}_m} u_i(x, t) & \text{if } x \in \mathcal{V}_m, \\ \gamma & \text{otherwise.} \end{cases},$$

and

$$\underline{u}_i^m(x,t) = \begin{cases} \inf_{x \in \mathcal{V}_m} u_i(x,t) & \text{if } x \in \mathcal{V}_m, \\ 0 & \text{otherwise.} \end{cases}.$$

$\overline{u} = (\overline{u}_1, \overline{u}_2, \ldots, \overline{u}_I)$ and $\underline{u} = (\underline{u}_1, \underline{u}_2, \ldots, \underline{u}_I)$, where $\overline{u}_i$ and $\underline{u}_i$ are the upper and lower envelopes of u_i, respectively, defined for each player i as

$$\begin{aligned} \overline{u}_i(x,t) &= \inf_{m \geq 1} \overline{u}_i^m(x,t), \text{ and} \\ \underline{u}_i(x,t) &= \sup_{m \geq 1} \underline{u}_i^m(x,t). \end{aligned}$$

Similarly, define

$$\overline{\pi}_i^m(x,t) = \begin{cases} \sup_{x \in \mathcal{V}_m} \sup_{u \in \mathcal{U}} \overline{u}_i(x,t) & \text{if } x \in \mathcal{V}_m, \\ \gamma & \text{otherwise.} \end{cases},$$

and

$$\underline{\pi}_i^m(x,t) = \begin{cases} \inf_{x \in \mathcal{V}_m} \inf_{u \in \mathcal{U}} \underline{u}_i(x,t) & \text{if } x \in \mathcal{V}_m, \\ 0 & \text{otherwise.} \end{cases}.$$

We then define $\overline{\pi}_i$ and $\underline{\pi}_i$ to be the upper and lower envelopes across all payoff selections, respectively, defined for each player i as

$$\begin{aligned} \overline{\pi}_i(x,t) &= \inf_{m \geq 1} \overline{\pi}_i^m(x,t), \text{ and} \\ \underline{\pi}_i(x,t) &= \sup_{m \geq 1} \underline{\pi}_i^m(x,t). \end{aligned}$$

We state the following preliminary result that is used repeatedly in what follows.

Lemma 1. *For all i, the functions $\overline{u}_i, \underline{u}_i, \overline{\pi}_i$, and $\underline{\pi}_i$ are $\mathcal{B}(X) \otimes \mathcal{T}$-measurable.*

The proof of this lemma and all others not appearing in the main text are provided in the Appendix A.

The set of action profiles for which at least one player's payoff is not maximal is $\Sigma(u) = \{(x,t) \in X \times T : u(x,t) < \overline{\pi}(x,t)\}$ and we define $\mathcal{U}_u = \{v \in \mathcal{U} : \Sigma(v) \subset \Sigma(u)\}$. The set of pure strategy equilibrium shared by the class of Bayesian games with payoffs $v \in \mathcal{U}_u$ is denoted by $IE(\mathcal{U}_u) = \bigcap_{v \in \mathcal{U}_u} EQ(v)$. We call this the set of invariant pure strategy equilibrium for the games $\mathcal{U}_u$.

3. Main Results

In this section, we establish sufficient conditions for equilibrium invariance of a class of games in terms of pure strategies and behavioral strategies.

3.1. Invariance

We begin by formally defining the notion of random matching, which serves as a basis for the conditions developed within this paper. For the purpose of the following definition, we denote a profile of payoffs by $\phi : T \to \mathbb{R}$. Given an information vector $t \in T$, a player i can *match* a payoff $\phi(t)$ at a strategy profile if player i can deviate and receive a payoff that is either greater than $\phi(t)$ or arbitrarily close. Player i can *random match* the profile ϕ if player i can match the payoff $\phi(t)$ for almost all t. This notion is formalized in the following definition.

Definition 1. *Given a Bayesian game $G \in \mathcal{G}$, player i can* random match *ϕ at $s \in S$ if, for any $\varepsilon > 0$, there exists an $s_i' \in S_i$ such that $u_i(s_i'(t_i), s_{-i}(t_{-i}), t) \geq \phi(t) - \varepsilon$ for λ-almost all $t \in T$.*

Next, we define our primary matching condition.

Definition 2. *A Bayesian game $G \in \mathcal{G}$ satisfies* random superior payoff matching (RSPM) *if each player $i \in N$ can random match $\overline{\pi}_i(s) = (\overline{\pi}_i(s(t), t))_{t \in T}$ for any $s \in S$.*

Now we state our result pertaining to the invariance of pure strategy equilibrium across a class of games.

Theorem 1. *Let $\mathcal{G}$ be a class of games, and $u \in \mathcal{U}$ be such that $G(u)$ satisfies RSPM, then $EQ(u) = IE(\mathcal{U}_u)$.*

The following lemma is used in the proofs of Theorems 1 and 2.

Lemma 2. *Suppose that each player i can random match a measurable function $\phi_i(s, t)$ for any $s \in S$. Then, in any equilibrium $s^* \in EQ(u)$, each player's equilibrium payoff $U_i(s^*) \geq \int_T \phi_i(s^*, t)\lambda(dt)$. Consequently, if a Bayesian game $G \in \mathcal{G}$ satisfies RSPM, then $U(s^*) = \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$ in any equilibrium $s^* \in EQ(u)$.*

Proof of Lemma 2. Let $s^* \in EQ(u)$ and suppose that $U_i(s^*) < \int_T \phi_i(s^*(t), t)\lambda(dt)$ for some player i. Choose $\varepsilon \in (0, \int_T \phi_i(s^*(t), t)\lambda(dt) - U_i(s^*))$. Since player i can random match ϕ_i, each player has a deviation $s_i \in S_i$ such that

$$u_i(s_i'(t_i), s_{-i}(t_{-i}), t) \geq \overline{\pi}_i(s(t), t) - \varepsilon$$

for λ-almost all $t \in T$. Since u_i and $\overline{\pi}_i$ are $\mathcal{B}(X) \otimes \mathcal{T}$-measurable, it follows that for this deviation that

$$\begin{aligned} \int_T u_i(s_i'(t_i), s_{-i}(t_{-i}), t)\lambda(dt) &\geq \int_T \phi(s(t), t)\lambda(dt) - \varepsilon \\ &> U_i(s^*). \end{aligned} \tag{4}$$

This contradicts s^* as an equilibrium.

