JOURNAL OF OPTIMIZATION THEORY AND APPLICATIONS: Vol. 98, No. 1, pp. 109–129, JULY 1998

Limit Consistent Solutions in Noncooperative Games[1]

A. Perea y Monsuwé[2] and H. Peters[3]

Communicated by G. Papavassilopoulos

Abstract. Strong and limit consistency in finite noncooperative games are studied. A solution is called strongly consistent if it is both consistent and conversely consistent (Ref. 1). We provide sufficient conditions on one-person behavior such that a strongly consistent solution is nonempty. We introduce limit consistency for normal form games and extensive form games. Roughly, this means that the solution can be approximated by strongly consistent solutions. We then show that the perfect and proper equilibrium correspondences in normal form games, as well as the weakly perfect and sequential equilibrium correspondences for extensive form games, are limit consistent.

Key Words. Noncooperative games, consistency, Nash equilibrium refinements, extensive form games.

1. Introduction

The concept of consistency has been widely applied in game theory. Globally speaking, a game-theoretic solution is called consistent if it is invariant in the reduced game for a subset of the players given that the other players are kept at the solution outcome in the original game. Here, an outcome may be a payoff vector (usually in a cooperative game) or a strategy vector (usually in a noncooperative game). Obviously, this global definition leaves much freedom for formalization, by varying the definition of the reduced game. For many different solutions, reduced games have been proposed with respect to which the solution is consistent.

[1]The authors are indebted to two anonymous referees whose suggestions improved the presentation of this paper.
[2]Postdoctoral Researcher, Department of Economics, Autònoma University, Barcelona, Spain.
[3]Professor, Department of Quantitative Economics, University of Maastricht, Maastricht, Netherlands.

The paper by Peleg and Tijs (Ref. 1) was the first to apply the consistency concept to noncooperative games. For a strategy profile in a normal form game and a coalition, a natural reduced game arises by fixing the players outside the coalition at their strategies as given in the profile. Obviously, a Nash equilibrium will induce a Nash equilibrium in the reduced game. Conversely, if in every reduced game the induced strategy profile is a Nash equilibrium, then the original profile must be a Nash equilibrium. We express this by saying that the Nash equilibrium correspondence is consistent as well as conversely consistent. For a survey on reduced noncooperative games, consistency, and converse consistency, see Peleg and Tijs (Ref. 1). In the present paper, the combination of consistency and converse consistency is termed strong consistency. Once a solution is strongly consistent, it is completely determined by its behavior on one-person games. Hence, the Nash equilibrium correspondence for normal form games is the unique strongly consistent solution assigning the set of expected payoff maximizing strategies to one-player games. This is the main result in Peleg and Tijs (Ref. 1). By varying the behavior of a solution on one-person games, other strongly consistent solutions are obtained. Patrone et al. (Ref. 2) call these solutions personalized Nash equilibria. Basically, by defining one-person behavior and imposing strong consistency, a best response correspondence is defined and extended to an equilibrium concept.

In this paper, we aim to contribute to the study of consistency in finite noncooperative games in the following ways. First, we provide sufficient conditions on one-person behavior such that a corresponding strongly consistent solution is nonempty. Observe that the pure Nash equilibrium correspondence cannot satisfy these conditions, because it is strongly consistent but may be empty. In fact, Norde et al. (Ref. 3) have proved that there exists no proper refinement of the Nash equilibrium correspondence that is consistent and nonempty—let alone strongly consistent. This result applies in particular to familiar refinements like perfect and proper Nash equilibrium, which are both nonempty and therefore cannot be consistent. Second, as an approach to this problem, we introduce the concept of limit consistency. A solution is limit consistent, roughly, if it can be approximated by strongly consistent solutions. A limit consistent solution is generated by a class of one-person solutions rather than by a unique one-person solution. It will be shown that both the perfect and the proper equilibrium correspondences are limit consistent. Third, we extend the concept of limit consistency to solutions for games in extensive form. We show that the correspondences of weakly perfect equilibrium and sequential equilibrium (Kreps and Wilson, Ref. 4) are limit consistent.

The organization of the paper is as follows. In Section 2, we introduce strong consistency for solutions in normal form games; in Section 3, among other things, we provide conditions for nonemptiness of strongly consistent solutions. In Section 4, we introduce limit consistency and provide a necessary and sufficient condition for nonemptiness of limit consistent solutions. The perfect and proper equilibrium correspondences are proved to be limit consistent in Section 5. In Section 6, extensive form games and the corresponding extensions of strong and limit consistency are defined. Sections 7 and 8 give the mentioned results on weakly perfect and sequential equilibria, respectively.

2. Strongly Consistent Solutions in Normal Form Games

2.1. Normal Form Games and Reduced Games. Let $\Gamma = \langle N, M, v\rangle$ be a normal form game with player set N, where $M = \times_{i\in N} M_i$ denotes the space of pure strategy profiles and $v = (v_i)_{i\in N}$ is a collection of payoff functions $v_i : M \to \mathbb{R}$. We assume that the pure strategy spaces M_i are all finite. In the remainder of this article, we write $\langle M, v\rangle$ instead of $\langle N, M, v\rangle$ if there can be no misunderstanding about the player set. Moreover, we write "game" instead of "normal form game," if it is clear that we are talking about normal form games.

A mixed strategy for player i is a probability distribution p_i on the set M_i of pure player i strategies. A combination $p = (p_i)_{i\in N}$ of mixed strategies p_i is called a mixed strategy profile (MSP). If the MSP p is played, the expected payoff for player i, is denoted by $v_i(p)$.

An MSP p is called completely mixed if

$$p_i \in \Delta^0(M_i), \qquad \text{for every } i.$$

Here, $\Delta^0(M_i)$ denotes the set of probability distributions on M_i which put strictly positive weight on every $m_i \in M_i$.

By $B_i(p, \Gamma)$, we denote the set of player i's pure best responses against the MSP p, i.e.,

$$B_i(p, \Gamma) := \{m_i \mid v_i(p\backslash m_i) \geq v_i(p\backslash l_i), \text{ for every } l_i \in M_i\}.$$

Here, $p\backslash m_i$ is the MSP which we obtain if the mixed strategy p_i is replaced by the pure strategy m_i in p.

Definition 2.1. Let $\Gamma = \langle M, v\rangle$ be a game, p an MSP, and T a coalition of players. By $\Gamma^{T,p}$, we denote the reduced game $\langle T, M', v'\rangle$, where $M' = \times_{i\in T} M_i$ and $v'_i(p') = v_i(p', p_{N\backslash T})$, for all $i \in T$ and all MSP $p' = (p'_i)_{i\in T}$.

