Perfect Information Games
with Upper Semicontinuous Payoffs

Roger A. Purves, William D. Sudderth
University of California, Berkeley, University of Minnesota, Minneapolis

Abstract. It was shown by Flesch et al (2010) that every n-person, perfect information game with lower semicontinuous payoffs has a subgame perfect ϵ-equilibrium in pure strategies. Here the same is proved when the payoffs are upper semicontinuous.

1 Introduction

The games treated here are sequential with perfect information, infinitely many stages, and no chance moves. There is a finite number of players and one of them is assigned to choose an action at every stage of the game. The payoff to each player is a function of the infinite sequence of actions chosen by the players.

Flesch et al (2010) showed that, if the payoff functions are bounded and lower semicontinuous, then such a game always has a pure, subgame perfect ϵ-equilibrium for $\epsilon > 0$. Here we prove the same result for bounded, upper semicontinuous payoffs. Moreover, Example 3 in Solan and Vieille (2003) shows that if one player has a lower semicontinuous payoff and another player has an upper semicontinuous payoff, then such an equilibrium need not exist.

The proof of Flesch et al uses an intricate, transfinite construction. The proof given here uses techniques from the Dubins and Savage (1976) theory of gambling. It also uses certain methods from Secchi and Sudderth (2001), who studied stochastic games with upper semicontinuous payoffs.

See the introduction to Flesch et al (2010) for historical details and references for perfect information games.

2 The Model and Main Result

The model here and much of the notation will be the same as in Flesch et al (2010).

Let $I = \{1, 2, \ldots, \nu\}$ be the set of players and let A, an arbitrary non-empty set, be the set of actions for a sequential game. Denote by $\mathcal{H}$ the set of all finite sequences of actions including the empty sequence e. Elements of $\mathcal{H}$ are called *histories*. A given function $i : \mathcal{H} \mapsto I$ assigns an active player to each history.

Play begins at the empty history e and the player $i(e)$ selects an action a_1 from A. If in the first n stages the players have selected the history $h_n = (a_1, a_2, \ldots, a_n) \in \mathcal{H}$, then player $i(h_n)$ selects the next action a_{n+1} so that the next history is the concatenation $h_{n+1} = h_n a_{n+1} = (a_1, a_2, \ldots, a_n, a_{n+1})$. The game is one of *perfect information* in that the player selecting a_{n+1} knows the history h_n at every stage n. (When a history $h = (a)$ has only one coordinate, we often write a rather than (a).)

By continuing to select actions at every stage the players generate an infinite history or *play* $p = (a_1, a_2, \ldots) \in A^{\mathbb{N}}$. Each player $j \in I$ has a bounded payoff function $u^j : A^{\mathbb{N}} \mapsto \mathbb{R}$ and receives $u^j(p)$ when the play is p.

For $j \in I$, let

$$\mathcal{H}^j \;=\; i^{-1}(j) \;=\; \{h \in \mathcal{H} \,:\, i(h) = j\}$$

be the set of finite histories where j is the active player.

A *(pure) strategy* σ^j for player j is a mapping $\sigma^j : \mathcal{H}^j \mapsto A$. A ν-tuple $\sigma = (\sigma^j)_{j \in I}$ consisting of a strategy for each player is called a *profile*. Every profile σ determines a unique play $p = p(\sigma) = (a_1, a_2, \ldots)$ where

$$a_1 = \sigma^{i(e)}(e) \qquad \text{and} \qquad a_{n+1} = \sigma^{i(a_1, a_2, \ldots, a_n)}(a_1, a_2, \ldots, a_n),\ n \geq 1.$$

The payoff to player j from the profile σ is $u^j(\sigma) = u^j(p(\sigma))$. Let $u = (u^j)_{j \in I}$ be the vector of payoff functions, and write $\mathcal{G}(u, i)$ for the game with payoff functions u and assignment function i.

Let σ be a profile and, for $j \in I$, write σ^{-j} for the vector $(\sigma^k)_{k \in I \setminus \{j\}}$ of strategies for the set of all players except j.

