Notes On The Theory of Strategic Games[1]

Tilman Börgers, Department of Economics, University of Michigan, July 22, 2013

1. Definitions

Definition 1. *A* strategic game *is a list* $(N, (A_i)_{i\in N}, (u_i)_{i\in N})$ *where* N *is a finite set of the form* $N = \{1, 2, \ldots, n\}$ *(with* $n \geq 2$*), for every* $i \in N$ A_i *is a non-empty, finite set, and for every* $i \in N$ u_i *is a real-valued function with domain* $A = \times_{i\in N} A_i$.

N is the set of players. A_i is the set of player i's actions, and u_i is player i's von Neumann Morgenstern utility function. We write a_i for elements of A_i. We refer to elements of A_i also as *pure* (as opposed to *mixed*) actions of player i. We write a for elements of A. We write A_{-i} for $\times_{j\in N, j\neq i} A_j$. We write a_{-i} for elements of A_{-i}.

Definition 2. *A mixed action* α_i *for player* i *is a probability distribution over* A_i.

We identify the mixed action α_i that places probability 1 on action a_i with this action itself. For any finite set X we write $\Delta(X)$ for the set of probability distributions over X, and thus $\Delta(A_i)$ is the set of mixed actions of player i. We write α for elements of $\times_{i\in N}\Delta(A_i)$. We write α_{-i} for elements of $\times_{j\in N, j\neq i}\Delta(A_j)$. The expected utility of player i when players choose mixed action profile α is:

$$U_i(\alpha) = \sum_{(a_1,a_2,\ldots,a_n)\in A} \left(u_i(a_1, a_2, \ldots, a_n) \prod_{j\in N} \alpha_j(a_j) \right) \tag{1}$$

Note that this definition implicitly assumes that different players' randomizations when implementing their mixed action are independent. We shall make this assumption throughout.

2. Best Replies and Actions that are not Strictly Dominated

Two basic ideas about rationality in games are (i) rational players only choose actions that maximize expected utility for *some* belief about the other players' actions, and (ii) rational players only choose actions that are not strictly dominated by other actions. Here, we formalize these two ideas, and then demonstrate that they imply exactly the same predictions about the behavior of rational players.

[1]Definitions and notation in these notes are very similar, although not completely identical, to the definitions and notation in Osborne and Rubinstein [1].

Definition 3. *A* belief μ_i of player i *is an element of* $\Delta(A_{-i})$. *An action* $a_i \in A_i$ *is a* best reply to a belief μ_i *if*

$$a_i \in \arg\max_{a_i' \in A_i} \sum_{a_{-i} \in A_{-i}} \left(u_i(a_i', a_{-i}) \mu_i(a_{-i}) \right). \tag{2}$$

Note that we do not require beliefs to be the products of their marginals, that is, we don't require player i to believe that the other players' actions are stochastically independent. If we did, Proposition 1 below would not be true.

Definition 4. *An action* $a_i \in A_i$ *is* strictly dominated *by a mixed action* $\alpha_i \in \Delta(A_i)$ *if*

$$U_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i}) \quad \forall a_{-i} \in A_{-i}. \tag{3}$$

Note that we allow the strictly dominating action to be a mixed action. If we did not, then Proposition 1 below would not be true.

Proposition 1. *An action* $a_i^* \in A_i$ *is a best reply to some belief* $\mu_i \in \Delta(A_{-i})$ *if and only if* a_i^* *is not strictly dominated.*

Proof. Step 1: We prove the "only if" part, that is, we assume that a_i^* is a best reply to a belief $\mu_i \in \Delta(A_{-i})$, and infer that a_i^* is not strictly dominated. The proof is indirect. Suppose a_i^* were strictly dominated by $\alpha_i \in \Delta(A_i)$. Then, obviously, α_i yields strictly higher expected utility given the belief μ_i than a_i^*:

$$\sum_{a_{-i} \in A_{-i}} U_i(\alpha_i, a_{-i}) \mu_i(a_{-i}) > \sum_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i}) \mu_i(a_{-i}). \tag{4}$$