The second conclusion of the lemma follows directly from the fact that $u \leq \overline{\pi}$. □

Proof of Theorem 1. Let $s^* \in EQ(u)$. Lemma 2 implies that $U(s^*) = \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$. We will use this fact to argue that $EQ(u) \subset EQ(v)$ for all $v \in \mathcal{U}_u$.

Let $v \in \mathcal{U}_u$ and suppose $s^* \notin EQ(v)$. Since $u \leq \overline{\pi}$ and $U(s^*) = \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt)$, it must be that $u(s^*(t), t) = \overline{\pi}(s^*(t), t)$ for λ-almost all $t \in T$. By definition of $\mathcal{U}_u$, $v(s^*(t), t) = \overline{\pi}(s^*(t), t)$ whenever $u(s^*(t), t) = \overline{\pi}(s^*(t), t)$, and thus $v(s^*(t), t) = \overline{\pi}(s^*(t), t)$ for λ-almost all $t \in T$. It follows that

$$V(s^*) = \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt). \tag{5}$$

Since $s^* \notin EQ(v)$, there exists a player $i \in I$ with strategy $s_i' \in S_i$ such that

$$\begin{aligned} V_i(s_i', s_{-i}^*) &> V_i(s^*) \\ &= \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt). \end{aligned}$$

Let $\varepsilon \in (0, V_i(s_i', s_{-i}^*) - \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt))$. From RSPM, there exists an $s_i'' \in S_i$ such that

$$\begin{aligned} U_i(s_i'', s_{-i}^*) &\geq \int_T \overline{\pi}_i(s_i'(t_i), s_{-i}^*(t), t)\lambda(dt) - \varepsilon \\ &\geq V_i(s_i', s_{-i}^*) - \varepsilon \\ &> V_i(s_i', s_{-i}^*) - \left(V_i(s_i', s_{-i}^*) - \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt)\right) \\ &= \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt) \\ &= U_i(s^*). \end{aligned}$$

This contradicts $s^* \in EQ(u)$. We conclude that $s^* \in EQ(v)$ and $IE(\mathcal{U}_u) = \bigcap_{v \in \mathcal{U}_u} EQ(v) = EQ(u)$. □

At this point, we turn to establishing invariance results for the set of behavioral strategy equilibrium.

Definition 3. *A Bayesian game $G \in \mathcal{G}$ satisfies* random uniform superior payoff matching (RUSPM) *if each player $i \in N$ can random match $\overline{\pi}_i(s) = (\overline{\pi}_i(s(t), t))_{t \in T}$ using the same $\bar{s}_i = (\bar{s}_i(t_i))_{t_i \in T_i} \in S_i$ for every $s_{-i} \in S_{-i}$.*

The following theorem shows that RUSPM is sufficient for the invariance of the set of behavioral strategy equilibrium.

Theorem 2. *Let $\mathcal{G}$ be a class of games; take $u \in \mathcal{U}$ and suppose that $G(u)$ satisfies RUSPM, then $\widetilde{EQ}(u) = \widetilde{IE}(\mathcal{U}_u)$.*

The proof of Theorem 2 is similar to that of Theorem 1; however, additional care must be taken to avoid the necessity of defining the less intuitive analogues of $\overline{u}$ and $\overline{\pi}$ in the mixed extension.

Proof of Theorem 2. First, we show that each player i can random match $\int_T \int_{X_{-i}} \overline{\pi}_i \delta(dx|t)\lambda(dt)$ at any $\delta \in M$. Second, we argue that $\boldsymbol{U}(\delta^*) = \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$ at any $\delta^* \in \widetilde{EQ}(u)$. Lastly, we argue that $\widetilde{EQ}(u) \subset \widetilde{EQ}(v)$.

Let $\delta \in M$ and $i \in I$. From Lemma A2 (in Appendix A), there is a $\mathcal{T}$-measurable selection g_i' such that

$$\int_T \int_{X_{-i}} \overline{\pi}_i(g_i'(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt) \geq \int_T \int_X \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt).$$

From RUSPM, for all $\varepsilon > 0$, there exists $s_i' \in S_i$ such that

$$u_i(s_i'(t_i), x_{-i}, t) \geq \overline{\pi}_i(g_i'(t_i), x_{-i}, t) - \varepsilon,$$

for all $x_{-i} \in X_{-i}$, and λ-almost all $t \in T$. It follows that for all $\varepsilon > 0$,

$$\begin{aligned} \int_T \int_{X_{-i}} u_i(s_i'(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt) &\geq \int_T \int_{X_{-i}} \overline{\pi}_i(g_i'(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt) - \varepsilon \\ &\geq \int_T \int_X \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt) - \varepsilon. \end{aligned} \tag{6}$$

We conclude that each player i can random match $\int_T \int_{X_{-i}} \overline{\pi}_i \delta(dx|t)\lambda(dt)$ at any $\delta \in M$.

Lemma 2 then implies that $\boldsymbol{U}(\delta^*) \geq \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$ at any $\delta^* \in \widetilde{EQ}(u)$. Combining the statement of the previous sentence with the fact that $u \leq \overline{\pi}$ implies that $\boldsymbol{U}(\delta^*) = \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$.

We now show that $\widetilde{EQ}(u) \subset \widetilde{EQ}(v)$. Let $\delta^* \in \widetilde{EQ}(u)$ and suppose to the contrary that $\delta^* \notin \widetilde{EQ}(v)$. Then there exists a player $i \in I$ with behavioral strategy $\delta_i' \in S_i$ such that

$$\begin{aligned} \boldsymbol{V}_i(\delta_i', \delta_{-i}^*) &> \boldsymbol{V}_i(\delta^*) \\ &= \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt). \end{aligned}$$

Let $\varepsilon \in (0, \boldsymbol{V}_i(\delta_i', \delta_{-i}^*) - \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)\ )$ and applying the condition (6), there must be a $\delta_i'' \in S_i$ such that

$$\begin{aligned}
\boldsymbol{U}_i(\delta_i'', \delta_{-i}^*) &\geq \int_T \int_X \overline{\pi}_i(x,t)\delta_i'(dx_i|t_i)\delta_{-i}^*(dx_{-i}|t_{-i}) - \varepsilon \\
&\geq \boldsymbol{V}_i(\delta_i', \delta_{-i}^*) - \varepsilon \\
&> \boldsymbol{V}_i(\delta_i', \delta_{-i}^*) - \left(\boldsymbol{V}_i(\delta_i', \delta_{-i}^*) - \int_T \int_X \overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt)\right) \\
&= \int_T \int_X \overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt) \\
&= \boldsymbol{U}_i(\delta^*).
\end{aligned}$$