Here, $p_{N\backslash T}$ denotes the restriction of the MSP p on the players in $N\backslash T$. Hence, $\Gamma^{T,p}$ is the game with player set T which we obtain if we assume that players outside T play according to p.

A family $\mathscr{G}$ of normal form games is called closed if, for every Γ in $\mathscr{G}$, every nonempty coalition $T \subset N$, and every MSP p, the reduced game $\Gamma^{T,p}$ is also in $\mathscr{G}$.

2.2. Strongly Consistent Solutions. A solution on a family of normal form games is a correspondence which assigns to every game in this family a set of MSPs in this game. For a game Γ in this family, we call $\varphi(\Gamma)$ the set of solution profiles in this game.

Let $\mathscr{G}$ be a closed family of normal form games, and let φ be a solution on $\mathscr{G}$. For a game Γ in $\mathscr{G}$ containing at least two players, we define

$$\tilde{\varphi}(\Gamma) := \{p \mid p_T \in \varphi(\Gamma^{T,p}), \text{ for every } \varnothing \subsetneq T \subsetneq N\}.$$

In other words, $\tilde{\varphi}(\Gamma)$ is the set of MPSs such that, for every nonempty coalition, the reduced profile is a solution profile of the reduced game.

For one-player games Γ, we define $\tilde{\varphi}(\Gamma) := \varphi(\Gamma)$. The reason for this definition is the fact that there are no coalitions $\varnothing \subsetneq T \subsetneq N$ in one-player games.

Definition 2.2. A solution φ is called consistent (CONS) if $\varphi(\Gamma) \subset \tilde{\varphi}(\Gamma)$ for every Γ in $\mathscr{G}$.

This means that, for every solution profile in Γ and every nontrivial subset of players, the reduced solution profile is a solution of the reduced game.

Definition 2.3. A solution φ is called conversely consistent (COCONS) if $\tilde{\varphi}(\Gamma) \subset \varphi(\Gamma)$ for every Γ in $\mathscr{G}$ with at least two players.

This means that an MSP which has the property that, for every nontrivial subset of players, the reduced profile is a solution profile of the reduced game should be a solution profile itself.

Note that properties CONS and COCONS together are equivalent to the property

$$\varphi(\Gamma) = \tilde{\varphi}(\Gamma), \qquad \text{for every } \Gamma \text{ in } \mathscr{G}.$$

Definition 2.4. We call a solution strongly consistent if it satisfies CONS and COCONS.

3. Properties of Strongly Consistent Solutions

In Patrone et al. (Ref. 2), it is shown that every strongly consistent solution is completely determined by its behavior on one-person games. In fact, the behavior of a strongly consistent solution on the class of one-person games induces an optimal response correspondence, which assigns to every MSP p a set of MSPs which are optimal responses against p. Here, we use the expression "optimal response" instead of "best response", since the latter term is reserved to denote payoff-maximizing responses. Hence, best responses are in our notation a special case of optimal responses. In the following lemma, it is shown that the class of strongly consistent solutions consists exactly of those solutions which assign to a game the set of MSPs which are optimal responses against themselves.

Lemma 3.1. A solution φ is strongly consistent if and only if, for every game Γ,

$$\varphi(\Gamma) = \{p \mid p_i \in \varphi(\Gamma^{\{i\},p}), \text{ for every } i\}.$$

There is a close relationship between this lemma and a result in Patrone et al. (Ref. 2), which states that a solution is a personalized Nash equilibrium correspondence if and only if it satisfies personalized one-person rationality, CONS, and COCONS. This relationship stems from the fact that the class of personalized Nash equilibrium correspondences contains exactly those solutions assigning to every game Γ the set of MSPs

$$\{p \mid p_i \in \varphi_i^1(\Gamma^{\{i\},p}), \text{ for every } i\},$$

where φ_i^1 is an arbitrary but fixed solution on one-person games for every i.

Proof of Lemma 3.1. (a) First, we show the implication from left to right. Let φ be a strongly consistent solution. We define the solution ψ by

$$\psi(\Gamma) = \{p \mid p_i \in \varphi(\Gamma^{\{i\},p}), \text{ for every } i\},$$

for every game Γ and show that $\varphi = \psi$. Because φ is consistent, it follows that

$$\varphi(\Gamma) \subset \psi(\Gamma), \qquad \text{for every game } \Gamma.$$

The other inclusion will be shown by induction on the number of players in Γ. If Γ is a one-player game, the inclusion

$$\psi(\Gamma) \subset \varphi(\Gamma)$$

holds by definition. Now, assume that $\psi(\Gamma') \subset \varphi(\Gamma')$ for every game Γ with strictly less than n players and consider a game Γ with n players. Choose an arbitrary profile $p \in \psi(\Gamma)$ and let

$$\varnothing \neq T \subsetneq N.$$

By definition, we have $p_T \in \psi(\Gamma^{T,p})$. Using the induction assumption, it follows that $p_T \in \varphi(\Gamma^{T,p})$. Since this holds for every such T, the converse consistency of φ implies that $p \in \varphi(\Gamma)$.

(b) The implication from right to left is obvious. □

Intuitively, the lemma says that a solution is strongly consistent if and only if it assigns to every game the set of strategy profiles in which every strategy is an optimal response. Therefore, the behavior of a strongly consistent solution is similar to that of the Nash equilibrium correspondence. This is a reason why strongly consistent solutions are called personalized Nash equilibria in Patrone et al. (Ref. 2).

The observations above enable us to construct a mechanism which produces strongly consistent solutions. The mechanism works in the following way.

First, choose a one-person solution φ^1 which assigns to every one-person game a set of mixed strategies.

Definition 3.1. For every game Γ, the optimal response correspondence induced by the one-person solution φ^1 is the function f assigning to every player i and every MSP p the set $\varphi^1(\Gamma^{\{i\},p})$.

Finally, define the solution φ to be the solution which assigns to every normal form game the set of MSPs in which every mixed strategy is an optimal response (w.r.t. φ^1) against the strategies of the other players. Formally, we obtain the following definition.

Definition 3.2. The solution φ given by

$$\varphi(\Gamma) := \{p \mid p_i \in \varphi^1(\Gamma^{\{i\},p}), \text{ for every } i\}$$

is called the strongly consistent solution generated by the one-person solution φ^1.

The Nash equilibrium correspondence, for example, is a strongly consistent solution generated by the one-person solution which assigns to every one-person game the set of payoff maximizers. A proof for this result can be found in Peleg and Tijs (Ref. 1).

In general, the solution φ generated by the mechanism described above may be empty in some games, even when the one-person solution φ^1 is always nonempty. To illustrate this fact, consider the one-person solution φ^1 which assigns to every one-person game the set of pure strategies with maximal payoff. Obviously, the solution generated by φ^1 assigns to every game the set of pure Nash equilibria. However, it is well known that pure Nash equilibria do not always exist.