Definition 1. *For $\epsilon \geq 0$, a profile $\sigma_* = (\sigma_*^j)_{j \in I}$ is an* ϵ-equilibrium *if $u^j(\sigma_*) \geq u^j(\sigma_*^{-j}, \sigma^j) - \epsilon$ for every player j and every strategy σ^j for player j.*

Here the interest is in the stronger notion of a subgame perfect ϵ-equilibrium; that is, a profile that induces an ϵ-equilibrium in every subgame. To define a subgame of $\mathcal{G}(u, i)$, let $h = (a_1, \ldots, a_n)$ be a finite history and let $uh = (u^j h)_{j \in I}$, where for each j the function $u^j h : A^{\mathbb{N}} \mapsto \mathbb{R}$ is the h-section of u^j defined, for $p = (b_1, b_2, \ldots) \in A^{\mathbb{N}}$ by

$$(u^j h)(p) \;=\; u^j(hp) \;=\; u^j(a_1, \ldots, a_n, b_1, b_2, \ldots).$$

Similarly, let $ih : \mathcal{H} \mapsto I$ be the h-section of i defined by $(ih)(h') = i(hh')$ for $h' \in \mathcal{H}$. Here hh' denotes the history consisting of the elements of h followed by those of h'.

The *subgame* $\mathcal{G}(u, i|h)$ of $\mathcal{G}(u, i)$ associated with the finite history $h = (a_1, \ldots, a_n)$ is now defined to be the game $\mathcal{G}(uh, ih)$. Intuitively, this subgame is just the continuation of the original game after the history h has occurred. It is also natural to view this subgame as the *conditional game* given the history h.

Let $\sigma = (\sigma^j)_{j \in I}$ be a profile for the original game $\mathcal{G}(u, i)$ and let $h = (a_1, \ldots, a_n) \in \mathcal{H}$. The *conditional profile* $\sigma[h] = (\sigma^j[h])_{j \in I}$ is the profile consisting of the *conditional strategies* $\sigma^j[h]$ where, for $h' \in \mathcal{H}$ and $hh' \in \mathcal{H}^j$, $\sigma^j[h](h') = \sigma^{i(hh')}(hh')$ for all j. Thus $\sigma[h]$ chooses the same action at h' that σ chooses at hh' and determines the play $p(\sigma[h]) = (b_1, b_2, \ldots)$ where

$$b_1 = \sigma^{i(h)}(h) = \sigma^{i(a_1, \ldots, a_n)}(a_1, \ldots, a_n)$$

and

$$b_{k+1} = \sigma^{(ih)(b_1, \ldots, b_k)}(h(b_1, \ldots, b_k)) = \sigma^{i(a_1, \ldots, a_n, b_1, \ldots, b_k)}(a_1, \ldots, a_n, b_1, \ldots, b_k)$$

for $k \geq 1$.

Here is a simple equality that will be used later. Let σ be a profile and let $a_1 = \sigma^{i(e)}(e)$ be its initial action. Then

$$u^j(\sigma) \;=\; (u^j a_1)(\sigma[a_1])), \; j \in I. \tag{2.1}$$

To verify this equality, let $p(\sigma) = (a_1, a_2, \ldots)$. The left side is $u^j(a_1, a_2, \ldots)$. But $p(\sigma[a_1]) = (a_2, \ldots)$. So the right side is $(u^j a_1)(a_2, \ldots) = u^j(a_1, a_2, \ldots)$.

It is also useful to observe that a profile σ can be specified by naming its first action $\sigma^{i(e)}(e)$ and the collection of all the conditional profiles $\sigma[a]$ for $a \in A$.

The *conditional payoff* to player j from the profile σ given the history h is

$$u^j(\sigma|h) \;=\; (u^j h)(\sigma[h]) \;=\; (u^j h)(p(\sigma[h])).$$

Definition 2. *Let $\epsilon \geq 0$. The profile $\sigma_* = (\sigma_*^j)_{j \in I}$ is a* subgame perfect ϵ-equilibrium *for the game $\mathcal{G}(u, i)$ if, for every $h \in \mathcal{H}$, the conditional profile $\sigma_*[h] = (\sigma_*^j[h])_{j \in I}$ is an ϵ-equilibrium for the subgame $\mathcal{G}(u, i|h)$. When σ_* is a subgame perfect 0-equilibrium, we say simply that it is* subgame perfect.

