Robust Equilibria under Linear Tracing Procedure

Hailin Liu
Department of Philosophy
Carnegie Mellon University, Pittsburgh, USA
hailinl@andrew.cmu.edu

Abstract

In Harsanyi and Selten's equilibrium selection theory, the linear tracing procedure has been used to model the hypothetical reasoning process of expectation formation. This paper reconsiders the linear tracing procedure from the perspective of the relationship between priors and Nash equilibria. A prior belongs to the source set of a Nash equilibrium if the linear tracing procedure based on this prior leads to that equilibrium. We show that for any Nash equilibrium, its source set is always nonempty and closed, but not generally convex. This paper also constructs an approach of iterative application of the linear tracing procedure to the auxiliary games that are used to model the hypothetical reasoning under the procedure. We present a notion of robustness of Nash equilibria based on this idea, by replacing uncertainty modelled by a single probability measure with uncertainty modelled by sets of probability measures. This approach attempts to capture the fact that players may not be sufficiently confident in the available information in order to single out one probability distribution that represents their initial beliefs about the other players' possible strategy choices.

Keywords. Equilibrium refinement, linear tracing procedure, stability, robustness, sets of probabilities.

1 Introduction

There are a variety of nontrivial games, with important applications in economics, which generate (sometimes infinitely) many different Nash equilibria. In game theory, a strategy profile is a Nash equilibrium if each player's choice is an optimal response to other players' choices. The fundamental assumption behind this definition is that one player's optimal choice maximizes her own expected utility given the other players' choices. The fact that there are typically multiple Nash equilibria seems to suggest that the equilibrium solution concept is too weak a criterion for predicting the players' behavior. Therefore, a great deal of effort has been devoted to refining the concept of Nash equilibrium by providing more stringent strategy-selection criteria. Examples of suggested equilibrium refinement concepts are Harsanyi and Selten's risk dominance ([2]), Kohlberg and Mertens's stability ([5]), Kreps and Wilson's sequentiality ([6]), and Selten's perfectness ([8]).

Harsanyi and Selten's idea of *risk dominance* captures the idea that, without knowing which equilibrium would be played, the players undergo an introspective process of expectation formation, which may eventually single out one equilibrium as less risky than another. This process is fully modelled by the so-called *linear tracing procedure*, which is thus the mathematical foundation of risk dominance. One of the basic assumptions of this model is that the uncertainty among all players' likely strategy choices is represented by a common prior strategy. However, it could be the case that the uncertainty among the Nash equilibria in question cannot be completely resolved as the assumed reasoning process proceeds. Nevertheless, the linear tracing procedure itself is a mathematical mechanism for modelling the players' hypothetical deliberation process about uncertainty. We shall later return to the linear tracing procedure and describe it in detail.

Moreover, we extend the framework of the linear tracing procedure to accommodate sets of probabilities as a representation of uncertainty. We then examine the possibility of iteratively applying the linear tracing procedure to a sequence of auxiliary games. This may be regarded as a minor generalization of the traditional game-theoretic framework, by only dropping the so-called "dogma of precision" ([9]), namely, that uncertainty should always be represented by a single probability measure. This enables us to assess the robustness of Nash equilibria in the traditional game-theoretic context, where uncertainty is represented in a more realistic manner.

To explain the basic ideas, consider the two-person coordination game described by Figure 1. In this game, player 1 has two pure strategies denoted by s_{11} and s_{12}, while player 2 also has two pure strategies denoted by s_{21} and s_{22}.

	s_{21}	s_{22}
s_{11}	$1,1$	$0,0$
s_{12}	$0,0$	$3,3$

Figure 1: Coordination Game

The game has two Nash equilibria in pure strategies, namely $E_1 = (s_{11}, s_{21})$ and $E_2 = (s_{12}, s_{22})$. It also has one Nash equilibrium in mixed strategies, $E_3 = ((\frac{3}{4}, \frac{1}{4}), (\frac{3}{4}, \frac{1}{4}))$, where the first and second pairs of numbers denote the probabilities assigned to player 1's and player 2's two pure strategies respectively. For convenience, the strategy space of the game can be described by the square $ABCD$ in Figure 2. Any point X of this square will represent a mixed strategy profile $\delta = ((q_{11}, q_{12}), (q_{21}, q_{22}))$. In particular, the horizontal distances XY and XZ will represent the probabilities q_{11} and $q_{12} = 1 - q_{11}$ respectively, and the vertical distances XV and XU will represent the probabilities q_{21} and $q_{22} = 1 - q_{21}$ respectively. Accordingly, the three Nash equilibria of the game can be represented by the corner points A and C, and the point E, as shown in Figure 2.

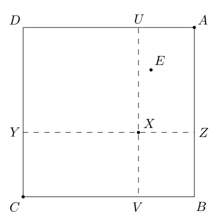

Figure 2: The Strategy Space

A natural question concerning this game is which of the three equilibria would be played. Harsanyi and Selten attempt to answer this question by employing the linear tracing procedure to examine the risk associated with different Nash equilibria when belief-uncertainty is represented by a single probability distribution. Here, we want to investigate the viability of Nash equilibria under the recursive application of the linear tracing procedure when uncertainty is modelled by *sets of probabilities.* We thereby hope to shed light on how traditional solution concepts can be informed by the notion of imprecise probabilities.

The remainder of the paper is structured as follows. Section 2 provides a formal description of the linear tracing procedure and some characterization results concerning source sets. In Section 3 we describe an approach which involves iterative application of the linear tracing procedure to a self-generated sequence of hypothetical games, where uncertainty is represented by sets of probabilities. On the basis of such a recursively applied linear tracing procedure, we then formalize and investigate a notion of *stability,* which measures the tenability of a prior strategy with respect to a certain Nash equilibrium under this procedure. The rest of this section extends the analysis of the linear tracing procedure to allow for representing uncertainty by sets of probabilities, and proposes a notion of *robustness* of Nash equilibria. Section 4 consists of concluding remarks and suggestions for future work along these lines.

2 Linear Tracing Procedure and Source Sets

The linear tracing procedure is a mathematical tool first introduced by Harsanyi ([3]), which underpins the equilibrium refinement concept proposed by Harsanyi and Selten ([2]). Informally speaking, it models how the players gradually update their strategy plans in light of what they know about the opponents' strategic reactions to their own expectations. The procedure can be regarded as a rational deliberation process of expectation formation, after which each player comes to choose a particular Nash equilibrium and to expect the others to make the same choice. The linear tracing process begins with a common probability distribution over all players' strategies, which represents their initial expectations about the other players' likely strategy choices. This way of setting up the initial belief state is often called the *Harsanyi doctrine* or, alternatively, the common prior assumption. Under such an assumption, it would seem that all players should adopt the best responses against the assumed common prior. And this typically gives rise to a different strategy combination that generally does not constitute a Nash equilibrium. Throughout the linear tracing procedure, all players gradually change their own tentative strategy plans, as well as their expectations about the other players' possible strategies, until they arrive at a certain Nash equilibrium. It has been shown ([2]) that the linear tracing procedure determines a unique Nash equilibrium for almost every game. In this section we shall explore the linear tracing procedure from a different perspective, focusing on characterizing the set of priors associated with a certain Nash equilibrium under the linear tracing procedure.