We thus have that α_i is a better response to μ_i than a_i^*, which is almost what we want to obtain, but not quite. To obtain the desired contradiction, we want to find a *pure* action that is a better response to μ_i than a_i^*. This can be done as follows. We re-write the left hand side of (4) as follows:

$$\begin{aligned} \sum_{a_{-i} \in A_{-i}} U_i(\alpha_i, a_{-i}) \mu_i(a_{-i}) &= \sum_{a_{-i} \in A_{-i}} \sum_{a_i \in A_i} \alpha_i(a_i) u_i(\alpha_i, a_{-i}) \mu_i(a_{-i}) \\ &= \sum_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \alpha_i(a_i) u_i(\alpha_i, a_{-i}) \mu_i(a_{-i}) \\ &= \sum_{a_i \in A_i} \alpha_i(a_i) \left(\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \mu_i(a_{-i}) \right). \qquad (5) \end{aligned}$$

Combining (4) and (5) we have:

$$\sum_{a_i \in A_i} \alpha_i(a_i) \left(\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \mu_i(a_{-i}) \right) > \sum_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i}) \mu_i(a_{-i}). \tag{6}$$

The left hand side of (6) is a convex combination of the expressions in large brackets in that term. This convex combination can be larger than the right hand side of (6) only if one of the expressions in large brackets is larger than the right hand side of (6), i.e., for some $a_i \in A_i$:

$$(7) \qquad \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})\mu_i(a_{-i}) > \sum_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i})\mu_i(a_{-i}).$$

and thus a_i is a better response to μ_i than a_i^*, which contradicts the assumption a_i^* is a best response to μ_i among all pure actions.

STEP 2:[2] We prove the "if" part, that is, we assume that a_i^* is not strictly dominated, and we show that there is a belief $\mu_i \in \Delta(A_{-i})$ to which a_i^* is a best response. The proof is constructive. We define two subsets, X and Y, of the set $\mathbb{R}^{|A_{-i}|}$, that is, the Euclidean space with dimension equal to the number of elements of A_{-i}. We shall think of the elements of these sets as payoff vectors. Each component indicates a payoff that player i receives when the other players choose some particular $a_{-i} \in A_{-i}$.

Now pick any one-to-one mapping $f : A_{-i} \longrightarrow \{1, 2, \ldots, |A_{-i}|\}$. For any action a_i of player i, we write $u_i(a_i, \langle a_{-i} \rangle) \in \mathbb{R}^{|A_{-i}|}$ for the vector of payoffs that player i receives when playing a_i, and when the other players play their various action combinations a_{-i}. Specifically, the k-th entry of $u_i(a_i, \langle a_{-i} \rangle)$ is the payoff $u_i(a_i, f^{-1}(k))$ where f^{-1} is the inverse of f. Intuitively, f defines an order in which we enumerate the elements of A_{-i}, and $u_i(a_i, \langle a_{-i} \rangle)$ lists the payoffs of player i when he plays a_i and the other players play a_{-i} in the order defined by f.

The set X is:

$$(8) \qquad X = \{x \in \mathbb{R}^{|A_{-i}|} | x > u_i(a_i^*, \langle a_{-i} \rangle)\}.$$

Here, ">" is to be interpreted as: "strictly greater in every component." Therefore, the set X is the set of payoff vectors that are strictly greater in every component than $u_i(a_i^*, \langle a_{-i} \rangle)$, that is the payoff vector that corresponds to the undominated action a_i^*.

The set Y is:

$$(9) \qquad Y = co\{y \in \mathbb{R}^{|A_{-i}|} | \exists a_i \in A_i : y = u_i(a_i, \langle a_{-i} \rangle)\}.$$

Here, "*co*" stands for "convex hull." The payoff vectors in Y are the payoff vectors that player i can achieve by choosing some mixed action. The weight that the convex combination that defines an element of y places on each element of $\{x \in \mathbb{R}^{|A_{-i}|} | \exists a_i \in A_i : x = u_i(a_i, \langle a_{-i} \rangle)\}$ corresponds to the probability which the mixed action assigns to each pure action $a_i \in A_i$.

It is obvious that both sets are nonempty and convex sets. Moreover, their intersection is empty. If X and Y overlapped, then every common element would correspond to the payoffs arising from a mixed action of player i that

[2] An example and a graph that illustrate STEP 2 follow after the end of the proof.

strictly dominates a_i^*. Because by assumption no such mixed action exists, X and Y cannot have any elements in common.