This contradicts $\delta^* \in \widetilde{EQ}(u)$. We conclude that $\delta^* \in \widetilde{EQ}(v)$. Since this is true for all $v \in \mathcal{U}_u$, it follows that $\widetilde{IE}(\mathcal{U}_u) = \bigcap_{v \in \mathcal{U}_u} \widetilde{EQ}(v) = \widetilde{EQ}(u)$. □

3.2. Existence

In this section, we demonstrate that RSPM (RUSPM), along with a weak efficiency condition on the payoffs, is sufficient for the normal form game of a Bayesian game (a mixed extension of the normal form) to satisfy better-reply security, as introduced by [3]. Reny shows that better-reply security is a sufficient condition for a compact, quasiconcave game to have a Nash equilibrium. RSPM and RUSPM can therefore be used as alternative conditions for verifying the existence of a pure and behavioral strategy equilibrium, respectively.

Before we present our results, we must first define better-reply security. A player $i \in I$ can *secure* a payoff of α at a strategy profile $s \in S$ if there exists an $\overline{s}_i \in S_i$ and neighborhood $\mathcal{N}(s_{-i})$ of s_{-i} such that $U_i(\overline{s}_i, s_{-i}') \geq \alpha$ for all $s_{-i}' \in \mathcal{N}(s_{-i})$.

Definition 4. *A game G_0 is* better-reply secure *if whenever $(s^*, U^*) \in clG_0$ and s^* is not a Nash equilibrium of G_0, there is some player i that can secure a payoff strictly higher than U_i^* at s^*.*

In order to connect RSPM to better-reply security, we will need to introduce a weak efficiency condition. Let $G_0(t)$ denote the t-section for the game G, that is, $G_0(t) = (X, u(\cdot, t))$, the normal form of the game with a fixed-type profile t. Define the set of actions $\Psi(t)$ for which jointly maximal payoffs are simultaneously feasible for all players. That is, $\Psi(t) = \{x \in X : (x, \overline{\pi}(x,t)) \in clG_0(t)\}$.

Definition 5. *A Bayesian game $G \in \mathcal{G}$ satisfies* random weak efficiency (RWE) *if whenever $s \in S$ is such that $s(t) \in \Psi(t)$ for λ-almost all $t \in T$, then $u(s(t), t) = \overline{\pi}(s(t), t)$ for λ-almost all $t \in T$.*

A game satisfies RWE if, given fixed action and type profiles x and t, all players receive the maximal payoffs $\overline{\pi}(x,t)$ if such an allocation is feasible. As the contest model in Section 4 clarifies, the distinction between all t and λ-almost t is significant in application.

The following theorem shows that RSPM and RWE can be used to verify that the normal form of a Bayesian game is better-reply secure.

Theorem 3. *If $G(\underline{u})$ satisfies RSPM and $G(u)$ satisfies RWE, then $G_0(u)$ is better-reply secure.*

Proof of Theorem 3. The proof is done in two parts. First, we show that if $G(\underline{u})$ satisfies RSPM, then in $G(u)$ each player i can secure a payoff of $\int_T \overline{\pi}_i(s(t), t)\lambda(dt) - \varepsilon$ for any $\varepsilon > 0$ at any strategy profile $s \in S$. Second, we use this security condition along with RWE of $G(u)$ to show that $G_0(u)$ is better-reply secure.

Let $\varepsilon > 0$ and $s \in S$. From RSPM, as shown in (4) of the proof of Lemma 2, each player i has a strategy $s_i' \in S_i$ such that

$$\int_T \underline{u}_i(s_i'(t), s_{-i}(t), t)\lambda(dt) > \int_T \overline{\pi}_i(s(t), t)\lambda(dt) - \frac{\varepsilon}{2}.$$

By construction, each $\underline{u}_i$ is lower semicontinuous. Therefore, there is a neighborhood $\mathcal{N}(s_{-i}(t_{-i})))$ such that

$$\underline{u}_i(s_i(t_i), s'(t), t) > \underline{u}_i(s_i(t_i), s_{-i}(t_{-i}), t) - \frac{\varepsilon}{2}$$

for all $s_{-i}'(t) \in \mathcal{N}(s_{-i}(t_{-i}))$. Define $\mathcal{N}(s_{-i}) = (\mathcal{N}(s_{-i}(t_{-i})))_{t_{-i}\in T_{-i}}$ and observe that $\mathcal{N}(s_{-i})$ is a neighborhood of s_{-i} such that

$$\int_T \underline{u}_i(s_i'(t), s_{-i}'(t), t)\lambda(dt) > \int_T \underline{u}_i(s_i'(t), s_{-i}(t), t)\lambda(dt) - \frac{\varepsilon}{2} \tag{7}$$

for all $s_{-i}' \in \mathcal{N}(s_{-i})$. It follows that

$$\int_T \underline{u}_i(s_i'(t), s_{-i}'(t), t)\lambda(dt) > \int_T \overline{\pi}_i(s(t), t)\lambda(dt) - \varepsilon \tag{8}$$

for all $s_{-i}' \in \mathcal{N}(s_{-i})$. The fact that $u \geq \underline{u}$ then implies that each player i can secure a payoff of $\int_T \overline{\pi}_i(s(t), t)\lambda(dt) - \varepsilon$ in the game $G(u)$.

We now show that $G_0(u)$ is better-reply secure. Let $(s^*, U^*) \in \mathrm{cl}G_0(u)$ and suppose that s^* is not a Nash equilibrium of $G_0(u)$. Observe first that the upper semicontinuity of $\int_T \overline{\pi}(s(t), t)\lambda(dt)$ implies that $U^* \leq \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$. We consider two cases corresponding to whether $U(s^*) = \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$ or $U(s^*) \neq \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$.

Case 1: $U(s^*) = \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$

Since $s^* \notin EQ(u)$, there is a player i with strategy $s_i' \in S_i$ such that $U_i(s_i', s_{-i}^*) > U_i(s^*) \geq U_i^*$. Let $\varepsilon > 0$ be such that $\varepsilon < U_i(s_i', s_{-i}^*) - U_i^*$. From the security condition above, player i can secure a payoff of

$$\begin{aligned} \int_T \overline{\pi}_i(s_i'(t), s_{-i}^*(t), t)\lambda(dt) - \varepsilon &\geq U_i(s_i', s_{-i}^*) - \varepsilon \\ &> U_i^*. \end{aligned}$$

Thus, the game is better-reply secure.