The set of actions A is given the discrete topology and the space $A^{\mathbb{N}}$ the corresponding product topology. A function $f : A^{\mathbb{N}} \mapsto \mathbb{R}$ is *upper semicontinuous* if, for every real number r, the set $\{p \in A^{\mathbb{N}} : f(p) \geq r\}$ is closed. Here now is the main result of the paper.

Theorem 1. *If the payoff function u^j is bounded and upper semicontinuous for every player $j \in I$, then the game $\mathcal{G}(u, i)$ has a subgame perfect ϵ-equilibrium for every $\epsilon > 0$.*

The rest of the paper is devoted to the proof of Theorem 1. The next section presents a useful lemma based on the stop rule methods of Dubins and Savage (1976). Section 4 shows how to approximate an upper semicontinuous function by a finite sum of indicators of nested closed sets. The proof of Theorem 1 is completed in section 5.

3 A Stop Rule Lemma

A *stop rule* is a function $t : A^{\mathbb{N}} \mapsto \{0, 1, \ldots\}$ such that, given plays p, p' in $A^{\mathbb{N}}$, if $t(p) = n$ and p and p' agree in their first n coordinates, then $t(p') = n$. This definition of a stop rule agrees with the more conventional one that requires that, for all n, the set $\{p : t(p) \leq n\}$ belong to the sigma-field generated by the first n coordinate functions on $A^{\mathbb{N}}$. However, notice that we require that $t(p)$ be finite for *all* p.

It follows from the definition that a stop rule is either everywhere strictly positive or is identically equal to 0. We write $\mathbf{0}$ for the identically zero stop rule.

If t is a stop rule, $a \in A$, and t is not $\mathbf{0}$, then $t(a, a_1, a_2, \ldots) \geq 1$ for all $p = (a_1, a_2, \ldots)$ and the function $t[a]$ defined on $A^{\mathbb{N}}$ by

$$t[a](p) \;=\; t[a](a_1, a_2, \ldots) \;=\; t(a, a_1, a_2, \ldots) - 1$$

is easily seen to be a stop rule itself. It is called the *conditional stop rule* given a.

Dubins and Savage (1976) proved many results with a technique that might be called *stop rule induction*. Here is a formalization of the method.

Lemma 1. *Let $\Phi(t)$ be a proposition for every stop rule t. Assume (a) $\Phi(\mathbf{0})$ holds and (b) if t is not $\mathbf{0}$ and $\Phi(t[a])$ holds for all $a \in A$, then $\Phi(t)$ holds. Then $\Phi(t)$ holds for all stop rules t.*

This lemma is Theorem 2.3.1, page 10, in Maitra and Sudderth (1996).

Associated to every stop rule t is the mapping $h_t : A^{\mathbb{N}} \mapsto \mathcal{H}$ defined for $p = (a_1, a_2, \ldots) \in A^{\mathbb{N}}$ by $h_t(p) = (a_1, a_2, \ldots, a_{t(p)})$. Thus $h_t(p)$ is the history consisting of the first $t(p)$ coordinates of p. If $t = \mathbf{0}$, then $h_t(p) = e$ for all p.

Lemma 2. *Assume that each of the payoff functions $u^j, j \in I$, has a finite range. Let t be a stop rule and suppose that, for every history h in the range of h_t, there exists a subgame perfect equilibrium σ_h for the conditional game $\mathcal{G}(u, i|h)$. Then there is a subgame perfect equilibrium σ_* for $\mathcal{G}(u, i)$.*

Proof. The proof is an application of Lemma 1 in which $\Phi(t)$ is the assertion of Lemma 2. So it suffices to verify conditions (a) and (b) of Lemma 1.

If $t = \mathbf{0}$, then the range of h_t is the singleton e, and , by hypothesis, there is a subgame perfect equilibrium for the conditional game $\mathcal{G}(u, i|e)$. But this conditional game is just the original game $\mathcal{G}(u, i)$. So (a) holds.