In the previous paragraph we have checked all the assumptions of the separating hyperplane theorem: we have two nonempty and convex sets that have no elements in common. The separating hyperplane theorem (Theorem 1.68 in Sundaram [4]) then says that there exist some row vector $\pi \in \mathbb{R}^{|A_{-i}|}$ which is not equal to zero in every component, and some $\delta \in \mathbb{R}$, such that:

$$\pi \cdot x \geq \delta \quad \forall x \in X \tag{10}$$

and

$$\pi \cdot y \leq \delta \quad \forall y \in Y. \tag{11}$$

Here "$\cdot$" stands for the scalar product of two vectors. We treat all vectors in X and Y as column vectors. Therefore, the above scalar products are well-defined.

We now make two observations. The first is:

$$\delta = \pi \cdot u_i(a_i^*, \langle a_{-i} \rangle). \tag{12}$$

To show this note that by definition $u_i(a_i^*, \langle a_{-i} \rangle) \in Y$, and therefore, by (11), $\pi \cdot u_i(a_i^*, \langle a_{-i} \rangle) \leq \delta$. Next, for every $n \in \mathbb{N}$ define $x_n = u_i(a_i^*, \langle a_{-i} \rangle) + \varepsilon^n \cdot \iota$, where $\varepsilon \in (0,1)$ is some constant and ι is the column vector in $\mathbb{R}^{|A_{-i}|}$ in which all entries are "1". Observe that for every $n \in \mathbb{N}$ we have $x_n \in X$, so that by (10) we have: $\pi \cdot x_n \geq \delta$ for every n. On the other hand, we have: $\lim_{n \to \infty} x_n = u_i(a_i^*, \langle a_{-i} \rangle)$. By the continuity of the scalar product of vectors therefore: $\pi \cdot u_i(a_i^*, \langle a_{-i} \rangle) \geq \delta$. This, combined with our earlier observation $\pi \cdot u_i(a_i^*, \langle a_{-i} \rangle) \leq \delta$ implies (12).

Our second observation is:

$$\pi \geq 0 \tag{13}$$

where we interpret "$\geq$" to mean "greater or equal in every component," but not identical, and 0 stands for the vector consisting of zeros only. That π is not equal to 0 is part of the assertion of the separating hyperplane theorem. We shall prove indirectly that no component of π can be negative. Assume that π has a negative component. Without loss of generality, assume that it is the first. Now we consider the vector $x = u_i(a_i^*, \langle a_{-i} \rangle) + (1, \varepsilon, \varepsilon, \ldots, \varepsilon)$ where $\varepsilon > 0$ is some number. Clearly, the vector x that we define like this is contained in X. Moreover,

$$\pi \cdot x = \pi \cdot u_i(a_i^*, \langle a_{-i} \rangle) + \pi \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon) = \delta + \pi \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon) \tag{14}$$

where for the second equality we have used (12). We now want to evaluate $\pi \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon)$. Unfortunately, we need some additional notation: we shall use for the k -th component of the vector π the symbol: π_k. Here,

$k \in \{1, 2, \ldots, |A_{-i}|\}$. We then obtain:

$$\pi \cdot (1, \varepsilon, \varepsilon, \ldots, \varepsilon) = \pi_1 + \varepsilon \sum_{k=2}^{|A_{-i}|} \pi_k \tag{15}$$

Now observe that by the contrapositive assumption of our indirect proof $\pi_1 < 0$. Therefore, for small enough ε the right hand side of (15) is negative. Using this fact, we obtain from (14):

$$\pi \cdot x < \delta \tag{16}$$

which contradicts $x \in X$ and (10). This completes our indirect proof of (13).