Case 2: $U(s^*) \neq \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$

We first argue that $s^*(t) \notin \Psi(t)$ for some $T' \subset T$ with λ-positive measure. Suppose to the contrary that $s^*(t) \in \Psi(t)$ for λ-almost all t. Then RWE implies that $u(s^*(t), t) = \overline{\pi}(s^*(t), t)$ for λ-almost all t, and thus that $U(s^*) = \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$, a violation of the assumption of this case. We conclude that $s^*(t) \notin \Psi(t)$ for some $T' \subset T$ with λ-positive measure.

Let $s^k \to s^*$ be such that $U(s^k) \to U^*$. Define $u^*(t) = \limsup_k u(s^k(t), t)$, noting that $U^* \leq \int_T u^*(t)\lambda(dt)$. Since $s^*(t) \notin \Psi(t)$, there is a λ-positive measure set $T' \subset T$ of types for which $u^*(t) \neq \overline{\pi}(s^*(t), t)$; since the set of players is finite, there must be at least one player i such that $u_i^*(t) \neq \overline{\pi}_i(s^*(t), t)$ for some $T'' \subset T'$ with λ-positive measure. Further, since $u \leq \overline{\pi}$, this implies that $\int_T u^*(t)\lambda(dt) \neq \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$, and thus that $U^* \neq \int_T \overline{\pi}(s^*(t), t)\lambda(dt)$. It follows that there is some player i such that $U_i(s^*) < \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt)$.

Let $\varepsilon > 0$ be such that $\varepsilon < \int_T \overline{\pi}_i(s^*(t), t)\lambda(dt) - U_i^*$. Again from the security condition above, player i can secure a payoff of $\int_T \overline{\pi}_i(s_i'(t), s_{-i}^*(t), t)\lambda(dt) - \varepsilon > U_i^*$. We conclude that the game $G_0(u)$ is better-reply secure. □

The following theorem extends our analysis to behavioral strategies. Specifically, the following theorem demonstrates that RUSPM and RWE together can be used to show that the mixed extension of the normal form of a Bayesian game is better-reply secure. This is particularly useful since RUSPM and RWE are conditions on the primitives of the Bayesian game, and thus better-reply security and the existence of behavioral strategy equilibrium can be verified without any computations in the mixed extension.

Theorem 4. *If $G(\underline{u})$ satisfies RUSPM and $G(u)$ satisfies RWE, then $\widetilde{G}_0(u)$ is better-reply secure. Thus, $\widetilde{EQ}(u) \neq \emptyset$.*

Proof of Theorem 4. The proof follows the same basic structure as that of Theorem 3. First, we show that if $G(\underline{u})$ satisfies RUSPM, then in $\widetilde{G}_0(u)$ each player i can secure a payoff of $\int_T \int_{X_{-i}} \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt) - \varepsilon$ for any $\varepsilon > 0$ at any strategy profile $\delta \in M$. Second, we use this security condition along with RWE of $G(u)$ to show that $\widetilde{G}_0(u)$ is better-reply secure.

Let $\varepsilon > 0$ and $\delta \in M$. From RUSPM and condition (6) in the proof of Theorem 2, there exists for each player i a strategy $s'_i \in S_i$ such that

$$\int_T \int_{X_{-i}} \underline{u}_i(s'_i(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt) \geq \int_T \int_X \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt) - \frac{\varepsilon}{2}.$$

Next, from Lemma A3 (in the Appendix A), $\int_T \int_{X_{-i}} \underline{u}_i(s'_i(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt)$ is lower semicontinuous in δ_{-i}. As such, there exists a neighborhood $\mathcal{N}(\delta_{-i})$ such that

$$\int_T \int_{X_{-i}} \underline{u}_i(s'_i(t_i), x_{-i}, t_i, t_{-i})\delta'_{-i}(dx_{-i}|t_{-i})\lambda(dt) > \int_T \int_{X_{-i}} \underline{u}_i(s'_i(t_i), x_{-i}, t_i, t_{-i})\delta_{-i}(dx_{-i}|t_{-i})\lambda(dt) - \frac{\varepsilon}{2}$$

for all $\delta'_{-i} \in \mathcal{N}(\delta_{-i})$. Combining these inequalities, we get

$$\int_T \int_{X_{-i}} \underline{u}_i(s'_i(t_i), x_{-i}, t_i, t_{-i})\delta'_{-i}(dx_{-i}|t_{-i})\lambda(dt) > \int_T \int_X \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt) - \varepsilon \quad (9)$$

for all $\delta'_{-i} \in \mathcal{N}(\delta_{-i})$. The fact that $u \geq \underline{u}$ then implies that each player i can secure a payoff of $\int_T \int_{X_{-i}} \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt) - \varepsilon$ in the game $\widetilde{G}_0(u)$.

We now show that $\widetilde{G}_0(u)$ is better-reply secure. Let $(\delta^*, \boldsymbol{U}^*) \in \mathrm{cl}\widetilde{G}_0(u)$ and suppose that and δ^* is not a Nash equilibrium of $\widetilde{G}_0(u)$. Observe first that the upper semicontinuity of $\int_T \int_X \overline{\pi}(x,t)\delta(dx|t)\lambda(dt)$ in δ from Lemma A3 implies that $\boldsymbol{U}^* \leq \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$. We consider two cases corresponding to whether $\boldsymbol{U}(\delta^*) = \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$ or $\boldsymbol{U}^*(\delta^*) \neq \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$.

Case 1: $\boldsymbol{U}(\delta^*) = \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$

Since δ^* is not an equilibrium, there is a player i with strategy $\delta'_i \in S_i$ such that $\boldsymbol{U}_i(\delta'_i, \delta^*_{-i}) > \boldsymbol{U}_i(\delta^*) \geq \boldsymbol{U}^*_i$. Let $\varepsilon > 0$ be such that $\varepsilon < \boldsymbol{U}_i(\delta'_i, \delta^*_{-i}) - \boldsymbol{U}^*_i$. From the security condition above, player i can secure a payoff of

$$\begin{aligned} \int_T \int_X \overline{\pi}_i(x,t)\delta'_i(dx|t)\delta^*_{-i}(dx_{-i}|t)\lambda(dt) - \varepsilon &\geq \boldsymbol{U}_i(s'_i, s^*_{-i}) - \varepsilon \\ &> U^*_i. \end{aligned}$$

Thus, the game is better-reply secure.