To check (b), suppose that $t \geq 1$, and assume that the assertion holds for the conditional stop rule $t[a]$ for every $a \in A$. Suppose $h = h_{t[a]}(p)$ for some play p and action a. Thus h is in the range of $h_{t[a]}$. Now

$$ah \;=\; ah_{t[a]}(p) \;=\; h_t(ap)$$

is in the range of h_t. By hypothesis, the conditional game $\mathcal{G}(u, i|ah)$ has a subgame perfect equilibrium σ_{ah}. However, the game $\mathcal{G}(u, i|ah)$ is the same as $\mathcal{G}(ua, ia|h)$. (Indeed, both are the same as $\mathcal{G}(uah, iah)$.) So, for every h in the range of $h_{t[a]}$, there is a subgame perfect equilibrium, namely σ_{ah}, for the game $\mathcal{G}(ua, ia|h)$. Thus, by the inductive hypothesis, there is, for every a, a subgame perfect equilibrium $\tilde{\sigma}_a$ for the game $\mathcal{G}(ua, ia)$.

To define the profile σ_*, first select the first action $a_1^* = \sigma_*^{i(e)}(e)$ to be the action a that maximizes $(u^{i(e)}a)(\tilde{\sigma}_a)$. (This expression achieves a maximum because of the assumption that payoff functions have a finite range.) Next define the conditional profile $\sigma_*[a]$ to be $\tilde{\sigma}_a$ for every a. It remains to be checked that σ_* is a subgame perfect equilibrium for $\mathcal{G}(u, i)$.

Consider a deviation by player j (say) from σ_*^j to another strategy σ^j and let σ be the profile $(\sigma_*^{-j}, \sigma^j)$. If player j moves first; that is, if $i(e) = j$, let $a_1 = \sigma^j(e)$ be the first move. Then use formula (2.1) and calculate:

$$u^j(\sigma) = (u^j a_1)(\sigma[a_1]) \leq (u^j a_1)(\sigma_*[a_1]) \leq (u^j a_1^*)(\sigma_*[a_1^*]) = u^j(\sigma_*).$$

Here the two equalities are instances of (2.1), the first inequality holds because $\sigma_*[a_1] = \tilde{\sigma}_{a_1}$ is subgame perfect for the game $\mathcal{G}(ua_1, ia_1)$, and the second inequality holds because of the choice of a_1^*. If player j does not move first, then the first move a_1^* is unaffected by the deviation and the calculation is shortened to

$$u^j(\sigma) = (u^j a_1^*)(\sigma[a_1]) \leq (u^j a_1^*)(\sigma_*[a_1^*]) = u^j(\sigma_*).$$

So far it has only been established that σ_* is an equilibrium for $\mathcal{G}(u, i)$.

Consider now a subgame $\mathcal{G}(uh, ih)$ where $h = (a_1, \ldots, a_n)$ with $n \geq 1$. To see that $\sigma_*[h]$ is an equilibrium for this subgame, note that $\sigma_*[h] = \tilde{\sigma}_{a_1}[(a_2, \ldots, a_n)]$ and $\tilde{\sigma}_{a_1}$ is subgame perfect for $\mathcal{G}(ua_1, ia_1)$.In particular, $\tilde{\sigma}_{a_1}[(a_2, \ldots, a_n)]$ is an equilibrium for $\mathcal{G}(ua_1, ia_1|(a_2, \ldots, a_n)) = \mathcal{G}(uh, ih)$. This completes the proof.

□

One can also prove a variation on Lemma 2 for subgame perfect ϵ-equilibria without the assumption that the payoff functions have finite range.

Remark 1. *The game $\mathcal{G}(u, i)$ is said to be* determined *if there is a stop rule t such that the values of the payoffs $u^j(p), j \in I$, depend only on $h_t(p)$ for every play p. Equivalently, the sections $u^j h$ are constant functions for h in the range of h_t. An easy corollary of Lemma 2 is that determined games have subgame perfect equilibria if the payoff functions have finite ranges. (They have subgame perfect ϵ-equilibria in general.)*

Stochastic games that are determined are used by Maitra and Sudderth (2007), where they are called finitary *games. Blackwell (1981) shows how to define the class of Borel subsets of the real line using certain two-person, zero-sum determined games. He remarks that the classical construction of the Borel sets uses the ordinals and that he has substituted stop rules. Likewise we use stop rules whereas Flesch et al (2010) used the ordinals.*