In the following theorem, we describe a sufficient condition to generate a nonempty, strongly consistent solution. By nonempty, we mean that the solution assigns to every game a nonempty set of MSPs. In order to state the theorem, we need the definition of upper hemicontinuity of a one-person solution.

Definition 3.3. A one-person solution φ^1 is called upper hemicontinuous if, for every one-person game $\Gamma = \langle M_i, v_i \rangle$, every sequence $\Gamma^k = \langle M_i, v_i^k \rangle$ of games converging to Γ, and every sequence p^k of MSPs converging to p with $p^k \in \varphi^1(\Gamma^k)$, we have $p \in \varphi^1(\Gamma)$.

In the proof of the theorem, we also need another kind of upper hemicontinuity, namely the upper hemicontinuity of optimal response correspondences.

Definition 3.4. The optimal response correspondence induced by φ^1 is upper hemicontinuous if, for every sequence p^k of MSPs converging to p and every sequence $q_i^k \in \varphi^1(\Gamma^{\{i\}, p^k})$ converging to q_i for all i, it holds that $q_i \in \varphi^1(\Gamma^{\{i\}, p})$ for all i.

In words, if a sequence of strategies consists of optimal responses against a sequence of MSPs, then the limit strategy is an optimal response against the limit MSP.

Observe that both types of upper hemicontinuity are based on the usual closed-graph definition of upper hemicontinuity of a correspondence. A one-person solution, however, is regarded as a correspondence on a class of games, whereas the optimal response correspondence varies with the strategy profiles within the same game.

Theorem 3.1. Let φ be a strongly consistent solution generated by the one-person solution φ^1. If φ^1 is upper hemicontinuous and assigns to every one-person game a nonempty, convex, and compact set of strategies, then φ is nonempty.

Proof. Let $\Gamma = \langle M, v \rangle$ be a normal form game. First, we show that the optimal response correspondence in Γ induced by φ^1 is upper-hemicontinuous. Let p^k be a sequence of MSPs converging to p and $q_i^k \in \varphi^1(\Gamma^{\{i\}, p^k})$ a sequence of optimal responses converging to q_i. For every k, define the game $\Gamma^{i,k}$ by $\Gamma^{i,k} := \Gamma^{\{i\}, p^k}$. Obviously, the sequence $\Gamma^{i,k}$ converges to $\Gamma^{\{i\}, p}$.

By the upper hemicontinuity of φ^1, we know that $q_i \in \varphi^1(\Gamma^{\{i\},p})$, which implies that the optimal response correspondence is upper hemicontinuous.

Since, by assumption, the optimal response correspondence assigns to every MSP p a nonempty, convex, and compact set of MSPs, we can apply Kakutani's theorem to ensure the existence of a fixed point of the optimal response correspondence in the game Γ. Since the solutions in $\varphi(\Gamma)$ are exactly those fixed points, the nonemptiness of φ is established. □

Example 3.1. Let φ^1 be the one-person solution assigning to every one-person game $\Gamma = \langle M_i, v_i \rangle$ the set of MSPs p with $v_i(p)$ minimal. Since it can be checked easily that φ^1 satisfies all the conditions in the theorem above, it follows that the strongly consistent solution generated by φ^1 is nonempty.

A solution which assigns to every game a subset of the set of Nash equilibria is called a Nash equilibrium refinement. The following result, which has been proved in Norde et al. (Ref. 3), shows that it is impossible to find a nonempty, consistent Nash equilibrium refinement other than the Nash equilibrium correspondence itself.

Theorem 3.2. Let φ be a nonempty Nash equilibrium refinement satisfying CONS. Then, φ is equal to the Nash equilibrium correspondence.

This theorem implies, in particular, that there does not exist a nonempty, strongly consistent solution which is a strict refinement of the Nash equilibrium correspondence. This is an important reason to introduce a new mechanism in the next section, which enables us to generate so-called limit consistent solutions. In particular, Nash equilibrium refinements such as perfect equilibria, proper equilibria, weakly perfect and sequential equilibria can be generated by this mechanism.

4. Limit Consistent Solutions

In this section, we concentrate on solutions which need not be strongly consistent, but have the property that every solution point can be approximated by a sequence of solution points of strongly consistent solutions. We call such solutions limit consistent.

For a formal definition of limit consistent solutions, we need the notion of a so-called one-person solution function.

Definition 4.1. A one-person solution function is a function Φ^1 which assigns to every $\epsilon > 0$ a one-person solution $\varphi^1 = \Phi^1(\epsilon)$.

Intuitively, the number ϵ can be interpreted as a perturbation factor, and the one-person solution $\Phi^1(\epsilon)$ can be seen as a perturbed one-person solution.

For every $\epsilon > 0$, the unique strongly consistent solution generated by the one-person solution $\Phi^1(\epsilon)$ is denoted by $\Phi(\epsilon)$. The function Φ is called a solution function.

Definition 4.2. A solution φ is called limit consistent if there is a one-person solution function Φ^1 such that

$$\varphi(\Gamma) = \{p \mid \text{there is a sequence } \epsilon^k > 0 \text{ converging to 0 and a sequence } p^k \in \Phi(\epsilon^k)(\Gamma) \text{ converging to } p\}.$$

In this case, we say that the limit consistent solution is generated by the one-person solution function Φ^1.

Every one-person solution function Φ^1 induces a solution function $\bar{\Phi}$ given by

$$\bar{\Phi}(\epsilon)(\Gamma) := \bigcup_{\delta \in (0, \epsilon]} \Phi(\delta)(\Gamma),$$

for every game Γ and every $\epsilon > 0$.

Theorem 4.1. The following two statements are equivalent:

(i) φ is a limit consistent solution generated by the one-person solution function Φ^1;
(ii) $\varphi(\Gamma) = \bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma))$ for every game Γ.

Here, $\mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma))$ denotes the closure of the set $\bar{\Phi}(\epsilon)(\Gamma)$.

Proof. We only show the implication from (i) to (ii); the other implication can be shown in a similar way.

Let φ be a limit consistent solution generated by Φ^1, and let Γ be a game. We prove the coincidence of both sets in (ii) by the double inclusion argument.