Case 2: $\boldsymbol{U}(\delta^*) \neq \int_T \int_X \overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt)$

We first argue that $(x,t) \notin \Psi(t) \times \{t\}$ for some $\lambda \diamond \delta^*$-positive measure subset of $X \times T$. Suppose to the contrary that $(x,t) \in \Psi(t) \times \{t\}$ for $\lambda \diamond \delta^*$-almost all $x \times t$. Then RWE implies that $u(x,t) = \overline{\pi}(x,t)$ for $\lambda \diamond \delta^*$-almost all $x \times t$, and thus that $U(\delta^*) = \int_T \int_X \overline{\pi}(x,t)\delta^*(dx|t)\lambda(dt)$, a violation of the assumption of this case. We conclude that $(x,t) \notin \Psi(t) \times \{t\}$ for some $\lambda \diamond \delta^*$-positive measure subset of $X \times T$.

Define $\mathrm{A}(x,t) = \sum_{i\in I} u_i(x,t)$ and define $\overline{\mathcal{A}}(x,t)$ as we defined $\overline{u}_i$ for the function u_i. Clearly, $\overline{\mathcal{A}}$ is $\mathcal{B}(X)\otimes\mathcal{T}$-measurable and upper semicontinuous in x. Since $u \le \overline{\pi}$ and $(x,t) \notin \Psi(t)\times\{t\}$ for some $\lambda\diamond\delta^*$-positive measure subset of $X\times T$, then it must be that

$$\int_T\int_X \overline{\mathcal{A}}(x,t)\delta^*(dx|t)\lambda(dt) < \int_T\int_X \sum_{i\in I}\overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt).$$

From Lemma A3, $\int_T\int_X \overline{\mathcal{A}}(x,t)\delta^*(dx|t)\lambda(dt)$ is upper semicontinuous in δ. Thus,

$$\begin{aligned}\sum_{i\in I}\boldsymbol{U}_i^* &\le \int_T\int_X \overline{\mathcal{A}}(x,t)\delta^*(dx|t)\lambda(dt)\\ &< \int_T\int_X \sum_{i\in I}\overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt).\end{aligned}$$

it follows that $\boldsymbol{U}_i^* < \int_T\int_X \overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt)$ for some player i.

Let $\varepsilon > 0$ be such that $\varepsilon < \int_T\int_X \overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt) - \boldsymbol{U}_i^*$. Again using the security condition above, player i can secure a payoff of $\int_T\int_X \overline{\pi}_i(x,t)\delta^*(dx|t)\lambda(dt) - \varepsilon > \boldsymbol{U}_i^*$. We conclude that the game $\widetilde{G}_0(u)$ is better-reply secure.

Finally, since the mixed extension of the normal form of the Bayesian game is better-reply secure, Corollary 5.2 of [3] implies that the normal form game has a mixed strategy equilibrium. □

Remark 1. *If a Bayesian game satisfies RUSPM and RWE, then Theorem 4 allows for the application of Corollary 5.2 of [3] to the mixed extension to get the existence of behavioral strategy equilibrium in the Bayesian game. Purification results offer an avenue to apply Theorem 4 (combined with Corollary 5.2 of [3]) to get the existence of pure strategy equilibrium in a Bayesian game without the restrictive assumption of its own payoff quasiconcavity. Ref. [4] explicitly show the conditions for applying purification results adopting the "relative diffuseness" conditions of [13] for a Bayesian game satisfying a uniform payoff security condition. Ref. [14] provide new purification results based on the "decomposable coarser payoff-relevant information" condition.*

4. Contest with Bid Caps

The basic structure of the contest is similar to [15,16] with the addition of incomplete information[5]. The inclusion of bid caps complicates the verification of existence and provides a good illustration of why the "λ-almost all" sufficient conditions are important for application[6].

Consider a contest with a set of players $I = \{1,\ldots,n\}$ and m identical prizes, where $I > m > 1$. Each player i has a space of types $T_i = [\underline{t}_i, \overline{t}_i]$, where $0 \le \underline{t}_i < \overline{t}_i$. Each player i has a valuation of winning denoted by the measurable function $w_i : X_i \times T \to \mathbb{R}$, and player i's valuation of losing is denoted by the measurable function $l_i : X_i \times T \to \mathbb{R}$. Each player i observes their type t_i and picks a score $x_i \in X_i(t_i) = [0, t_i]$. We make the following six assumptions on the primitives of this model:

(i) The common prior λ is absolutely continuous with respect to $\otimes_{i\in I}\lambda_i$, and each λ_i is atomless on T_i.
(ii) For all $t \in T$, $w_i(x_i,t)$ is upper semicontinuous and nonincreasing in x_i.
(iii) $w_i(x_i,t) \ge l_i(x_i,t)$ for all $(x_i,t) \in X_i \times T$.
(iv) $l_i(x_i,t) \le l_i(0,t)$ for all $(x_i,t) \in X_i \times T$.
(v) For each i and t, there exists $r_i(t)$ such that $w_i(r_i(t),t) = l_i(0,t)$, $w_i(x_i,t) > l_i(0,t)$ for all $x_i < r_i(t)$, and $w_i(x_i,t) \le l_i(0,t)$ for all $x > r_i(t)$.
(vi) For every player i and any score x_i, the set of t such that $r_i(t) = x_i$ is λ-measure zero.