4 A Reduction to Simple Payoff Functions

It was assumed in Theorem 1 that the payoff functions $u^j, j \in I$ are bounded and upper semicontinuous. Clearly, there is no real loss of generality in assuming, as we now do, that the range of each u^j is contained in the unit interval $[0, 1]$. Our object in this section is to show that it suffices to prove Theorem 1 when each u^j has the special form

$$u^j_m = 1_{C_{j,1}} + 1_{C_{j,2}} + \cdots + 1_{C_{j,m}}, \tag{4.1}$$

where, for each $j \in I$, the sets $C_{j,k}, k = 1, \ldots, m$ are closed subsets of $A^{\mathbb{N}}$ that are nested in the sense that

$$C_{j,1} \supseteq C_{j,2} \supseteq \cdots \supseteq C_{j,m}. \tag{4.2}$$

(For sets $C \subseteq A^{\mathbb{N}}$, 1_C denotes the indicator function that equals 1 on C and 0 on the complement of C.)

To see that this simplification is possible, let m be a positive integer, and, for each $j \in I$ and $k = 1, \ldots, m$, define

$$C_{j,k} = \{p \in A^{\mathbb{N}} : u^j(p) \geq \frac{k-1}{m}\}.$$

Notice that, for every $j \in I$,

$$C_{j,1} = A^{\mathbb{N}}, \tag{4.3}$$

since $u^j \geq 0$. Let u^j_m be given by (4.1).

Lemma 3. *For all $j \in I$, $\sup_{p \in A^{\mathbb{N}}} |u^j(p) - \frac{1}{m} u^j_m(p)| \leq \frac{1}{m}$.*

This lemma is identical with Lemma 2.3 of Secchi and Sudderth (2001), and is also easy to prove directly.

Let $\epsilon > 0$ be from the statement of Theorem 1, and choose m so that $\frac{1}{m} < \epsilon$ so that the functions $\frac{1}{m} u^j_m$ are, by the lemma, uniformly within distance ϵ of the u^j. So it will suffice for Theorem 1 to prove that there is a subgame perfect equilibrium for the game with payoffs the functions $\frac{1}{m} u^j_m, j \in I$. In the next section, it is shown that there does exist a subgame perfect equilibrium for the game with payoffs the $u^j_m, j \in I$, but this is equivalent.

5 Completion of the Proof of Theorem 1

Consider a game $\mathcal{G}(u_m, i)$ where the payoff functions $u_m = (u^j_m)_{j \in I}$ are as in (4.1). Let $\mathcal{C}$ be the collection of closed sets $\{C_{j,k} : j \in I, k = 1, \ldots, m\}$ satisfying (4.2) and (4.3). In this section, the notation $\mathcal{G}(\mathcal{C}, i)$ is used for the game $\mathcal{G}(u_m, i)$. As was mentionned at the end of the previous section, Theorem 1 will be established once $\mathcal{G}(u_m, i) = \mathcal{G}(\mathcal{C}, i)$ is seen to have a subgame perfect equilibrium.

Here is a simple fact about closed sets.

Lemma 4. *Let C be a closed subset of $A^{\mathbb{N}}$ and let $p = (a_1, a_2, \ldots) \in A^{\mathbb{N}}$ be a play such that, for all n, the section $C(a_1, \ldots, a_n)$ is not empty. Then $p \in C$.*

Proof. For each n, there is, by hypothesis, a play q_n such that the play $p_n = (a_1, \ldots, a_n)q_n \in C$. Now $p_n \to p$ as $n \to \infty$. Hence, $p \in C$. $\square$

For a finite history $h \in \mathcal{H}$, let $\mathcal{C}h$ be the collection of sets $\{C_{j,k}h\ ,:\ C_{j,k} \in \mathcal{C}\}$ where, for $C \subseteq A^{\mathbb{N}}$, $Ch = \{hp \in C\ :\ p \in A^{\mathbb{N}}\}$ is the h - section of C. The subgame $\mathcal{G}(u_m h, ih)$ of $\mathcal{G}(u_m, i)$ is, in the notation of this section, the same as $\mathcal{G}(\mathcal{C}h, ih)$. Theorem 1 will follow from the next lemma. Its proof uses ideas from the proof of Lemma 4.2 in Secchi and Sudderth (2001).