(a) We show that $\varphi(\Gamma) \subset \bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma))$. Let $p \in \varphi(\Gamma)$, and let $\epsilon > 0$ be given. By definition of limit consistency, there is a sequence $\epsilon^k > 0$ converging to 0 and a sequence $p^k \in \Phi(\epsilon^k)(\Gamma)$ converging to p. Obviously,

$$p^k \in \bar{\Phi}(\epsilon)(\Gamma), \qquad \text{for } k \text{ large enough,}$$

which implies that $p \in \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma))$. Since this holds for every $\epsilon > 0$, it follows that

$$p \in \bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)).$$

(b) We show that $\bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)) \subset \varphi(\Gamma)$. Let

$$p \in \bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)).$$

Then,

$$p \in \mathrm{cl}(\bar{\Phi}(1/k)(\Gamma)), \qquad \text{for every } k,$$

so for every k there is a sequence $(p^{k,l})_{l \in \mathbb{N}} \in \bar{\Phi}(1/k)(\Gamma)$, with $p = \lim_{l \to \infty} p^{k,l}$. By definition, there is a number $\epsilon(k, l) \in (0, 1/k]$ for every k, l with $p^{k,l} \in \Phi(\epsilon(k, l))(\Gamma)$. For every k, choose an integer $l(k)$ with $\|p - p^{k,l(k)}\| \leq 1/k$. By construction, we obtain

$$p^{k,l(k)} \in \Phi(\epsilon(k, l(k)))(\Gamma) \quad \text{and} \quad p = \lim_{k \to \infty} p^{k,l(k)}.$$

Because $p^{k,l(k)}$ converges to p and $\epsilon^{k,l(k)}$ converges to 0 as $k \to \infty$, we conclude that $p \in \varphi(\Gamma)$. □

We say that a solution is closed if it assigns to every game a closed set of MSPs.

From Theorem 4.1, it follows that a limit consistent solution assigns to every game Γ the set $\bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma))$, which is obviously closed. Hence, we obtain the following corollary.

Corollary 4.1. Every limit consistent solution is closed.

It turns out that nonemptiness of limit consistent solutions can be characterized by nonemptiness of the corresponding solution functions $\bar{\Phi}$.

Theorem 4.2. Let φ be a limit consistent solution generated by the one-person solution function Φ^1. Then, φ is nonempty if and only if $\bar{\Phi}(\epsilon)$ is nonempty for every $\epsilon > 0$.

Proof. First, we prove the implication from left to right. Let φ be a nonempty limit consistent solution generated by Φ^1 and Γ an arbitrary game. By Theorem 4.1, we know that

$$\varphi(\Gamma) = \bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)),$$

which implies that

$$\mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)) \neq \varnothing, \qquad \text{for every } \epsilon > 0.$$

This is only possible when

$$\bar{\Phi}(\epsilon)(\Gamma) \neq \varnothing, \qquad \text{for every } \epsilon > 0.$$

Now, consider the direction from right to left. Let the solution $\bar{\Phi}(\epsilon)$ be nonempty for every $\epsilon > 0$, and let Γ be a game. Then, the set $\mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma))$ is compact for every $\epsilon > 0$. Since, by construction,

$$\mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)) \subset \mathrm{cl}(\bar{\Phi}(\epsilon')(\Gamma)), \qquad \text{whenever } \epsilon \leq \epsilon',$$

it follows by Theorem 4.1 that

$$\varphi(\Gamma) = \bigcap_{\epsilon > 0} \mathrm{cl}(\bar{\Phi}(\epsilon)(\Gamma)) \neq \varnothing. \qquad \square$$

This theorem and the definition of $\bar{\Phi}(\epsilon)(\Gamma)$ imply a characterization of the class of nonempty limit consistent solutions in terms of the solutions $\Phi(\epsilon)$.

Corollary 4.2. Let φ be a limit consistent solution generated by the one-person solution function Φ^1. Then, φ is nonempty if and only if, for every $\epsilon > 0$, there is a $\delta \in (0, \epsilon]$ with $\Phi(\delta)$ nonempty.

Example 4.1. Consider the one-person solution function Φ^1 which assigns to every $\epsilon > 0$ and every one-person game $\Gamma = \langle M_i, v_i \rangle$ the set

$$\Phi^1(\epsilon)(\Gamma) := \begin{cases} \{p \mid v_i(p) \text{ minimal}\}, & \text{if } \epsilon \in \mathbb{Q}, \\ \varnothing, & \text{otherwise.} \end{cases}$$

The strongly consistent solution $\Phi(\epsilon)$ generated by $\Phi^1(\epsilon)$ is nonempty if $\epsilon \in \mathbb{Q}$. Hence, in view of the corollary above, the limit consistent solution φ generated by Φ^1 is nonempty despite the fact that $\Phi(e)$ is empty for all $\epsilon \notin \mathbb{R}$.

5. Perfect and Proper Equilibria

One of the possible ways to define perfect equilibria is by saying that a perfect equilibrium can be approximated by a sequence of so called ϵ-perfect strategy profiles.

Definition 5.1. (a) Let $\epsilon > 0$ and Γ a normal form game. An MSP p is called ϵ-perfect if p is completely mixed and $v_i(p \backslash m_i) < v_i(p \backslash m_i')$ implies $p_i(m_i) \leq \epsilon$.

(b) An MSP p is a perfect equilibrium if there is a sequence ϵ^k of positive numbers converging to 0 and a sequence p^k of ϵ^k-perfect profiles converging to p.

Theorem 5.1. The perfect equilibrium correspondence is a limit consistent solution generated by the one-person solution function Φ^1 given by

$$\Phi^1(\epsilon)(\Gamma) = \{p_i \in \Delta^0(M_i) \mid v_i(m_i) < v_i(m_i') \text{ implies } p_i(m_i) \leq \epsilon\},$$

for every $\epsilon > 0$ and every one-person game $\Gamma = \langle M_i, v_i \rangle$.

The proof of this result is straightforward.

Definition 5.2. (a) An MSP p is called an ϵ-proper profile if p is completely mixed and $v_i(p \backslash m_i) < v_i(p \backslash m_i')$ implies $p_i(m_i) \leq \epsilon p_i(m_i')$.

(b) An MSP p is called a proper equilibrium if there is a sequence ϵ^k converging to 0 and a sequence p^k of ϵ^k-proper profiles converging to p.

Similarly to perfect equilibria, we can show that the proper equilibrium correspondence is limit consistent.

Theorem 5.2. The proper equilibrium correspondence is a limit consistent solution generated by the one-person solution function Φ^1 given by

$$\Phi^1(\epsilon)(\Gamma) = \{p_i \in \Delta^0(M_i) \mid v_i(m_i) < v_i(m_i') \text{ implies } p_i(m_i) \leq \epsilon p_i(m_i')\},$$

for every $\epsilon > 0$ and every one-person game $\Gamma = \langle M_i, v_i \rangle$.

6. Strongly and Limit Consistent Solutions in Extensive Form Games

6.1. Extensive Form Games. An extensive form structure is a quintet $\mathscr{S} = \langle K, P, H, A, \tau \rangle$ characterized by (i) to (v) below.