Let us begin with some basic notations and concepts. Let $G = \langle I, \{S_i\}, \{u_i\} \rangle_{i \in I}$ be a finite non-cooperative strategic form game, where I denotes a finite set of players, and S_i denotes the finite set of pure strategies of player i, and $u_i : S \to \mathbb{R}$ denotes a continuous payoff function of this player (where $S = \prod_{i \in I} S_i$).

As usual, we can extend the strategy space to include mixed strategies. In general, we let Δ_i represent the set of mixed strategies of player i, and similarly $\Delta = \prod_{i \in I} \Delta_i$. Likewise, the payoff function of a given player i can be extended in the standard way to the set of all mixed strategy combinations Δ, and we usually write $u_i(\delta)$ for the expected payoff of player i when $\delta \in \Delta$ is played. Let δ_{-i} denote the strategy combination $(\delta_1, \ldots, \delta_{i-1}, \delta_{i+1}, \ldots, \delta_n)$. For any $\delta_{-i} \in \Delta_{-i}$, the set of player i's best responses given δ_{-i} is defined as $B_i(\delta_{-i}) = \{\delta_i \in \Delta_i : u_i(\delta_i, \delta_{-i}) \geq u_i(\delta'_i, \delta_{-i}) \text{ for all } \delta'_i \in \Delta_i\}$. A strategy profile $\delta^* \in \Delta$ is a Nash equilibrium of G if and only if each player's strategy is a best response to the other players' strategies, i.e., $\delta^*_i \in B_i(\delta^*_{-i})$ for every player $i \in I$. Henceforth, the set of all Nash equilibria of the game G will be denoted by $NE(G)$. Also, we shall assume that some finite non-cooperative game G is already given.

The linear tracing procedure is a mapping φ from the strategy space Δ into the equilibrium set $NE(G)$. It transforms each strategy profile into a certain Nash equilibrium of the game G. In order to define the linear tracing procedure, consider a one-parameter family of auxiliary games $\Gamma^{t,p}$ with $t \in [0,1]$ and $p \in \Delta$. Each game $\Gamma^{t,p}$ is of the same structure as the original game G, except for the payoff functions. In $\Gamma^{t,p}$, for each $\delta \in \Delta$, each player i's payoff function $u_i^{t,p}$ satisfies

$$u_i^{t,p}(\delta_i, \delta_{-i}) = tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, p_{-i})$$

where u_i is player i's payoff function in the original game G. Obviously, $u_i^{1,p}(\delta_i, \delta_{-i}) = u_i(\delta_i, \delta_{-i})$, which implies that $\Gamma^1 = G$. While, for $t = 0$ the game $\Gamma^{0,p}$ decomposes into n independent and separate one-person maximization problems, one for each player. Now consider the equilibrium correspondence $\psi : t \to NE(\Gamma^{t,p})$ for $t \in [0,1]$ and $p \in \Delta$:

$$\psi = \{(t, \delta) \mid t \in [0,1], \delta \in NE(\Gamma^{t,p})\}$$

Let $\varphi = \varphi(G, p)$ be the graph of the correspondence ψ. Thus, any point x of graph φ has the mathematical form $x = (t, \delta)$, where t is a specific t value whereas δ is an equilibrium point of the corresponding auxiliary game Γ^t. Note that the graph is not always a function. It can be shown that the graph φ contains at least one distinguished path L, the so-called *feasible path*, which connects a starting point $x_0 = (0, \delta_0)$ with an end point $x_1 = (1, \delta^*)$. Hence, for a given game G and for a given prior strategy $p \in \Delta$, we call Θ a *linear tracing procedure* if it consists in selecting an outcome q^* by tracing a feasible path L from its starting point $x_0 = (0, \delta_0)$ to its end point $x_1 = (1, \delta^*)$. And δ^* will be called the *outcome* of the linear tracing procedure Θ. Figure 3 shows the graph of a linear tracing procedure for the game in Figure 1. For this linear tracing procedure, $B'''B''C''C'C$ is the unique feasible path and the equilibrium E_3 (point C) is the outcome.

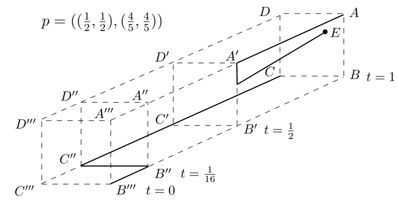

Figure 3: The linear tracing procedure based on p

The linear tracing procedure will always lead to at least one equilibrium, and select one unique equilibrium as the solution for "almost all" games[1]. The linear tracing procedure is called *feasible* if the graph $\varphi = \varphi(G, p)$ contains at least one feasible path L, and is called *well-defined* if X contains exactly one feasible path L. It can be shown that, for any possible pair (G, p), the linear tracing procedure is always feasible but is not always well-defined. In light of its fundamental importance, we state this result as follows.

Proposition 1. *([2]) For any possible choice of game G and prior vector p, the linear tracing procedure is always feasible. However, for some choices of G and p, the linear tracing procedure is not well-defined.*[2]

It is worth noting that the above proposition tells us nothing about the set of priors associated with a certain Nash equilibrium. There are several interesting questions that are worthy of further investigation. For instance, is the set non-empty, closed or convex? Before considering these problems, we first define the concept of *source sets* as follows.

Definition 2. For a given game G and a strategy $\delta^* \in NE(G)$, the *source set* for δ^*, denoted by $\Phi(\delta^*)$, is defined as the set of all prior strategies, based on which the linear tracing procedure yields the Nash equilibrium δ^* as outcome.

Our next proposition shows that for each Nash equilibrium δ^*, its source set always includes itself as an

[1] See Harsanyi and Selten ([2]) for a more detailed explanation of the term "almost all".

[2] The proof provided by Harsanyi and Selten is heavily dependent on the result showing that the logarithmic tracing procedure (also introduced by them) is always well defined. It is worth pointing out that a mathematical proof of feasibility of the linear tracing procedure can be easily derived from a theorem given in [7]. Using techniques from the field of algebraic geometry, Schanuel et al. first show that the logarithmic tracing procedure always connects the prior strategy to exactly one equilibrium point. Based on this result, one can argue that the feasibility of the linear tracing procedure is exactly a limit case of the feasibility of the logarithmic tracing procedure. More recently, Herings ([4]) directly shows the feasibility of the linear tracing procedure without appealing to the logarithmic tracing procedure. The two simple proofs provided by Herings are based on theorems related to the fixed-point theorems of Brouwer and Kakutani.

element, and is thus non-empty.