Lemma 5. *The game $\mathcal{G}(\mathcal{C}, i)$ has a subgame perfect equilibrium $\sigma = (\sigma^j)_{j \in I}$.*

Proof. The proof is by induction on the integer $\lambda(\mathcal{C})$ defined to be the number of sets $C_{j,k} \in \mathcal{C}$ that are nonempty, proper subsets of $A^{\mathbb{N}}$. Notice that $\lambda(\mathcal{C}h') \geq \lambda(\mathcal{C}h)$ whenever h and h' are histories such that h' is an initial segment of h. The reason is that a section of the empty set or the whole space is always equal to the set itself. Thus the number of proper, nonempty sets cannot increase as play proceeds along a history. Also, if a set C is empty or equal to the whole space $A^{\mathbb{N}}$, then the indicator function 1_C is a constant. Thus, if $\lambda(\mathcal{C}) = 0$, the game $\mathcal{G}(\mathcal{C}, i)$ has constant payoff functions and every profile is a subgame perfect equilibrium.

Assume now that $\lambda(\mathcal{C}) = \lambda_0 > 0$ and make the inductive assumption that the assertion holds for all games $\mathcal{G}(\mathcal{C}', i')$ where $\mathcal{C}'$ is another such collection of closed sets with $\lambda(\mathcal{C}') < \lambda_0$.

To define the profile σ, an action $\sigma^{i(h)}(h)$ must be assigned to every history $h \in \mathcal{H}$. Three cases will be considered.

Case 1. $\lambda(\mathcal{C}h) < \lambda_0$.

Let $h = (a_1, \ldots, a_n)$, and let l be the least positive integer in $\{1, \ldots, n\}$ such that, for $h' = (a_1, \ldots, a_l)$, $\lambda(\mathcal{C}h') < \lambda_0$. By the inductive hypothesis, there is a subgame perfect equilibrium $\sigma_{h'}$ for the game $\mathcal{G}(\mathcal{C}h', ih')$. Define

$$\sigma^{i(h)}(h) \;=\; \sigma_{h'}^{i(h)}(a_{l+1}, \ldots, a_n) \;=\; \sigma_{h'}^{(ih')(a_{l+1},\ldots,a_n)}(a_{l+1}, \ldots, a_n).$$

Indeed, the definition of $\sigma^{i(h'')}(h'')$ is made in the same way for every history h'' for which h' is an initial segment. To be more specific, if $h'' = (a_1, \ldots, a_l, b_1, \ldots, b_r)$, then h'' also satisfies Case 1 and we set

$$\sigma^{i(h'')}(h'') \;=\; \sigma_{h'}^{i(h'')}(b_1, \ldots, b_r).$$

This consistently defines $\sigma^{i(h)}(h)$ for all h satisfying Case 1 in such a way that the conditional profile $\sigma[h] = \sigma_{h'}[(a_{l+1}, \ldots, a_n)]$ is subgame perfect for

$\mathcal{G}(\mathcal{C}h, ih) = \mathcal{G}(\mathcal{C}h'(a_{l+1}, \ldots, a_n), ih'(a_{l+1}, \ldots, a_n)$. In particular, $\sigma[h]$ is an equilibrium for $\mathcal{G}(\mathcal{C}h, ih)$.

To specify the remaining two cases, let, for each $j \in I = \{1, \ldots, \nu\}$, the integer k_j be the largest k in $\{1, \ldots, m\}$ such that the set $C_{j,k}$ is not empty. Then a play p in C_{j,k_j} results in the largest possible reward to player j, namely

$$1_{C_{j,1}}(p) + \cdots + 1_{C_{j,k_j}}(p) \;=\; k_j.$$

Case 2. $\lambda(\mathcal{C}h) = \lambda_0$ and $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})h \neq \emptyset$.