(i) K is a rooted tree with root x_0. The set of terminal nodes is denoted by Z, and the set of nonterminal nodes except x_0 is denoted by X. The unique sequence of nodes and edges connecting the root x_0 and a terminal node z is called the path from x_0 to z. For an $x \in X \cup \{x_0\}$, $E(x)$ is the set of edges leaving the node x.

(ii) $P: X \to \{0, 1, 2, \ldots, n\}$, and $P(x)$ is the player that controls the node x. If $P(x) = 0$, then the node x represents a situation where a chance move occurs. Such nodes will be called chance nodes.

(ii) H is an n-tuple $(H_1, H_2, \ldots, H_n)$, where H_i is a partition of the set $P^{-1}(i)$ into information sets (or agents) such that, for every information set h, the following conditions are satisfied:

(a) every path from x_0 to (an element of) Z intersects h at most once;
(b) $|E(x)| = |E(y)|$, for all $x, y \in h$.

We also use the letter H for the set $\bigcup_i H_i$ of all information sets.

(iv) For every information set h, $A(h)$ is a partition of the set of edges $\bigcup_{x \in h} E(x)$ leaving the information set h into actions such that

$$|E(x) \cap a| = 1,$$

for any action $a \in A(h)$ and $x \in h$. We will assume that

$$|A(h)| \geq 2, \qquad \text{for all } h \in H.$$

(v) τ is a function which defines at every chance node a strictly positive probability distribution on the set of edges which leave the chance node. These edges represent the different chance moves.

An extensive form game is a pair $\langle \mathcal{S}, u \rangle$, where $\mathcal{S}$ is an extensive form structure and u is an n-tuple $(u_1, u_2, \ldots, u_n)$, where $u_i : Z \to \mathbb{R}$. The function u_i is the payoff function for player i.

We assume that the extensive form games considered have perfect recall, which means that two paths leading to the same player i information set contain the same player i actions.

A behavior strategy for player i is a function σ_i which assigns to every player i information set h a mixed strategy for agent h. This means that

$$\sigma_i = (\sigma_h)_{h \in H_i}, \qquad \text{with } \sigma_h \in \Delta(A(h)), \text{ for every } h \in H_i.$$

A vector $\sigma = (\sigma_i)_{i \in N}$ consisting of behavior strategies is called a behavior strategy profile (BSP). Note that a BSP assigns to every information set h a mixed strategy σ_h for agent h. Therefore, we can write a BSP σ also in the form $\sigma = (\sigma_h)_{h \in H}$.

A belief system is a function β which assigns to every information set h a probability distribution on the nodes in this information set. Formally,

$$\beta = (\beta_h)_{h \in H}, \qquad \text{with } \beta_h \in \Delta(h), \text{ for every } h \in H.$$

Intuitively, a belief system β reflects at every information set the subjective probabilities assigned by the corresponding player to the nodes in this information set.

A combination (σ, β) of a BSP and a belief system is called an assessment.

6.2. Reduced Extensive Form Games, Strongly and Limit Consistent Solutions. Let Γ be an extensive form game, $T \subset N$ a subset of players, and

σ a BSP. Then, the reduced game $\Gamma^{T,\sigma}$ is the extensive form game which we obtain by splitting every information set h controlled by a player outside T into a collection of chance nodes at which the moves of nature are given by the probability distribution σ_h. If an action is played with probability zero in σ_h, then the corresponding edges are deleted in the game tree. In this way, we assure that the probability distributions of the moves of nature are strictly positive. The remaining players in the reduced game are the players in T.

A solution for extensive form games is a correspondence which assigns to every extensive form game a set of BSPs.

An assessment solution for extensive form games is a correspondence which assigns to every extensive form game a set of assessments.

Let φ be a solution. For a game Γ in $\mathscr{G}$, we define

$$\tilde{\varphi}(\Gamma) := \{\sigma \mid \sigma_T \in \varphi(\Gamma^{T,\sigma}), \text{ for every } \varnothing \subsetneq T \subsetneq N\},$$

where σ_T is the restriction of σ on the information sets controlled by players in T.

If φ is an assessment solution, we define

$$\tilde{\varphi}(\Gamma) := \{(\sigma, \beta) \mid (\sigma_T, \beta_T) \in \varphi(\Gamma^{T,\sigma}), \text{ for every } \varnothing \subsetneq T \subsetneq N\},$$

where (σ_T, β_T) is the restriction of (σ, β) on the information sets controlled by players in T.

In the obvious way, we define consistency (CONS) and converse consistency (COCONS) for solutions and assessment solutions. Again, we call a solution (or assessment solution) strongly consistent, if it is both consistent and conversely consistent.

The definition of limit consistent solutions (assessment solutions) for extensive form games is given in the obvious way.

7. Weakly Perfect Equilibria

Definition 7.1. A BSP σ is called a weakly perfect equilibrium in an extensive form game $\Gamma = \langle \mathscr{S}, u \rangle$ if there is a sequence of games $\Gamma^k = \langle \mathscr{S}, u^k \rangle$ converging to Γ and a sequence of completely mixed BSPs σ^k converging to σ such that $a \notin B_h(\sigma^k, \Gamma^k)$, for some k, implies $\sigma_h(a) = 0$, for all $a \in A(h)$ and $h \in H_i$. Here, $\beta_h(\sigma^k, \Gamma^k)$ denotes the set of pure best responses (actions) of agent h against σ^k in Γ^k.

This definition of weakly perfect equilibria can be found in Kreps and Wilson (Ref. 4). It can be seen easily that every perfect equilibrium of the agent normal form of Γ is weakly perfect by choosing Γ^k equal to Γ.

In the following theorem, we show that the weakly perfect equilibrium concept is limit consistent. In order to prove this result, we need two more definitions.

For two extensive form games $\Gamma = \langle \mathscr{S}, u \rangle$ and $\Gamma' = \langle \mathscr{S}, u' \rangle$ with the same extensive form structure, we define the distance $d(\Gamma, \Gamma')$ by

$$d(\Gamma, \Gamma') := \max_{i,z} |u_i(z) - u'_i(z)|.$$

Definition 7.2. In an extensive form game Γ, the BSP σ is called an ϵ-perfect equilibrium if σ is completely mixed and $U_i(\sigma \backslash a) < U_i(\sigma \backslash b)$ implies $\sigma_h(a) \leq \epsilon$, for every $h \in H_i$ and all a, $b \in A(h)$.

Here, U_i denotes the expected payoff.