Proposition 3. *Let G be a finite non-cooperative game. For each Nash equilibrium δ^* of game G, δ^* belongs to its own source set, i.e., $\delta^* \in \Phi(\delta^*)$.*[3]

Next, we might ask whether the source sets are closed under the topology of pointwise convergence. To characterize the closure property of source sets, we must first introduce the concept of *pointwise convergence* on the strategy space of a game G, as well as that of a *limit point* of a set. Recall that we are considering only games with a finite number of pure strategies. Thus, the topology that we are considering is relatively easy to work with. We now define pointwise convergence as follows.

Definition 4. Let Δ be the strategy space of a finite game $G = \langle I, \{S_i\}, \{\pi_i\}\rangle$. A sequence $\{\delta^r\}$ *converges pointwise* to $\delta \in \Delta$, denoted by $\{\delta^r\} \to \delta$, if for each player $i \in I$, all $s_i \in S_i$, and all $\epsilon > 0$, there exist some k such that $|\delta_i^j(s_i) - \delta_i(s_i)| < \epsilon$ for each $j \geq k$. And δ is called the *limit point* of the sequence $\{\delta^r\}$.

Let us compare pointwise convergence and uniform convergence defined in the following sense. We say that a sequence $\{\delta^r\}$ *converges uniformly* over players' strategies to δ if for all $\epsilon > 0$, there exists some k such that for each player i, all $s_i \in S_i$, and all $j \leq k$, it holds that $|\delta_i^j(s_i) - \delta_i(s_i)| < \epsilon$. Clearly, uniform convergence is a stronger concept, and always implies pointwise convergence, but not vice versa. In our framework, however, pointwise convergence implies uniform convergence, since the set of players is finite, as well as each player's set of pure strategies.

In this paper, a point $p \in \Delta$ is called a *limit point* of the source set $\Phi(\delta^*)$ if there exists some sequence $\{p^r\}$ such that each element of $\{p^r\}$ belongs to $\Phi(\delta^*)$ and $\{p^r\} \to p$. We shall employ the notion of limit points to obtain a characterization of closed sets. Loosely speaking, a set A is *closed* in a space X if it contains all its limit points. Our main result is that the source sets of Nash equilibria are always closed. More formally:

Proposition 5. *Let G be a finite non-cooperative game and δ^* be a Nash equilibrium of G. If $p \in \Delta$ is a limit point of the source set $\Phi(\delta^*)$, then $p \in \Phi(\delta^*)$.*

We now give some definitions and lemmas that will be used in the proof of the foregoing proposition.

Definition 6. Let $G^r = \langle N^r, (S_i^r), (\pi_i^r)\rangle$ be a finite non-cooperative game with $r = 1, 2, \ldots$. A sequence of games $\{G^r\}$ converges to a game G if all the games in question have the same set of players $N^r = N$ and the identical set of pure strategies $S_i^r = S_i$, and the payoff function π_i^r converges uniformly to π_i, that is, for all $\epsilon > 0$, there exists some k such that for each player $i \in I$, all $s \in S$ and for all $j \geq k$, it holds that $|\pi_i^j(s) - \pi_i(s)| < \epsilon$.

Obviously, it follows from the definition that the sequence $\{G^r\}$ converges to G if all games under consideration share the same set of players $N^r = N$ and strategy space $\Delta^r = \Delta$ and, moreover, the payoff function u_i^r converges uniformly to u_i. We say that a game G is the limit game of a sequence of games $\{G^r\}$ if the sequence converges to G. The following lemma can be regarded as a special version of the well-known result ([1]) in game theory, which relates the Nash equilibria of a convergent sequence of games to the Nash equilibria of the limit game.

Proposition 7. *Let $\{G^r\}$ be a sequence of finite non-cooperative games converging to G. If the strategy profiles δ^r are Nash equilibria of G^r respectively with $\{\delta^r\} \to \delta$, then δ is a Nash equilibrium of game G.*

Now consider a sequence of prior strategies $\{p^r\}$, which converges to a prior strategy p. It is easy to verify that for each $t \in [0, 1]$ the sequence of games $\{\Gamma_{p^r}^t\}$ converges to the game Γ_p^t. In order to see this, let us recall that in game $\Gamma_{p^r}^t$ the payoff function $u_{i,p^r}^t : \Delta \to \mathbb{R}$ is given by $u_{i,p^r}^t(\delta_i, \delta_{-i}) = tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, p_{-i}^r)$, where u_i denotes the payoff function of the original game G. Since the payoff function u_i is assumed to be continuous, it directly follows from $\{p^r\} \to p$ that $\{u_{i,p^r}^t\}$ converges to $u_{i,p}^t$. Moreover, if a sequence $\{t^m\}$ converges to t where $t^m, t \in [0, 1]$, then it still holds that the sequence of games $\{\Gamma_{p^r}^{t^m}\}$ converges to the game Γ_p^t, since the sequence $\{u_i^{t^m, r}\}$ of payoff functions converges to u_i^t. Thus, it follows from the above lemma that the limit of Nash equilibria relative to the sequence of games is the Nash equilibrium of the limit game in both cases. These noteworthy facts turn out to play a significant role in the proof of the closure property of source sets. We now state the foregoing results as follows.

Corollary 8. *Suppose that $\{p^r\} \to p$ and $t \in [0, 1]$ where $p^r, p \in \Delta$. If δ^r are Nash equilibria relative to game $\Gamma_{p^r}^t$ with $\{\delta^r\} \to \delta$, then δ is a Nash equilibrium of game Γ_p^t.*

Corollary 9. *Suppose that $\{p^r\} \to p$ and $\{t^m\} \to t$ where $t^m, t \in [0, 1]$ and $p^r, p \in \Delta$. If $\delta^{m,r}$ are Nash equilibria relative to game $\Gamma_{p^r}^{t^m}$ with $\{\delta^{m,r}\} \to \delta$, then δ is a Nash equilibrium of game Γ_p^t.*

Let us turn to the convexity of source sets. In some games, the source sets are convex, in the sense that any mixture combination between two strategies from a source set also belongs to the source set. It is not the

[3] Proofs not given in the main text can be found in the Appendix.

case, however, that convexity holds in general. This point can be easily illustrated by considering the coordination game in Figure 1. As mentioned before, this game has three Nash equilibria, i.e., E_1, E_2, and E_3. It can be verified by simple computation that the source sets of these three Nash equilibria can be described as shown in Figure 4. In particular, the source set of E_1 consists of all points within the area $AHEF$, the source set of E_2 consists of all points within the area $BCDFEH$, and the source set of E_3 consists of all points lying on the border segment HEF. Clearly, the source sets of the equilibria E_1 and E_3 are not convex.

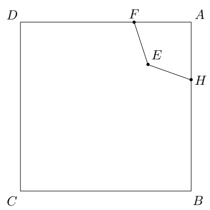

Figure 4: The source Sets

3 Robustness to Sets of Probabilities

As we mentioned above, the purpose of the linear tracing procedure is to provide a rational and effective mechanism for selecting a Nash equilibrium as the solution of non-cooperative game. Now let us recall how the linear tracing procedure works. The linear tracing procedure is not merely an examination of the game itself. Instead, we invoke a sequence of hypothetical games to investigate how the equilibria of the original game behave in auxiliary games. It is worthwhile to note that these auxiliary games are also non-cooperative, and typically resemble the original game in other respects. In other words, the auxiliary games themselves are also amenable to the linear tracing procedure. Therefore, it seems reasonable to apply the linear tracing procedure recursively to solve these auxiliary games. In this section, we will investigate such recursive applications of the linear tracing procedure.