Denote the probability of player i winning a prize given the vector of scores $x = (x_1, \ldots, x_n)$, by $P_i : X \times T \to [0,1]$. Formally,

$$P_i(x,t) = \begin{cases} 0 & \text{if } x_i < x_j \text{ for } m \text{ or more players } j \neq i, \\ 1 & \text{if } x_i > x_j \text{ for } n-m \text{ or more players } j \neq i, \\ \alpha_i(x,t) \in [0,1] & \text{otherwise.} \end{cases}.$$

Thus, we can write the payoff of player i for the fixed action profile x as

$$u_i(x,t) = P_i(x,t)w_i(x_i,t_i) + (1 - P_i(x,t))l_i(x_i,t_i).$$

The set of payoffs for the class of games $\mathcal{U}$ is the set of all payoffs u for all α such that $\alpha_i(x,t) \in [0,1]$ and $\sum_{i=1}^n P_i(x,t) \leq m$. Notice for every u in the class $\mathcal{U}$, $\underline{u}$ is the payoff section in which for all i, x, and t, $\alpha_i(x,t) = 0$. Finally, assume that for any profile x with a relevant tie ($x_i = x_j$ for some $i \neq j$ with no more than $m-1$ higher bidders and no more than $n-m-1$ lower bidders), if $w_i(x_i,t) = l_i(x_i,t)$ and the number K of players k with either $x_k > x_i$ or $x_k = x_i$ and $w_k(x_k,t) > l_k(x_k,t)$ is such that $K \leq m$, then $P_k(x_k,t) = 1$ for each such player k. That is, if at a tie, one bidder prefers winning to losing, and another prefers losing to winning, the player who prefers winning to losing must "win" that tie and prizes must be allocated at least to players that prefer to win, unless more than m players bid higher or tie and prefer to win.

To apply our existence result to a payoff selection in the class $\mathcal{U}$ (any u with a particular tie-breaking rule), we need to add an additional payoff to the class and leverage the invariance results. This is because we are unable to show that there is a $u \in \mathcal{U}$ that satisfies RWE. Define the measurable function l_i' such that $l_i'(x,t) = l_i(x,t)$ for all $x \leq r_i(t)$, and $l_i'(x,t) < w_i(x,t)$ for all $x > r_i(t)$. Then, for all i,

$$u_i'(x,t) = P_i(x,t)w_i(x_i,t_i) + (1 - P_i(x,t))l_i'(x_i,t_i).$$

Then the class of games $\mathcal{U}'$ is defined by the payoffs u' at all tie-breaking rules α. Notice for the class $\mathcal{V} = \mathcal{U} \cup \mathcal{U}'$, $\underline{u}' = \underline{\pi}$.

Proposition 1. *In the all-pay contest with incomplete information, $\underline{\pi}$ satisfies RUSPM and all payoffs u satisfying the assumptions satisfy RWE.*

Proof of Propositon 1. First, we argue that the game $G(\underline{\pi})$ (this is the game with $\underline{u}'$ and for all i, $\alpha_i(x,t) = 0$ for all x and t) satisfies RUSPM. Note that the set of discontinuity points for each player i's payoff is the same for all tie-breaking rules. At any tie for player i at a score x for every type t_i such that $t_i \neq x$, the sequence $x^k = x + 1/k$ gives $\lim_k \underline{\pi}_i(x^k, x_{-i}, t) = w_i(x,t) = \overline{\pi}_i(x, x_{-i}, t)$. Since there is a unique type $t_i = x_i$ with the violation and the measure λ_i is non-atomic, the same sequence can payoff match $\overline{\pi}_i(x, x_{-i}, t)$ at x for all $x_{-i} \in X_{-i}$ and λ-almost all types t.

Second, we show that the game $G(u)$ satisfies RWE. Observe that $u_i(x,t) = \overline{\pi}_i(x,t)$ except possibly when there is a relevant tie. Suppose that $\underline{\pi}_i(x, x_{-i}, t) < \overline{\pi}_i(x, x_{-i}, t)$. Then $w_i(x_i,t) > l_i(x_i,t)$, and there must be a relevant tie at $x_i = z$, with $x_j > z$ for no more than $m-1$ other bidders j and $x_j < z$ for no more than $n-m-1$ other bidders j. Let K be the number of players k with either $x_k > z$ or $x_k = z$ and $w_k(x_k,t) > l_k(x_k,t)$ and let I_K be the set of these players. We consider two cases corresponding to whether the $K > m$ or $K \leq m$. If $K > m$, then $(x, \overline{\pi}(x,t)) \notin clG_0(t)$ since $\overline{\pi}(x,t)$ involves more than m prizes being allocated. Else, if $K \leq m$, then $i \in I_K$, so by assumption, $P_i(x,t) = 1$, so $u_i(x, x_{-i}, t) = \overline{\pi}_i(x, x_{-i}, t)$. It follows that further, since $P_k(x,t) > P_j(x,t)$ for any tied player k with $w_k(x_k,t) > \Psi(t) = \{x \in X : (x, \overline{\pi}(x,t)) \in clG_0(t)\}$Notice that for any player i the fixed t type, $\underline{\pi}_i(x,t) = \overline{\pi}_i(x,t)$ can only happen at a tie with player i involved at $x_i = r_i(t)$. This is because $\sum_{i=1}^n P_i(x,t) \leq m$ guarantees that at most m players can get the full upper bound payoff of winning a prize, for sure. For player i, the payoff $\underline{\pi}_i(x,t)$ at $r_i(t)$

is the same for winning or losing a prize only at $r_i(t)$. For any score x_i, the set of t such that $r_i(t) = x_i$ is λ-measure zero, that is $r_i(t) \neq x$ for λ-almost all types t. Therefore, for any sequence $s^k \to s$, $\lim_k u'(s^k(t), t) < \overline{\pi}(s(t), t)$ for λ-almost all types t. This makes the if statement in the definition of RWE never apply. Thus, $G(u)$ satisfies RWE. □

Based on the application of Theorem 4, $\widetilde{G}_0(\underline{\pi})$ is better replay secure, and thus we know $G(\underline{\pi})$ has a behavioral strategy equilibrium. Theorem 2 makes this behavioral strategy equilibrium an equilibrium for the entire class of games with all the payoff functions in $\mathcal{V}$.

Author Contributions: All phases of creating this research were done collaboratively by B.A.A. and J.J.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: No data was used for this research.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Technical Lemmas

The proofs of the lemmas below are adapted from results shown as part of the proof of Theorem 1 in [8].

Lemma A1. *For all i, the functions $\overline{u}_i, \underline{u}_i, \overline{\pi}_i$, and $\underline{\pi}_i$ are $\mathcal{B}(X) \otimes \mathcal{T}$-measurable.*

Proof of Lemma A1. We only do the proof for the lower envelope since the proof for the upper envelope follows the same lines. Since X is a compact metric space, it is second countable (see [12] Proposition 25, p. 204) and we can find a countable base $\{\mathcal{V}_m\}_{m\geq 1}$ for X. Let

$$\underline{u}_i^m(x,t) = \begin{cases} \inf_{x' \in V_m} u_i(x', t_i, t_{-i}) & \text{if } x \in \mathcal{V}_m \\ 0 & \text{o.w.} \end{cases}.$$

Clearly, $\underline{u}_i^m(\cdot, t)$ is lower semicontinuous on X for each fixed $t \in T$ and $m \geq 1$.