Theorem 7.1. The weakly perfect equilibrium correspondence is a limit consistent solution generated by the one-person solution function Φ^1 given by

$$\Phi^1(\epsilon)(\Gamma) := \{\sigma \mid \exists \Gamma' \text{ with } d(\Gamma, \Gamma') \leq \epsilon \text{ such that } \sigma \text{ is } \epsilon\text{-perfect in } \Gamma'\},$$

for every $\epsilon > 0$ and every one-person game Γ.

Proof. Let φ be the weakly perfect equilibrium correspondence, and let $\Phi(\epsilon)$ be the strongly consistent solution generated by $\Phi^1(\epsilon)$. First, we show that $\Phi(e)$ is equal to the solution $\Psi(e)$ given by

$$\Psi(\epsilon)(\Gamma) := \{\sigma \mid \exists \Gamma' \text{ with } d(\Gamma, \Gamma') \leq \epsilon \text{ such that } \sigma \text{ is } \epsilon\text{-perfect in } \Gamma'\},$$

for every game Γ. To this purpose, we first prove that

$$\Psi(\epsilon)(\Gamma) \subset \Phi(\epsilon)(\Gamma), \qquad \text{for each game } \Gamma.$$

Let $\sigma \in \Psi(\epsilon)(\Gamma)$, which means that there is a game Γ' with $d(\Gamma, \Gamma') \leq \epsilon$ such that σ is ϵ-perfect in Γ'. For every i, we have that

$$d(\Gamma^{\{i\},\sigma}, \Gamma'^{\{i\},\sigma}) \leq \epsilon$$

and σ_i is ϵ-perfect in $\Gamma'^{\{i\},\sigma}$. So, by definition,

$$\sigma_i \in \Phi(\epsilon)(\Gamma^{\{i\},\sigma}), \qquad \text{for every } i,$$

which implies that $\sigma \in \Phi(\epsilon)(\Gamma)$.

Next, we show that $\Phi(\epsilon)(\Gamma) \subset \Psi(\epsilon)(\Gamma)$. Let $\sigma \in \Phi(\epsilon)(\Gamma)$, so

$$\sigma_i \in \Phi(\epsilon)(\Gamma^{\{i\},\sigma}), \qquad \text{for every } i.$$

Hence, for every i, we can find a game Γ_i^* with $d(\Gamma^{\{i\},\sigma}, \Gamma_i^*) \leq \epsilon$ such that σ_i is ϵ-perfect in Γ_i^*. Let $u_i^*(z)$ be the payoffs in Γ_i^* at the terminal nodes

z. We define the game $\Gamma' = \langle \mathscr{S}, u' \rangle$ (with the same extensive form structure as Γ) by

$$u'_i(z) := u^*_i(z), \qquad \text{for every } i \text{ and every terminal node } z.$$

Then, by construction, $d(\Gamma, \Gamma') \leq \epsilon$ and σ is ϵ-perfect in Γ' which implies that $\sigma \in \Psi(\epsilon)(\Gamma)$. This leads to the conclusion that

$$\Phi(\epsilon) = \Psi(\epsilon), \qquad \text{for every } \epsilon.$$

Let ψ be the limit consistent solution generated by Φ^1. We show that $\varphi = \psi$.

First, we prove $\varphi(\Gamma) \subset \psi(\Gamma)$ for every game Γ. Let $\Gamma = \langle \mathscr{S}, u \rangle$ be an extensive form game, and let $\sigma \in \varphi(\Gamma)$. So, σ is a weakly perfect equilibrium in Γ supported by some sequences $\Gamma^k = \langle \mathscr{S}, u^k \rangle$ and σ^k. For every k, we define the number ϵ^k by

$$\epsilon^k := \max\{\sigma^k_h(a) \mid a \notin B_h(\sigma^k, \Gamma^k)\}, \qquad \text{if } B_h(\sigma^k, \Gamma^k) \neq A(h).$$

If $B_h(\sigma^k, \Gamma^k) = A(h)$, we define $\epsilon^k := 1/k$.

By definition of ϵ^k, σ^k is ϵ^k-perfect in Γ^k for every k and σ^k converges to σ. It remains to show that ϵ^k converges to 0. Let a be an action such that $a \notin B_h(\sigma^l, \Gamma^l)$ for some l. Since σ^k and Γ^k are supporting sequences for the weakly perfect equilibrium σ, we have $\sigma_h(a) = 0$. So,

$$\lim_{k \to \infty} \sigma^k(a) = 0, \qquad \text{for all } a \text{ with } a \notin B_h(\sigma^l, \Gamma^k), \text{ for some } l,$$

which implies that ϵ^k converges to 0.

For every k, we choose a number $l(k)$ such that $\epsilon^{l(k)} \leq 1/k$ and $d(\Gamma, \Gamma^{l(k)}) \leq 1/k$. Then, by construction, $\sigma^{l(k)}$ is $(1/k)$-perfect in $\Gamma^{l(k)}$ and $d(\Gamma, \Gamma^{l(k)}) \leq 1/k$, which implies that

$$\sigma^{l(k)} \in \Psi(1/k)(\Gamma), \qquad \text{for every } k.$$

Since $\Psi(1/k) = \Phi(1/k)$ and $\sigma^{l(k)}$ converges to σ, it follows that $\sigma \in \psi(\Gamma)$.

Finally, we show that $\psi(\Gamma) \subset \varphi(\Gamma)$. Let $\sigma \in \psi(\Gamma)$, which means that there is a sequence ϵ^k of positive numbers converging to 0 and a sequence $\sigma^k \in \Phi(\epsilon^k)(\Gamma)$ converging to σ. Since $\Phi(\epsilon^k) = \Psi(\epsilon^k)$, it follows that, for every k, there is a game Γ^k with $d(\Gamma, \Gamma^k) \leq \epsilon^k$ and σ^k is ϵ^k-perfect in Γ^k. Since the collection of information sets and the set of actions at each information set is finite, we may assume w.l.o.g. that $B_h(\sigma^k, \Gamma^k)$ is constant over k for each information set h. Now, let $a \notin B_h(\sigma^k, \Gamma^k)$ for some k. Then, by assumption, $a \notin B_h(\sigma^k, \Gamma^k)$ for all k. Since σ^k is ϵ^k-perfect in Γ^k,

$$\sigma^k_h(a) \leq \epsilon^k, \qquad \text{for every } k,$$

which implies that $\sigma_h(a) = 0$. Using the fact that Γ^k converges to Γ, we may conclude that σ is a weakly perfect equilibrium in Γ, so $\sigma \in \varphi(\Gamma)$.

Hence, the solutions φ and ψ are equal, implying that φ is the limit consistent solution generated by Φ^1.