Consider the finite non-cooperative strategic form game $G = \langle I, \{S_i\}, \{u_i\}\rangle_{i\in I}$ and the linear tracing procedure for G as described in Section 2. Let $p \in \Delta$ be one prior strategy and define a one-parameter family of auxiliary games Γ_p^t with $t \in [0,1]$. Generally speaking, any such game Γ_p^t will be a game of the same structure as the original G except for the payoff functions. More precisely, Γ_p^t can be specified as $\Gamma_p^t = \langle I, \{S_i\}, \{u_i^t\}\rangle_{i\in I}$, where, for each $\delta \in \Delta$, the payoff function u_i^t is defined by

$$u_i^t(\delta_i, \delta_{-i}) = tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, p_{-i}).$$

Now let us consider an application of the linear tracing procedure to the game Γ_p^t for some $t \in [0,1]$. That is, for some $t \in [0,1]$ assume that Γ_p^t is the original game, denoted by G^t. Define a one-parameter family of auxiliary games $\Lambda^{t'}$ with $t' \in [0,1]$ as follows. Given a prior strategy $p' \in \Delta$, game $\Lambda^{t'}$ can be defined as $\Lambda_{p'}^{t'} = \langle I, \{S_i\}, \{\mu_i^{t'}\}\rangle_{i\in I}$, where, for each $\delta \in \Delta$, the payoff function $\mu_i^{t'}$ satisfies

$$\mu_i^{t'}(\delta_i, \delta_{-i}) = t'u_i^t(\delta_i, \delta_{-i}) + (1-t')u_i^t(\delta_i, p'_{-i}).$$

It was shown in Proposition 2 that the source set of any equilibrium point for the original game G is not empty. In order to examine the robustness of an equilibrium, we focus on how its source set changes when applying the linear tracing procedure recursively to the auxiliary games.

Before entering into further analysis of the source set, we first consider one interesting case: what happens if, throughout the recursive application of the linear tracing procedure, we always use the same prior as a starting point? Suppose that δ^* is an equilibrium of game G, and p is an element of the source set of δ^*, that is, $p \in \Phi(\delta^*)$. Now consider the games Γ_p^t, which can be represented as $\Gamma_p^t = \langle I, \{S_i\}, \{u_i^t\}\rangle_{i\in I}$, where, for each $\delta \in \Delta$, the payoff function u_i^t satisfies

$$u_i^t(\delta_i, \delta_{-i}) = tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, p_{-i}).$$

Then apply the linear tracing procedure to game Γ_p^t with the same prior p. As mentioned above, we have to consider a new one-parameter class of auxiliary games $\Lambda_p^{t'} = \langle I, \{S_i\}, \{\mu_i^{t'}\}\rangle_{i\in I}$ with $t' \in [0,1]$, where, for each $\delta \in \Delta$, the payoff function $\mu_i^{t'}$ satisfies

$$\mu_i^{t'}(\delta_i, \delta_{-i}) = t'u_i^t(\delta_i, \delta_{-i}) + (1-t')u_i^t(\delta_i, p_{-i}).$$

Obviously, $\Lambda_p^0 = \Gamma_p^0$, since the payoff functions are identical, that is, $\mu_i^0 = u_i^0$. For the same reason, we have $\Lambda_p^1 = \Gamma_p^t$. Thus, the class of auxiliary games $\Lambda_p^{t'}$ is a subset of the family of auxiliary games Γ_p^t with respect to the game G. In other words, when considering the linear tracing procedure applied to the game Γ_p^t, we are in fact investigating a smaller subset of the family of auxiliary games generated by the linear tracing procedure applied to the *original* game. Hence, we can show that whenever δ^* is an equilibrium point of game Γ_p^t, the linear tracing procedure starting from p always feasibly leads to δ^*. On the basis of this observation, the following result is immediate:

Theorem 10. *Let $G = \langle I, \{S_i\}, \{u_i\}\rangle_{i\in I}$ be a finite non-cooperative strategic form game, and let δ^* be one equilibrium point of G. If $p \in \Phi(\delta^*)$ and δ^* is a Nash equilibrium of game Γ_p^t, then p is an element of the source set of δ^* with respect to game Γ_p^t.*

Now we can ask: given a certain equilibrium, under what constraint would a prior strategy belong to its source set throughout the recursive application of the linear tracing procedure? It turns out that, whenever the equilibrium δ^* under consideration maximizes the expected payoff for each player with respect to the prior strategy p, then p is always an element of the source set of δ^* pertaining to game Γ_p^t for any $t \in [0, 1]$. More precisely, we have:

Theorem 11. *Let $G = \langle I, \{S_i\}, \{u_i\}\rangle_{i\in I}$ be a finite non-cooperative strategic form game, and let δ^* be an equilibrium point of G. For any $t \in [0, 1]$, if δ^* maximizes all players' expected payoffs with respect to the prior strategy p, then $p \in \Phi^t(\delta^*)$ with respect to Γ_p^t.*

Clearly, the prior strategy determines how far into the recursive application of the procedure the prior p remains an element of a source set of the same equilibrium. This suggests that, when recursively applying the linear tracing procedure, the duration in which the prior strategy p belongs to the same source set can be considered a measure of the stability of p. According to the foregoing theorem, when a certain Nash equilibrium δ^* maximizes all the players' expected payoffs with respect to a prior p, then p is the most stable element of the source set of δ^*. This is because the linear tracing procedure that begins with p always points to the same equilibrium δ^*. Thus, we can say that such a prior strategy p is *maximally stable* with respect to δ^*. We now define the measure of stability.

Definition 12. Given a finite non-cooperative strategic form game G and one equilibrium δ^*, the *stability* of a prior strategy $p \in \Delta$ with respect to δ^* is a real-valued function γ on $\Phi(\delta^*)$, which is defined as $\gamma(p, \delta^*) = 1 - t^*$, where t^* is the smallest value of t such that $p \in \Phi^t(\delta^*)$. We say that p is *maximally stable* with respect to δ^* when $\gamma(p, \delta^*) = 1$. The set consisting of all such prior strategies is called the maximally stable source set of δ^*.

To illustrate the notion of stability with respect to Nash equilibria, consider the coordination game mentioned in section 1. The general description is shown in Figure 5. In particular, the source set of E_1 (the area $AHEF$) can be divided into two parts: the area $AGEI$ contains all the maximally stable priors with respect to the equilibrium E_1, and the remaining area consists of the priors with $\gamma(p, E_1) < 1$. Similarly, the source set of E_2 is composed of the maximally stable source set $ENCM$ and the rest of the non-maximally stable prior strategies with $\gamma(p, E_2) < 1$. In contrast, the maximally stable source set of the mixed strategy equilibrium E_3 consists of only *one* prior strategy, namely itself. All other prior strategies in its source set are not maximally stable with respect to E_3.