In this case, there must be an action b_1 such that $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})hb_1 \neq \emptyset$. Set $\sigma^{i(h)}(h) = b_1$. Similarly $b_2, b_3, \ldots$ are defined so that, for all n, $\sigma^{i(h(b_1,\ldots,b_n))}(h(b_1, \ldots, b_n)) = b_{n+1}$ and $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})h(b_1, \ldots, b_n)$ is nonempty. Then, by the previous lemma, the conditional profile $\sigma[h]$ results in a play $p = (b_1, b_2, \ldots)$ such that $hp \in C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu}$. Thus every player receives his or her maximum possible payoff and $\sigma[h]$ is an equilibrium for the game $\mathcal{G}(\mathcal{C}h, ih)$.

Case 3. $\lambda(\mathcal{C}h) = \lambda_0$ and $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})h = \emptyset$.

Each of the sets $C_{j,k_j}h$, $j \in I$, must be nonempty because $\lambda(\mathcal{C}h)$ would be smaller than λ_0 otherwise.

The history h of this case cannot have an initial segment h' satisfying Case 1 because, as was already pointed out, $\lambda(\mathcal{C}h') \geq \lambda(\mathcal{C}h)$ whenever h' is an initial segment of h. However, h could have initial segments satisfying Case 2. In general, if $h = (a_1, \ldots, a_n)$ is not the empty history e, then there is an integer l, $1 \leq l \leq n$ such that the histories $h_r = (a_1, \ldots, a_r)$ satisfy Case 2 for $r = 1, \ldots, l-1$ and Case 3 for $r = l, \ldots, n$. Thus $h' = (a_1, \ldots, a_l)$ is the first initial segment of h satisfying Case 3. So $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})h' = \emptyset$.

Let $p = (b_1, b_2, \ldots)$ be an arbitrary play. Then it cannot be the case that, for all $j \in I$ and all nonnegative integers q, the section $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})h'(b_1, \ldots, b_q)$ is nonempty. For if this were the case, it would follow from the previous lemma that p would belong to $(C_{1,k_1} \cap \cdots \cap C_{\nu,k_\nu})h'$, a contradiction.

Thus there is, for every play $p = (b_1, b_2, \ldots)$, some $j \in I$ and some nonnegative integer q such that the section $C_{j,k_j}h'(b_1, \ldots, b_q)$ is empty. (When $q = 0$, the history $(b_1, \ldots, b_q) = e$.) Let $t(p)$ be the least integer q for which this occurs. Then t is a stop rule and, for each p, the collection $\mathcal{C}_p = \mathcal{C}h'h_t(p)$ contains at least one additional empty set. Hence $\lambda(\mathcal{C}_p) < \lambda_0$. By the inductive hypothesis, the game $\mathcal{G}(\mathcal{C}h'h_t(p), ih'h_t(p))$ has a subgame perfect equilibrium for every p. By Lemma 2, the game $\mathcal{G}(\mathcal{C}h', ih')$ also has a subgame perfect equilibrium $\sigma_{h'}$. As in Case 1, define

$$\sigma^{i(h)}(h) = \sigma_{h'}^{i(h)}(a_{l+1}, \ldots, a_n).$$

If the history h is the empty history e, then, for each p, define $t(p)$ to be the least q such that some section $C_{j,k_j}(b_1, \ldots, b_q)$ is empty. Again t is a stop rule and the argument proceeds as above.

The profile σ is now completely defined. It follows from the construction that the conditional profile $\sigma[h]$ is an equilibrium for the game $\mathcal{G}(\mathcal{C}h, ih)$ for all $h \in \mathcal{H}$. □

References

[1] Blackwell, D. 1981. Borel sets via games. *Ann. Probab. 9* 321-322.

[2] Dubins, L.E., Savage, L.J. 1976. *Inequalities for Stochastic Processes.* Dover, New York.

[3] Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Solan, E., Vrieze, K. 2010. Borel games with lower-semi-continuous payoffs. Preprint.

[4] Maitra, A.P., Sudderth, W.D. 1996. *Discrete Gambling and Stochastic Games*, Springer, New York.

[5] Maitra, A.P., Sudderth, W.D. 2007. Subgame-perfect equilibria for stochastic games. *Math. Oper. Res. 32* 711-722.

[6] Secchi, P., Sudderth, W. 2001. N-person stochastic games with upper semicontinuous payoffs. *Int. J. Game Theory 30* 491-502.

[7] Solan, E., Vieille, N. 2002. Correlated equilibrium in stochastic games. *Games Econom. Behav. 38* 362-399.