8. Sequential Equilibria

An assessment (σ, β) is called sequentially rational if, at every information set, the corresponding player maximizes his expected payoff given his beliefs at this information set and given the strategies at the other information sets. Mathematically, this means that, for every player i and $h \in H_i$,

$$\sum_{x \in h} \beta_h(x) U_i(\sigma \mid x) \geq \sum_{x \in h} \beta_h(x) U_i(\sigma \backslash a \mid x),$$

for all $a \in A(h)$. Here, $U_i(\sigma \mid x)$ denotes the expected payoff for player i if the game would start in the node x and the players would play according to σ. The expression $U_i(\sigma \backslash a \mid x)$ is defined in a similar way.

For the sake of convenience, we write $U_i(\sigma, \beta \mid h)$ instead of $\sum_{x \in h} \beta_h(x) U_i(\sigma \mid x)$.

By $B_h(\sigma, \beta, \Gamma)$, we denote the set of actions at h which maximize the expected payoff of the player controlling h with respect to the beliefs β. Formally,

$$\begin{aligned} B_h(\sigma, \beta, \Gamma) := \{a \in A(h) \mid\ & U_i(\sigma \backslash a, \beta \mid h) \\ & \geq U_i(\sigma \backslash b, \beta \mid h), \text{ for all } b \in A(h)\}. \end{aligned}$$

An assessment (σ, β) is called Bayesian consistent if, at every information set which is reached with strictly positive probability, the beliefs are derived according to Bayes' rule. So, for every h with $\mathbb{P}_\sigma(h) > 0$, it must hold that

$$\beta_h(x) = \mathbb{P}_\sigma(x) / \mathbb{P}_\sigma(h),$$

for every $x \in h$. Here, $\mathbb{P}_\sigma(x)$ and $\mathbb{P}_\sigma(h)$ denote the probabilities that the node x and the information set h are reached, respectively, if σ is played.

The assessment (σ, β) is called consistent if there is a sequence (σ^k, β^k) of Bayesian consistent assessments such that σ^k is completely mixed and (σ^k, β^k) converges to (σ, β). Obviously, every consistent assessment is Bayesian consistent.

Definition 8.1. An assessment (σ, β) is called a sequential equilibrium if it is sequentially rational and consistent.

Theorem 8.1. (See Kreps and Wilson (Ref. 4).) In an extensive form game, a BSP σ can be extended to a sequential equilibrium if and only if σ is a weakly perfect equilibrium.

Definition 8.2. We call an assessment (σ, β) ϵ-perfect if σ is completely mixed, (σ, β) is Bayesian consistent, and $\sigma_h(a) \leq \epsilon$ whenever $a \notin B_h(\sigma, \beta, \Gamma)$.

Obviously, for every ϵ-perfect assessment (σ, β), it holds that σ is an ϵ-perfect BSP, as defined in the previous section.

In the following theorem, we show that every sequential equilibrium can be approximated by a sequence of ϵ-perfect assessments of perturbed games.

Theorem 8.2. Let $\Gamma = \langle \mathscr{S}, u \rangle$ be an extensive form game, and let (σ, β) be a sequential equilibrium in Γ. Then, there is a sequence ϵ^k of positive numbers converging to 0, a sequence $\Gamma^k = \{\mathscr{S}, u^k\rangle$ of games converging to Γ, and a sequence (σ^k, β^k) of ϵ^k-perfect assessments in Γ^k converging to (σ, β).

Proof. Let (σ, β) be a sequential equilibrium in the game Γ with supporting sequence (σ^k, β^k). W.l.o.g. we may assume that $B_h(\sigma^k, \beta^k, \Gamma)$ is constant over k for every h.

The proof consists of two steps. In the first step, we construct a game Γ^k for every k such that $B_h(\sigma, \beta, \Gamma) \subset B_h(\sigma^k, \beta^k, \Gamma^k)$. In the second step, we prove that (σ^k, β^k) is an ϵ^k-perfect assessment in Γ^k for some ϵ^k, ϵ^k converges to 0, and Γ^k converges to Γ.

Step 1. Let k be fixed. We define a function f which transforms Γ into a new game $f(\Gamma)$. The transformation works as follows.

Let H^* be the collection of information sets h for which $B_h(\sigma, \beta, \Gamma)$ is not a subset of $B_h(\sigma^k, \beta^k, \Gamma)$. For every i and every terminal node z, we define the payoff $u_i^*(z)$ in the following way.

If z does not follow any information set in H_i^*, we define $u_i^*(z) := u_i(z)$. Otherwise, there is exactly one player i information set h in H^* which precedes z and is not followed by any other player i information set in H^*. Since $B_h(\sigma, \beta, \Gamma)$ is not a subset of $B_h(\sigma^k, \beta^k, \Gamma)$, the set $B_h(\sigma, \beta, \Gamma) \backslash \beta_h(\sigma^k, \beta^k, \Gamma)$ is not empty. Let b be an arbitrary but fixed action in $B_h(\sigma^k, \beta^k, \Gamma)$. Since $B_h(\sigma^k, \beta^k, \Gamma)$ is constant over k, we can choose b independently of k. For every $a \in B_h(\sigma, \beta, \Gamma) \backslash B_h(\sigma^k, \beta^k, \Gamma)$, we define

$$u_i^*(z) := u_i(z) + [U_i(\sigma^k \backslash b, \beta^k \mid h) - U_i(\sigma^k \backslash a, \beta^k \mid h)], \qquad \text{for all } z \in Z(a),$$

where $Z(a)$ denotes the set of terminal nodes which follow the action a.

For all other terminal nodes z which follow the information set h, we define $u_i^*(z) := u_i(z)$.

In this way, we obtain the game $\langle \mathscr{S}, u^* \rangle$, which we call $f(\Gamma)$.

By construction of the payoffs u^*, $B_h(\sigma, \beta, \Gamma)$ is a subset of $\beta_h(\sigma^k, \beta^k, f(\Gamma))$ for all player i information sets h in H^* which are not followed by any other player i information sets in H^*.

If such a player i information set h in H^* is followed by another player i information set $h' \notin H^*$, by perfect recall there is a unique action at h which leads to h'. By definition of the payoffs u^*, the set $B_{h'}(\sigma^k, \beta^k, \Gamma)$ does not change when we go from the game Γ to the game $f(\Gamma)$.

Combining the two observations above leads to the conclusion that the set H^* becomes strictly smaller when we go from Γ to $f(\Gamma)$ (if H^* is not empty at Γ, of course). So, by subsequently applying the transformation f, the set H^* becomes empty after finitely many times.

We define Γ^k to be the game which we obtain by subsequently applying f until H^* is empty. Since H^* is empty, it holds that $B_h(\sigma, \beta, \Gamma)$ is a subset of $B_h(\sigma^k, \beta^k, \Gamma^k)$ for all h.