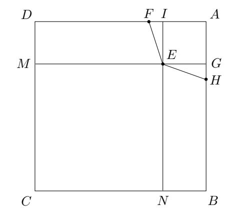

Figure 5: The Stability of Prior strategies

For each element p of the source set of a certain equilibrium, there is a measure of the stability of p, indicating how long the prior p remains associated with the same source set in the recursive application of the procedure. Recall from the previous section, that the source set of each equilibrium is non-empty and closed. Note that all the intermediate games invoked by the linear tracing procedure are closely related to the original game. Thus, the stability of the prior strategies under the recursive application of the linear tracing procedure indicates the robustness of the Nash equilibria with respect to the original game. On the basis of t-value as a measure for the stability of the priors, it is reasonable to employ the stability measure to compare the robustness of the Nash equilibria of a given game.

Before we define the measure of robustness, let us informally motivate the very idea of introducing sets of probabilities into the game-theoretic framework. As mentioned above, there are many games that have multiple Nash equilibria. This fact has given rise to a wide discussion of the equilibrium refinement problem in game theory. We believe that the linear tracing procedure is an appropriate mathematical mechanism for comparing Nash equilibria, since it accords with a common intuition regarding relative degrees of risk associated with different Nash equilibria. As mentioned above, the linear tracing procedure invokes a family of auxiliary games closely resembling the original game in question. Thus, by applying it recursively to the auxiliary games, we provide further information about the original game. In fact, it indicates the stability of one prior strategy with respect to a certain Nash equilibrium.

On the other hand, the linear tracing procedure assumes that all players employ the same probability distribution to represent their initial beliefs about the other players' likely strategy choices. In their analysis, Harsanyi and Selten choose a single probability distribution to express the uncertainty among players regarding which strategy the others would adopt. In decision theory, however, there are numerous suggested methods to represent decision makers' uncertainty besides using a single probability. Some salient approaches involve modelling uncertainty using sets of probabilities, upper and lower probabilities, upper and lower previsions, and belief functions ([9]).

Here we want to employ a non-trivial, convex set of probability measures $\mathcal{P}$ to represent all players' ignorance about the other players' likely behaviors. More precisely, we want to extend Harsanyi and Selten's framework framework by employing sets of prior strategies, rather than one single prior, to represent players' initial beliefs. Note that each of the prior strategies under consideration leads to a certain equilibrium under the linear tracing procedure, which simply means that it belongs to the source set of that equilibrium. Moreover, when we recursively apply the linear tracing procedure to the auxiliary games, we can determine the stability measure associated with each of the prior strategies with respect to a certain equilibrium. Based on these measures of stability, we can now define the robustness of equilibria with respect to a set of prior strategies as follows.

Definition 13. Let G be a finite non-cooperative strategic form game, and let the players' initial beliefs about the other players' possible behaviors be represented by a set of prior strategies $\mathcal{P}$. The robustness of an equilibrium δ^* with respect to $\mathcal{P}$ is defined as $R(\delta^*, \mathcal{P}) = \min_{p \in \mathcal{P}} \gamma(p, \delta^*)$, i.e., the minimum stability index associated with the priors with respect to $\mathcal{P}$.

This notion can be regarded as a further refinement of Nash equilibria based on the stability measures of the priors under the iterative application of the linear tracing procedure. Basically, one equilibrium is more robust than another if the least stability index associated with the elements of its source set is higher than the one associated with the other equilibrium under the recursive application of the linear tracing procedure. Given certain sets of prior strategies, we employ the maximin principle to assess the robustness of equilibria, where uncertainty is represented by sets of probabilities. That is, we select the equilibrium that maximizes the possible minimum stability of the prior strategies in its source set.

In order to illustrate the idea, let us consider an ϵ-*contaminated class* of probabilities given by $M = \{(1-\epsilon)P + \epsilon Q, Q \in \mathscr{P}\}$, where P is a particular prior distribution and ϵ is a fixed number in $[0, 1]$. $\mathscr{P}$ is the class of probability distributions that represents the possible deviations of the prior P. And the fixed ϵ represents the degree of contamination that players want to introduce into P.

Example 14. (ϵ-contamination under equilibria coordination) Consider the game described above. Suppose that all players believe that they will play the game in a coordination way. That is, the players collectively choose some mixed strategies involving the equilibria E_1, E_2, and E_3. Let $\mathscr{P} = \{Q : Q = p_1E_1 + p_2E_2 + p_3E_3,\ \text{where}\ p_1 + p_2 + p_3 = 1\}$. Figure 6 (the dark segment on the diagonal AC) illustrates the corresponding ϵ-contaminated class $\mathcal{P} = \{(1-\epsilon)P + \epsilon Q, Q \in \mathscr{P}\}$ when $P(E_1) = \frac{7}{10}$, $P(E_2) = \frac{1}{5}$, $P(E_3) = \frac{1}{10}$ and $\epsilon = 0.2$, which represents the players' initial beliefs. Observe that each prior in $\mathcal{P}$ is maximally stable with respect to either E_1, E_2, or E_3. Thus, $R(E_1, M) = R(E_2, M) = R(E_3, M) = 1$. This suggests that in this game when all players believe they will coordinate on an equilibrium, the notion of maximin robustness proposed here does not distinguish among these equilibria.

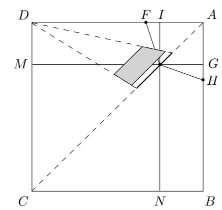

Figure 6: ϵ-contaminated class

Example 15. (Coordination failure) Reconsider the same game again. Suppose that all players initially believe that they will fail to coordinate with small probability. More precisely, the players believe that they will mostly choose a strategy from the ϵ-contaminated class $\mathcal{P}$, or otherwise adopt the strategy $D = (s_{12}, s_{21})$ with some probability in $[0.05, 0.2]$. In this case, the players' initial beliefs can be represented by $\mathcal{P}' = \{(1-\alpha)\mathcal{P} + \alpha D, 0.05 \leq \alpha \leq 0.2\}$, which is illustrated by the grey area in Figure 6. Simple calculation gives us the following result: for all $p \in \mathcal{P}$, $\frac{57}{64} = 0.8\bar{9} \leq \gamma(p, E_1) \leq 1$, $\frac{375}{482} = 0.7\bar{8} \leq \gamma(p, E_2) \leq 1$, and $\gamma(p, E_3) = 0$. Thus, $R(E_1, \mathcal{P}') = 0.8\bar{9}$, $R(E_2, \mathcal{P}') = 0.7\bar{8}$, and $R(E_3, \mathcal{P}') = 0$. Thus, according to the notion of maximin robustness we defined, E_1 is the most robust equilibrium with respect to $\mathcal{P}'$.