To show that $\underline{u}_i^m$ is jointly measurable we show that, for any $c \geq 0$, the set

$$\{(x,t) \in X \times T : \underline{u}_i^m(x_{-i}, t) < c\}$$

is $\mathcal{B}(X) \otimes \mathcal{T}$-measurable. Since u_i is jointly measurable and g_i is $\mathcal{T}_i$-measurable, the set

$$\{(x,t) \in \mathcal{V}_m \times T : u_i(x, t_i, t_{-i}) < c\}$$

is $\mathcal{B}(X) \otimes \mathcal{T}$-measurable. By the Projection Theorem the projection of this set on T, denoted as T_m, is $\mathcal{T}$-measurable[7]. Notice that

$$\{(x,t) \in X \times T : \underline{u}_i^m(x,t) < c\} = (\mathcal{V}_m \times T_m) \cup (\mathcal{V}_m^c \times T),$$

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

Since $\underline{u}_i(x,t) = \sup_{m\geq 1} \underline{u}_i^m(x,t)$, $\underline{u}_i(x,t)$ is the pointwise supremum of a sequence of lower semicontinuous functions, which is also lower semicontinuous (Theorem 3.1 in [3]). In addition $\underline{u}_i$ is the supremum of a sequence of $\mathcal{B}(X) \otimes \mathcal{T}$-measurable functions, which is also $\mathcal{B}(X) \otimes \mathcal{T}$-measurable. The proof for $\underline{\pi}_i$ is exactly the same as for $\underline{u}_i$ since the fact that $\mathcal{U}$ is countable is sufficient for $\inf_{u \in \mathcal{U}} \underline{u}_i(x,t)$ to be $\mathcal{B}(X) \otimes \mathcal{T}$-measurable. □

Lemma A2. *For any $\delta \in M$, player $i \in I$, and $\mathcal{B}(X) \otimes \mathcal{T}$-measurable u, there is a $\mathcal{T}$-measurable selection g_i' such that*

$$\int_T \int_{X_{-i}} u_i(g_i'(t_i), x_{-i}, t_i, t_{-i}) \delta_{-i}(dx_{-i}|t_{-i}) \lambda(dt) \geq \int_T \int_X u_i(x,t) \delta(dx|t) \lambda(dt)$$

Proof of Lemma A2. Fix a behavioral strategy profile $\delta \in M$, player i, and $\varepsilon > 0$. Let $S_i : T_i \to X_i$ be a correspondence defined by

$$S_i(t_i) = \left\{ x_i \in X_i : \begin{array}{l} \int_{T_{-i}} \int_{X_{-i}} u_i(x_i, x_{-i}, t_i, t_{-i}) \psi(t_i, t_{-i}) \delta_{-i}(dx_{-i}|t_{-i}) \otimes_{j \neq i} \lambda_i(dt_{-t}) \\ \geq \int_{T_{-i}} \int_X u_i(x_i, x_{-i}, t_i, t_{-i}) \psi(t_i, t_{-i}) \delta(dx|t_i t_{-i}) \otimes_{j \neq i} \lambda_i(dt_{-t}) \end{array} \right\}.$$

Clearly, for each t_i, $S_i(t_i) \neq \emptyset$. Since u_i is jointly measurable, and δ and ψ are measurable, the correspondence S_i has a $\mathcal{B}(X_i) \otimes \mathcal{T}_i$-measurable graph. By Aumann's Measurable Selection Theorem[8], S_i has a $\mathcal{T}_i$-measurable selection g_i' such that

$$\int_T \int_{X_{-i}} u_i(g_i'(t_i), x_{-i}, t_i, t_{-i}) \delta_{-i}(dx_{-i}|t_{-i}) \lambda(dt) \geq \int_T \int_X u_i(x,t) \delta(dx|t) \lambda(dt).$$

□

Lemma A3. *Let $f, g : X \times T \to \mathbb{R}$ be such that f (g) is upper (lower) semicontinuous in x. Then $\int_T \int_X f(x,t)\delta(dx|t)\lambda(dt)$ is upper semicontinuous in δ and $\int_T \int_X g(x,t)\delta(dx|t)\lambda(dt)$ is lower semicontinuous in δ. In particular, this implies that for each player i, the functions $\int_T \int_X \overline{u}_i(x,t)\delta(dx|t)\lambda(dt)$ and $\int_T \int_X \overline{\pi}_i(x,t)\delta(dx|t)\lambda(dt)$ are upper semicontinuous in δ, while the payoffs $\int_T \int_X \underline{u}_i(x,t)\delta(dx|t)\lambda(dt)$ and $\int_T \int_X \underline{\pi}_i(x,t)\delta(dx|t)\lambda(dt)$ are lower semicontinuous in δ.*

Proof of Lemma A3. Define a function $H_i^l : M \mapsto \mathbb{R}$ as follows: for any $\delta \in M$,

$$H_i^l(\delta) = \int_T \int_X \underline{u}_i(x,t)\psi(t) \otimes_{j \in I} \delta_j(dx_j|t_j) \otimes_{i \in I} \lambda_i(dt).$$

Let $\phi(x,t) = \int_{T_i} \underline{u}_i(x,t)\psi(t)\lambda_i(dt_i)$. Since $\underline{u}_i(x,t)\psi(t)$ is lower semicontinuous in x, jointly measurable, and integrably bounded, ϕ is also lower semicontinuous in x, jointly measurable, and integrably bounded. By Lemma 3 in [8], the functional $\delta \to \otimes_{j \in I}\delta_j$ from M to $\widetilde{M}$ is continuous. Then by Lemma 2 in [8], the functional

$$v \to \int_{T_{-i}} \int_X \phi(x, t_{-i}) v(dx|t_{-i}) \lambda_{-i}(dt_{-i}).$$

is lower semicontinuous on $\widetilde{M}$. Since H_i^l is a composition of these two functionals, it is lower semicontinuous. As a result for any $\varepsilon > 0$, there is an open neighborhood $\mathcal{N}(\delta) \subseteq M$ of δ such that for any $\delta' \in \mathcal{N}(\delta)$,

$$\int_T \int_X \underline{u}_i(x,t)\psi(t)\delta(dx|t) \otimes_{i \in I} \lambda_i(dt) \geq \int_T \int_X \underline{u}_i(x,t)\psi(t)\delta'(dx|t) \otimes_{i \in I} \lambda_i(dt) - \varepsilon.$$