Step 2. For every k, we define the number ϵ^k by

$$\epsilon^k := \max\{\sigma_h^k(a) \mid a \notin B_h(\sigma^k, \beta^k, \Gamma^k)\}, \qquad \text{if } \beta_h(\sigma^k, \beta^k, \Gamma^k) \neq A(h),$$

$$\epsilon^k := 1/k, \qquad \text{otherwise.}$$

By construction, (σ^k, β^k) is ϵ^k-perfect in Γ^k.

In order to show that ϵ^k converges to 0, choose an action a with $a \notin B_h(\sigma^k, \beta^k, \Gamma^k)$ for some k. Since $\beta_h(\sigma, \beta, \Gamma) \subset B_h(\sigma^k, \beta^k, \Gamma^k)$, it follows that $a \notin \beta_h(\sigma, \beta, \Gamma)$. Since (σ, β) is sequentially rational, $\sigma_h(a) = 0$. Hence, $\sigma_h^k(a)$ converges to 0, which implies that ϵ^k converges to 0.

It remains to show that Γ^k converges to Γ. Let k be fixed for the moment. In the step from Γ to $f(\Gamma)$, the player i payoff of a terminal node z only changes if z follows a player i action a such that $a \in B_h(\sigma, \beta, \Gamma)$, but $a \notin B_h(\sigma^k, \beta^k, \Gamma)$. If the payoff changes, the new payoff $u_i^*(z)$ is given by

$$u_i^*(z) := u_i(z) + [U_i(\sigma^k \backslash b, \beta^k \mid h) - U_i(\sigma^k \backslash a, \beta^k \mid h)],$$

where $b \in B_h(\sigma^k, \beta^k, \Gamma)$. Since we assumed that $B_h(\sigma^k, \beta^k, \Gamma)$ remains constant over k, it follows that $b \in B_h(\sigma, \beta, \Gamma)$. However, this implies that $U_i(\sigma \backslash b, \beta \mid h) = U_i(\sigma \backslash a, \beta \mid h)$. Using the fact that (σ^k, β^k) converges to (σ, β), it follows that

$$\lim_{k \to \infty} [U_i(\sigma^k \backslash b, \beta^k \mid h) - U_i(\sigma^k \backslash a, \beta^k \mid h)] = 0.$$

So, the difference $|u_i^*(z) - u_i(z)|$ tends to zero when k tends to infinity. Since the game Γ^k is obtained from Γ by applying the function f at most $|H|$ times, it follows that Γ^k converges to Γ.

This completes the proof. □

Theorem 8.3. The sequential equilibrium correspondence is a limit consistent assessment solution generated by the one-person solution function Φ^1 given by

$$\Phi^1(\epsilon)(\Gamma) := \{(\sigma, \beta) \mid \exists \Gamma' \text{ with } d(\Gamma, \Gamma') \leq \epsilon \text{ such that } (\sigma, \beta) \text{ is } \epsilon\text{-perfect in } \Gamma'\},$$

for every one-person game Γ and every $\epsilon > 0$.

Proof. Let φ be the sequential equilibrium correspondence, and let ψ be the limit consistent solution generated by Φ^1. Furthermore, we denote by $\Phi(\epsilon)$ the strongly consistent solution generated by $\Phi^1(\epsilon)$. Similarly to the proof of Theorem 7.1, it can be shown that

$$\Phi(\epsilon)(\Gamma) = \{(\sigma, \beta) \mid \exists \Gamma' \text{ with } d(\Gamma, \Gamma') \leq \epsilon \text{ such that } (\sigma, \beta) \text{ is } \epsilon\text{-perfect in } \Gamma'\},$$

for every game Γ.

Now, we show that $\varphi = \psi$.

First, we prove $\varphi(\Gamma) \subset \psi(\Gamma)$ for every game Γ.

Let $(\sigma, \beta) \in \varphi(\Gamma)$ for a game Γ, which means that (σ, β) is a sequential equilibrium in Γ. By Theorem 8.2, there is a sequence ϵ^k converging to 0, a sequence Γ^k converging to Γ, and a sequence (σ^k, β^k) of ϵ^k-perfect assessments in Γ^k converging to (σ, β). For every k, there is an $l(k)$ such that $d(\Gamma, \Gamma^{l(k)}) \leq 1/k$ and $\epsilon^{l(k)} \leq 1/k$. But then, obviously,

$$(\sigma^{l(k)}, \beta^{l(k)}) \in \Phi(1/k)(\Gamma), \qquad \text{for every } k.$$

Since $(\sigma^{l(k)}, \beta^{l(k)})$ converges to (σ, β), it follows that $(\sigma, \beta) \in \psi(\Gamma)$.

Finally, we show $\psi(\Gamma) \subset \varphi(\Gamma)$.

Let $(\sigma, \beta) \in \psi(\Gamma)$, so there is a sequence ϵ^k converging to 0 and a sequence $(\sigma^k, \beta^k) \in \Phi(\epsilon^k)(\Gamma)$ converging to (σ, β). So, for every k, there is a game Γ^k with $d(\Gamma, \Gamma^k) \leq \epsilon^k$ such that (σ^k, β^k) is ϵ^k-perfect in Γ^k. We show that (σ, β) is a sequential equilibrium.

The consistency of (σ, β) follows from the fact that (σ^k, β^k) is a sequence of completely mixed, Bayesian assessments converging to (σ, β).

It remains to show that (σ, β) is sequentially rational. To this purpose, consider an information set h and an action $a \in A(h)$ such that $a \notin B_h(\sigma, \beta, \Gamma)$. Since we may assume that $B_h(\sigma^k, \beta^k, \Gamma^k)$ remains constant over k, it follows that

$$a \notin B_h(\sigma^k, \beta^k, \Gamma^k), \qquad \text{for every } k.$$

Using the fact that (σ^k, β^k) is ϵ^k-perfect in Γ^k, it holds that

$$\sigma_h^k(a) \leq \epsilon^k, \qquad \text{for every } k,$$

implying that $\sigma_h(a) = 0$. It follows that (σ, β) is sequentially rational. □

References

1. Peleg, B., and Tijs, T., *The Consistency Principle for Games in Strategic Form*, International Journal of Game Theory, Vol. 25, pp. 13–34, 1996.
2. Patrone, F., Pieri, G., Tijs, S., and Torre, A., *On Consistent Solutions for Strategic Games*, Working Paper, Department of Mathematics, University of Genoa, Genoa, Italy, 1995.
3. Norde, H., Potters, J., Reijnierse, H., and Vermeulen, D., *Equilibrium Selection and Consistency*, Games and Economic Behavior, Vol. 12, pp. 219–225, 1996.
4. Kreps, D. M., and Wilson, R., *Sequential Equilibria*, Econometrica, Vol. 50, pp. 863–894, 1982.