4 Conclusion

Why should one Nash equilibrium be more likely to be played than any other in a game? There is a large literature providing different criteria for selecting a particular equilibrium among many. Harsanyi and Selten propose a notion of risk dominance based on the linear tracing procedure. Here we do not attempt to address the issue of whether risk dominance is an appropriate criterion for equilibrium comparison. Instead, we extend the manner in which the linear tracing procedure is applied, as well as to replace a single probability distribution with sets of priors to represent players' uncertainty about strategy choices.

We first showed that, for any Nash equilibrium, its source set is always nonempty and closed, but not necessarily convex. We then considered a recursive application of the linear tracing procedure onto sequences of hypothetical games generated by the procedure itself. Based upon this, we formalized a notion of stability of priors, as well as a sufficient condition for characterizing the set of maximally stable priors with respect to certain Nash equilibria. Finally, we introduced a notion of maximin robustness of equilibria by considering the recursively applied linear tracing procedure when uncertainty is represented by a set of probabilities rather a single probability measure. We employed the maximin criterion to measure the robustness index associated with the Nash equilibria related to certain sets of prior strategies under the recursive procedure. The approach considered here is thus meant to demonstrate how one might accommodate the idea of imprecise probabilities within the traditional game-theoretic framework.

In the future, we intend to continue our examination of robustness in more general games, for instance symmetric games. This would provide further characterization about how sets of probabilities can be incoporated within game-theoretic framework. We shall also compare this approach to other existing theories of equilibrium refinement to investigate the relationships among them. Moreover, we shall consider the possibility of developing a new solution concept based on sets of probabilities that possesses more appealing features.

Acknowledgements

I am particularly grateful to my supervisor, Professor Teddy Seidenfeld for valuable discussions, supervision and encouragement. And I would like to thank Assistant Professor Kevin Zollman and two anonymous referees for detailed and useful comments.

Appendix

Proof of Proposition 3: Let G be a finite noncooperative game. Assume that $\delta^* \in \Delta$ is a Nash equilibrium of G. We must show that $\delta^* \in \Phi(\delta^*)$. In fact, we need to show that for the prior strategy δ^* the linear tracing procedure will feasibly select δ^* as the outcome.

Consider the games Γ^t invoked by the linear tracing procedure based on δ^*. We will show that δ^* is a Nash equilibrium for any game Γ^t for $t \in [0,1]$, which of course suffices to establish the result that $\delta^* \in \Phi(\delta^*)$.

According to the definition of Nash equilibrium, we have that $\delta_i^* \in B_i(\delta_{-i}^*)$ for every player i. More precisely, it means that for any player i

$$\delta_i^* \in B_i(\delta_{-i}^*) = \{\delta_i \in \Delta_i \mid u_i(\delta_i, \delta_{-i}^*) \geq u_i(\delta_i', \delta_{-i}^*), \forall \delta' \in \Delta_i\}.$$

Consider first the separable game Γ^0. Since the strategy δ^* is the prior strategy, the payoff functions for each player i are given by $u_i^0(\delta_i, \delta_{-i}) = u_i(\delta_i, \delta_{-i}^*)$ for any $\delta_i \in \Delta_i$. Thus, for any $\delta \in \Delta$ the best response correspondence B_i^0 can be represented by $B_i^0(\delta_{-i}) = \{\delta_i \in \Delta_i \mid u_i(\delta_i, \delta_{-i}^*) \geq u_i(\delta_i', \delta_{-i}^*),$ for all $\delta_i' \in \Delta_i\}$. Clearly, it implies $B_i^0(\delta_{-i}) = B_i(\delta_{-i}^*)$ for any $\delta \in \Delta$. We then have that $\delta_i^* \in B_i^0(\delta_{-i}^*)$ for every player i, and thus $\delta^* \in NE_{\Gamma^0}$.

Now, let us consider the generic games Γ^t. First, the payoff functions for each player i are given by $u_i^t(\delta_i, \delta_{-i}) = tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, \delta_{-i}^*)$. Thus, the best response correspondence B_i^t can be represented as follows.

$$B_i^t(\delta_{-i}) = \{\delta_i \in \Delta_i \mid tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, \delta_{-i}^*) \geq tu_i(\delta_i', \delta_{-i}) + (1-t)u_i(\delta_i', \delta_{-i}^*), \forall \delta_i' \in \Delta_i\}.$$

In particular, $B_i^t(\delta_{-i}^*) = \{\delta_i \in \Delta_i \mid tu_i(\delta_i, \delta_{-i}^*) + (1-t)u_i(\delta_i, \delta_{-i}^*) \geq tu_i(\delta_i', \delta_{-i}^*) + (1-t)u_i(\delta_i', \delta_{-i}^*), \forall \delta_i' \in \Delta_i\}$. It follows that

$$B_i^t(\delta_{-i}^*) = \{\delta_i \in \Delta_i \mid u_i(\delta_i, \delta_{-i}^*) \geq u_i(\delta_i', \delta_{-i}^*), \forall \delta_i' \in \Delta_i\},$$

which is independent of the value of t. This means that $B_i^t(\delta_{-i}^*) = B_i(\delta_{-i}^*)$ for each player i.

Hence we have that $\delta_i^* \in B_i^t(\delta_{-i}^*)$ for every player i, and thus $\delta^* \in NE(\Gamma^t)$ for $t \in [0,1]$. We can therefore conclude that $\delta^* \in \Phi(\delta^*)$. ■

In order to prove our main result (Proposition 5), we need to show some properties concerning Nash equilibria of convergent sequences of games. Since the results are required in proving the main result, we first present their proofs.

Proof of Proposition 7: Suppose that $\{G^r\}$ converges to game G. Assume for contradiction that δ is not a Nash equilibrium of the limit game G. Then there exists some player i with some $t_i \in \Delta_i$ such that

$$u_i(\delta) < u_i(t_i, \delta_{-i}).$$

First, note that $\{G^r\}$ converges to G which thus implies that u_i^r converges uniformly to u_i. Thus we can find a continuous approximation of u_i, denoted by u_i^j, such that

$$u_i^j(\delta) < u_i^j(t_i, \delta_{-i}).$$

Moreover, we know that the sequence $\{\delta^r\}$ converges pointwise to δ, and thus converges uniformly to δ. Hence, when j is large enough, we have that

$$u_i^j(\delta^j) < u_i^j(t_i, \delta_{-i}^j),$$

which contradicts the assumption that δ^j is a Nash equilibrium of game G^j. Therefore, δ must be a Nash equilibrium of the limit game G. ■