That is,

$$\int_T \int_X \underline{u}_i(x,t)\delta(dx|t)\lambda(dt) \geq \int_T \int_X \underline{u}_i(x,t)\delta'(dx|t)\lambda(dt) - \varepsilon.$$

□

Notes

1. There is important previous literature focused on the existence of equilibrium in Bayesian games with continuous payoffs that includes the two seminal contributions of [1,2].
2. A game is better-reply secure if for every nonequilibrium strategy profile x^* and every limiting payoff vector u^* at x^*, there is a player i that has a strategy that gives a payoff that is strictly higher than u_i^* even when other players deviate slightly from x^*.
3. Introduced by [3] as one of two sufficient conditions for better reply security, a game satisfies reciprocal upper semicontinuity if whenever a strategy-payoff pair (x^*, u^*) is in the closure of the game and $u(x^*) \leq u^*$, then $u(x^*) = u^*$.
4. Since A is a compact metric space, by [12] Proposition 25 p. 204, it is second countable.
5. The literature addressing related contests includes [17–23].

[6] Ref. [24] show that asymmetric bid caps create an existence problem for most tie-breaking rules in all-pay contests with complete information. Two other recent papers, [25,26], include symmetric bidding constraints in all-pay auctions with incomplete information.

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

[8] 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 $\mathcal{S} \otimes B(X)-$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 $s \in S$.

References

1. Milgrom, P.R.; Weber, R.J. Distributional Strategies for Games with Incomplete Information. *Math. Oper. Res.* **1985**, *10*, 619–632. [CrossRef]
2. Balder, E.J. Generalized Equilibrium Results for Games with Incomplete Information. *Math. Oper. Res.* **1988**, *13*, 191–376. [CrossRef]
3. Reny, P.J. On the existence of pure and mixed strategy Nash equilibria in discontinuous games. *Econometrica* **1999**, *67*, 1029–1056. [CrossRef]
4. He, W.; Yannelis, N. Discontinuous games with asymmetric information: An extension of Reny's existence theorem. *Games Econ. Behav.* **2015**, *91*, 26–35. [CrossRef]
5. Carbonell-Nicolau, O.; McLean, R.P. On the Existence of Nash Equilibrium in Bayesian Games. *Math. Oper. Res.* **2018**, *43*, 100–129. [CrossRef]
6. Monteiro, P. K.; Page, F.H. Uniform Payoff Security and Nash Equilibrium in Compact Games.*J. Econ. Theory* **2007** *134*, 566–575. [CrossRef]
7. Prokopovych, P.; Yannelis, N. On the existence of mixed strategy Nash equilibria. *J. Mathatical Econ.* **2014**, *52*, 87–97. [CrossRef]
8. He, W.; Yannelis, N. Existence of Equilibrium in Discontinuous Bayesian Games. *J. Econ. Theory* **2016**, *162*, 181–194. [CrossRef]
9. Allison, B.A.; Lepore, J.J. Verifying payoff security in the mixed extension of discontinuous games. *J. Econ. Theory* **2014**, *152*, 291–303. [CrossRef]
10. Allison, B.A.; Bagh, A.; Lepore, J.J. Invariant Equilibria and Classes of Games. *Games Econ. Behav.* **2022**, *132*, 448–462. [CrossRef]
11. Allison, B.A.; Bagh, A.; Lepore, J.J. Sufficient conditions for weak reciprocal upper semi-continuity in mixed extensions of games. *J. Math. Econ.* **2018**, *74*, 99–107. [CrossRef]
12. Royden, H.L.; Fitzpatrick, P.M. *Real Analysis*, 4th ed.; Prentice Hall: Boston, MA, USA, 2010.
13. He, W.; Sun, X. On the diffuseness of incomplete information games. *J. Mathatical Econ.* **2014**, *54*, 131–137. [CrossRef]
14. He, W.; Sun, X.; Sun, Y.; Zeng, Y. Characterization of equilibrium existence and purification in general Bayesian games. *arXiv* **2021**, arXiv:2106.08563 .
15. Siegel, R. All-Pay Contests. *Econometrica* **2009**, *77*, 71–92.
16. Siegel, R. Asymmetric Contests with Conditional Investments. *Am. Econ. Rev.* **2010**, *100*, 2230–2260. [CrossRef]
17. Amann, E.; Leininger, W. Asymmetric all-pay auctions with incomplete information: The two player case. *Games Econ. Behav.* **1996**, *14*, 1–18. [CrossRef]
18. Lizzeri, A.; Persico, N. Uniqueness and existence of equilibrium in auctions with a reserve price. *Games Econ. Behav.* **2000**, *30*, 83–114. [CrossRef]
19. Parreiras, S.O.; Rubinchik, A. Contests with three or more heterogeneous agents. *Games Econ. Behav.* **2010**, *68*, 703–715. [CrossRef]
20. Siegel, R. Asymmetric all-pay auctions with interdependent valuations. *J. Econ. Theory* **2014**, *153*, 684–702. [CrossRef]
21. Rentschler, L.; Turocy, T.L. Two-bidder all-pay auctions with interdependent valuations, including the highly competitive case. *J. Econ. Theory* **2016**, *163*, 435–466. [CrossRef]
22. Lu, J.; Parreiras, S.O. Monotone Equilibrium of two-bidder all-pay auction Redux. *Games Econ. Behav.* **2017**, *104*, 78–91. [CrossRef]
23. Einy, E.; Goswami, M.P.; Haimanko, O.; Orzach, R.; Sela, A. Common-value all-pay auctions with asymmetric information. *Int. J. Game Theory* **2017**, *46*, 79–102. [CrossRef]
24. Pastine, I.; Pastine, T. *Equilibrium Existence and Expected Payoffs in All-Pay Contests with Caps*; Mimeo: Berlin, Germany , 2019.
25. Einy, E.; Haimanko, O.; Orzach, R.; Sela, A. Common-value all-pay auctions with asymmetric information and bid caps. *Int. J. Game Theory* **2016**, *45*, 63–88. [CrossRef]
26. Chen, B. On the Effects of Bid Caps in All-pay Auctions. *Econ. Lett.* **2019**, *177*, 60–65. [CrossRef]