Proof of Corollary 8: Suppose that $p^r \to p$ and $t \in [0,1]$. First, we show that the sequence of games $\{\Gamma_{p^r}^t\}$ converges to game Γ_p^t. As described in the linear tracing procedure, the sets of players and the strategy spaces, denoted by I and Δ respectively, are all the same as the original game G. Note that the payoff function of game $\Gamma_{p^r}^t$ is given by

$$u_{i,p^r}^t(\delta_i, \delta_{-i}) = tu_i(\delta_i, \delta_{-i}) + (1-t)u_i(\delta_i, p_{-i}^r)$$

where u_i is the player i's payoff function in the original game G. Note that the first term on the right side is independent of p^r. And since it is assumed that u_i is continuous, it thus follows from $\{p^r\} \to p$ that $\{u^t_{i,p^r}\}$ converge to $u^t_{i,p}$. Hence, the sequence $\{\Gamma^t_{p^r}\}$ converges to Γ^t_p. Moreover, it is assumed that δ^r are Nash equilibria relative to game $\Gamma^t_{p^r}$ with $\{\delta^r\} \to \delta$. Therefore by Proposition 7, δ is a Nash equilibrium of the limit game Γ^t_p. ■

Proof of Corollary 9: Suppose that $\{p^r\} \to p$ and $\{t^m\} \to t$ where $t^m, t \in [0,1]$ and $p^r, p \in \Delta$. In order to apply Proposition 7, we have to show that the sequence of games $\{\Gamma^{t^m}_{p^r}\}$ converges to game Γ^t_p. Similarly, we have that the sets of players and the strategy spaces are all the same as the original game G, denoted by I and Δ respectively. Now consider the payoff function of game $\Gamma^{t^m}_{p^r}$ which is defined as

$$u^{t^m}_{i,p^r}(\delta_i, \delta_{-i}) = t^m u_i(\delta_i, \delta_{-i}) + (1-t^m) u_i(\delta_i, p^r_{-i})$$

where u_i is the player i's payoff function in the original game G. Note that u_i is assumed to be continuous. And since $\{t^m\} \to t$ and $\{p^r\} \to p$, it implies that $\{u^{t^m}_{i,p^r}\}$ converge to $u^t_{i,p}$. Thus, according to the defintion of convergent sequence of games, we have that the sequence of games $\{\Gamma^{t^m}_{p^r}\}$ converges to game Γ^t_p. And since it is assumed that $\delta^{m,r}$ are Nash equilibria relative to game $\Gamma^{t^m}_{p^r}$ with $\{\delta^{m,r}\} \to \delta$, it follows from Proposition 7 that δ is a Nash equilibrium of the limit game Γ^t_p. ■

With the aid of the foregoing results, we can now present the proof of our main result in section 2.

Proof of Proposition 5: Let $p \in \Delta$ be a limit point of the source set $\Phi(\delta^*)$. This means that there exists some sequence of priors $\{p^r\}$ such that $p^j \in \Phi(\delta^*)$ for each $p^j \in \{p^r\}$ and $\{p^r\}$ pointwise converges to p, i.e., $\{p^r\} \to p$. According to the definition, $p^j \in \Phi(\delta^*)$ means that there exists a feasible path, denoted by L_{p^j}, connecting the starting point $\delta^0_{p^j}$ and the end point δ^*, where $\delta^0_{p^j}$ is a Nash equilibrium of the game $\Gamma^0_{p^j}$ corresponding to the separable game that used p^j as the prior strategy. Here for $t \in [0,1]$ and $p^j \in \{p^r\}$, we let $\Gamma^t_{p^j}$ denote the game generated by using p^j as the prior strategy, and let $\delta^t_{p^j}$ denote the Nash equilibrium point(s) of game $\Gamma^t_{p^j}$ appearing on the feasible path L_{p^j}.

We must show that there exists a feasible path L_p for p which connects some equilibrium point(s) of game Γ^0_p to δ^*. Clearly, the set of t-values T is totally bounded, and thus can be covered by finitely many sets, each of which is centered at a point of T with diameter at most ϵ, for any $\epsilon > 0$. Now let $\epsilon > 0$. The set T can then be written as the union of finitely many sets with diameters $< \epsilon$. Let us denote these sets by $T_1, \ldots, T_m$. To show the existence of such a feasible path L_p, let us consider whether infinitely many feasible paths of $\{L_{p^j}\}$ have continuous segments of equilibrium points for the corresponding games at these sets $T_1, \ldots, T_m$.

Case 1: There is no such set where infinitely many feasible paths of $\{L_{p^j}\}$ have continuous segments of equilibrium points. This implies that either all the feasible paths are straight lines or only finite many feasible paths have continuous segments of equilibrium points somewhere.

First, consider the former case. Since all the feasible paths L_{p^j} are straight lines passing from some points to the same point δ^*, these feasible paths thus can be fully characterized by the corresponding slopes of the lines. Note that we are considering only finite games. It thus follows that the strategy space can be viewed as a subset of a finite-dimensional Euclidean space R^n, and the slopes of the feasible paths must be bounded to a certain region. Recall that by the Bolzano–Weierstrass theorem, each bounded sequence in R^n has a convergent subsequence. Thus, there exists a convergent subsequence of the slopes of the straight feasible paths, which means that there exists some subsequence of the straight feasible paths converging to a straight line determined by the limit slope, denoted by L_p. And we know that each feasible path corresponds to a prior strategy in the sequence $\{p^r\}$, and, therefore, that convergent subsequence of the straight feasible paths also correspond to a convergent subsequence of $\{p^r\}$, denoted by $\{p^{r'}\}$. Of course, the subsequence $\{p^{r'}\}$ must converge to the same limit as $\{p^{r'}\}$, that is $\{p^{r'}\} \to p$.

Now consider the subsequence $\{p^{r'}\}$ converging to p. As pointed out above, for each $t \in [0,1]$ the sequence of games $\{\Gamma^t_{p^{r'}}\}$ converges to Γ^t_p. Since that subsequence of the straight feasible paths converges to a straight line L_p, the sequence of Nash equilibria $\{\delta^t_{p^{r'}}\}$ thus converges to δ^t_p for each $t \in [0,1]$. It thus follows from Corollary 9 that δ^t_p must be a Nash equilibrium of game Γ^t_p, which shows that each point of the straight line L_p is one Nash equilibrium of the corresponding game Γ^t_p. Hence, the straight line L_p is a feasible path for p which connects some starting point belonging to Γ^0_p to the end point δ^*.

Now, consider the latter case, where only finite many feasible paths have continuous segments of equilibrium points somewhere. We can always ignore such feasible paths and only consider the other infinitely many feasible paths that are straight lines. Since each feasible path is associated with a prior p_j, there thus exists some subsequence $\{p^{r'}\}$ corresponding to these infinitely many feasible paths. Hence, the above argument can be applied to this subsequence $\{p^{r'}\}$. Therefore, in this case we have that there exists one feasible path L_p for p as well.

Case 2: There exists one and only one set, say T_k, where infinitely many feasible paths of $\{L_{p^j}\}$ have continuous segments of equilibrium points. Assume that the set T_k is centered at t_k. Similarly, we have that there exists some subsequence $\{p^{r'}\}$ corresponding to these infinitely many feasible paths, and $\{p^{r'}\} \to p$. We are going to show that there exists a feasible path L_p for p.

Note that all or infinitely many feasible paths of $\{L_{p^j}\}$ do not have any continuous segments in the interval $(t_k, 1]$. This means that there are infinitely many feasible paths that are straight lines in $(t_k, 1]$. A similar argument as that of case 1 shows that these infinitely many feasible paths converge to L_p in $(t_k, 1]$.

Now consider the set T_k. As described above, T_k is a set centered at t_k with diameter$< \epsilon$ where ϵ is arbitrarily small. We have that $\{p^{r'}\} \to p$, and the corresponding feasible paths of $\{p^{r'}\}$ have continuous segments of equilibrium points at T_k. Then, according to the Bolzano–Weierstrass theorem, there exists a subsequence $\{p^{r''}\}$ of $\{p^{r'}\}$ such that $\{p^{r''}\}$ uniformly converges to p with $\{t^m\} \to t_k$. Thus the corresponding continuous segments of equilibrium points uniformly converge to one continuous segment of equilibrium points for the game $\Gamma_p^{t_k}$. This implies that there exists one continuous segment of equilibrium points of the game $\Gamma_p^{t_k}$. Next, we show that the coming-in and coming-out points are exactly the two endpoints of this continuous segment. The reason is that the coming-in and coming-out points should be the limits of the coming-in and coming-out points of the infinitely many paths corresponding to $\{p^{r''}\}$, which must coincide with the the limits of the endpoints of these infinitely many paths. So far we have established that there exists a continuous path from t_k to 1, which has a continous segment at t_k.

Note again that there are infinitely many feasible paths of $\{L_{p^j}\}$ that are straight line in $[0, t_k)$. By a similar argument as in case 1, these infinitely many feasible paths converge to L_p in $[0, t_k)$. Taking these together, we can therefore conclude that there exists a feasible path L_p for p.

We can employ the above argument to examine all the sets $T_1, \dots T_k$. Since these sets are finite, we know that there exists a feasible path L_p for p, which implies that $p \in \Phi(\delta^*)$. ∎

Proof of Theorem 11: Assume that δ^* is an equilibrium of the game G, and δ^* maximizes all players' expected payoff with respect to p. In order to check whether $p \in \Phi^t(\delta^*)$, let us regard Γ_p^t as the original game, which can be represented as $\Gamma_p^t = \langle I, \{S_i\}, \{u_i^t\} \rangle_{i \in I}$, where, for each $\delta \in \Delta$, the payoff function u_i^t is defined as

$$u_i^t(\delta_i, \delta_{-i}) = t u_i(\delta_i, \delta_{-i}) + (1-t) u_i(\delta_i, p_{-i}).$$

We then consider a new one-parameter class of auxiliary games $\Lambda_p^{t'} = \langle I, \{S_i\}, \{\mu_i^{t'}\} \rangle_{i \in I}$ with $t' \in [0,1]$, where, for each $\delta \in \Delta$, the payoff function $\mu_i^{t'}$ is given by

$$\mu_i^{t'}(\delta_i, \delta_{-i}) = t' u_i^t(\delta_i, \delta_{-i}) + (1-t') u_i^t(\delta_i, p_{-i}).$$

Obviously, $\Lambda_p^0 = \Gamma_p^0$, since the payoff functions are identical, that is, $\mu_i^0 = u_i^0$; and $\Lambda_p^1 = \Gamma_p^t$ for the same reason. In view of this, the class of auxiliary games $\Lambda_p^{t'}$ is a subset of the family of auxiliary games Γ_p^t with respect to the game G. In other words, when considering the linear tracing procedure with respect to game Γ_p^t, we are merely examining a small subset of the family of auxiliary games previously considered.

As was assumed, δ^* is an equilibrium point of G, that is, for each player i,

$$u_i(\delta_i^*, \delta_{-i}^*) \geq u_i(\delta_i, \delta_{-i}^*), \text{ for all } \delta_i \in \Delta_i.$$

Moreover, we assume that δ^* maximizes the expected payoff with respect to p, which means that $u_i(\delta_i^*, p_{-i}) \geq u_i(\delta_i, p_{-i})$ for each player i and each $\delta_i \in \Delta_i$. From these two conditions, it is easy to verify that $u_i^t(\delta_i^*, \delta_{-i}^*) \geq u_i^t(\delta_i, \delta_{-i}^*)$ for each player i and each $\delta_i \in \Delta_i$, which means that δ^* is an equilibrium of game Γ_p^t. Note that $u_i^t(\delta_i, p_{-i}) = u_i(\delta_i, p_{-i})$ for all $\delta_i \in \Delta_i$. Thus, we have that $u_i^t(\delta_i^*, p_{-i}) \geq u_i^t(\delta_i, p_{-i})$ for all $\delta_i \in \Delta_i$. Together, these two coniditions, which specfy the best response conditions for games Γ_p^t and Γ_p^0, guarantee the existence of a feasible path for the equilibrium δ^*. This point can be easily illustrated by the following inequality: for each player i and each $\delta_i \in \Delta_i$

$$\begin{aligned} \mu_i^{t'}(\delta_i^*, \delta_{-i}^*) &= t' u_i^t(\delta_i^*, \delta_{-i}^*) + (1-t') u_i^t(\delta_i^*, p_{-i}) \\ &\geq t' u_i^t(\delta_i, \delta_{-i}^*) + (1-t') u_i^t(\delta_i, p_{-i}) \\ &= \mu_i^{t'}(\delta_i, \delta_{-i}^*) \end{aligned}$$

Since this inequality holds for each player i and each $t' \in [0,1]$, it implies that there exists a feasible path continuously connecting game Λ_p^0 to Γ_p^t. We can therefore conclude that $p \in \Phi^t(\delta^*)$ for each $t \in [0,1]$. ∎

References

[1] D. Fudenberg and J. Tirole, *Game Theory*, The MIT Press, 1991.

[2] J. C. Harsanyi and R. Selten, *A General Theory of Equilibrium Selection in Games*, The MIT Press, 1988.

[3] J. C. Harsanyi, The Tracing Porcedure: a Bayesian approach to defining a solution for n-person noncooperative games, *International Journal of Game Theory* **4**, pp. 61-94, 1975.

[4] P. J. J. Herings, Two Simple Proofs of the Feasibility of the Linear Tracing Procedure, *Economic Theory* **15**, pp. 485-490, 2000.

[5] E. Kohlberg and J. F. Mertens, On the Strategic Stability of Equilibria, *Econometrica* **54**(5), pp. 1003-1037, 1986.

[6] D. M. Kreps and R. Wilson, Sequential Equilibria, *Econometrica* **50**, pp. 863-894, 1982.

[7] S. H. Schanuel, L. K. Simon, and W. R. Zame, The Algebraic Geometry of Games and the Tracing Procedure. In: Selten, R. (ed.) *Game Equilibrium Models II: methods, morals and markets*, pp. 9-43. Berlin Heidelberg New York: Springer 1991.

[8] R. Selten, A reexamination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory* **4**, pp. 25-55, 1975.

[9] P. Walley, *Statistical Reasoning with Impresise Probabilities*, Chapman and Hall, New York, 1991.