Now we denote by $\|\pi\|$ the Euclidean norm of π. Because $\pi \neq 0$, if we define μ_i by:

$$\mu_i = \frac{1}{\|\pi\|}\pi, \tag{17}$$

then the Euclidean norm of μ_i is 1, and thus $\mu_i \in \Delta(A_{-i})$, that is, μ_i is a belief of player i. We complete the proof by showing that a_i^* is a best response to μ_i. (11) and (12) together imply:

$$\pi \cdot y \leq \pi \cdot u_i(a_i^*, \langle a_{-i} \rangle) \quad \forall y \in Y. \tag{18}$$

Dividing this inequality by $\|\pi\|$ we get:

$$\mu_i \cdot y \leq \mu_i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \quad \forall y \in Y. \tag{19}$$

Now by definition of Y for every a_i: $u_i(a_i, \langle a_{-i} \rangle) \in Y$. Therefore:

$$\mu_i \cdot u_i(a_i, \langle a_{-i} \rangle) \leq \mu_i \cdot u_i(a_i^*, \langle a_{-i} \rangle) \quad \forall a_i \in A_i. \tag{20}$$

This is what we wanted to prove: a_i^* is a best response in A_i to μ_i. $\square$

We illustrate Step 2 of the proof of Proposition 1 with the following game in which player 1 chooses rows and player 2 chooses columns. Only the payoffs of player 1 are shown.

	L	R
T	2	2
MA	3	-2
MB	0	-1
MC	-1	0
B	-2	3

Example 1

Player 1's action T is not strictly dominated. We illustrate in Figure 1 the construction of beliefs to which action T is a best reply. The sets X and Y to which STEP 2 of the proof of Proposition 1 refers are shown in the figure.

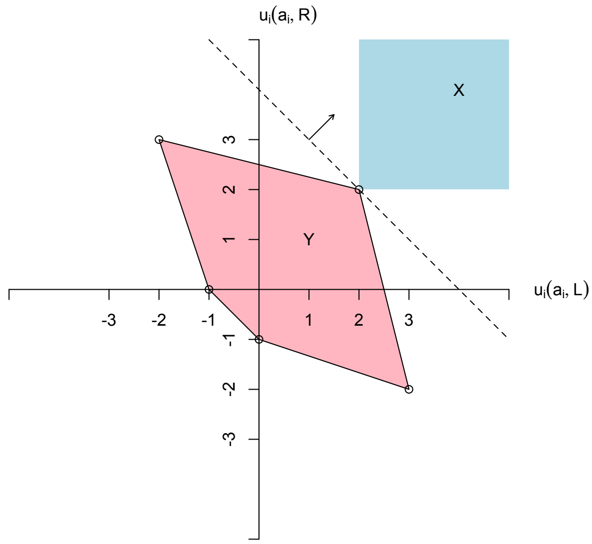

Figure 1

The hyperplane (straight line) separating X and Y is the dashed line in Figure 1. Figure 1 also shows the orthogonal vector for this hyperplane.

3. Cautious Best Replies and Actions that are not Weakly Dominated

A cautious player might hold a belief that attaches some strictly positive, although possibly arbitrarily small, probability to every action profile of the other players. Here we show that the set of best responses to such beliefs is exactly the set of actions that are not *weakly* dominated.

Definition 5. *An action $a_i \in A_i$ is* weakly dominated *by a mixed action $\alpha_i \in \Delta(A_i)$ if*

(21) $$U_i(\alpha_i, a_{-i}) \geq u_i(a_i, a_{-i}) \quad \forall a_{-i} \in A_{-i}$$

and

(22) $$U_i(\alpha_i, a_{-i}) > u_i(a_i, a_{-i}) \quad \exists a_{-i} \in A_{-i}.$$

Proposition 2. *In a finite strategic game an action a_i^* is a best reply among all pure actions of player i to a belief $\mu_i \in \Delta(A_{-i})$ with full support, i.e. $\mu_i(a_{-i}) > 0$ for every $a_{-i} \in A_{-i}$, if and only if a_i^* is not weakly dominated.*

Before proving Proposition 2 we shall illustrate by means of an example that the construction of STEP 2 of the proof of Proposition 1 does not necessarily yield a belief μ_i with full support and therefore can not be used to prove Proposition 1. Indeed, that construction can be done even if an action is weakly dominated. In this case the construction yields a belief that assigns zero probability to some of the action combinations of the other players. Our example is the same as example 1 except that we assume that player 1 has an additional action that yields payoffs 3 and 2, depending whether player 2 chooses L or R. This action thus weakly dominates T. We also assume that action B yields payoffs -2 and 2 rather than -2 and 3, as was the case before. This is to avoid that a mixed action strictly dominates T. The sets X and Y from STEP 2 of the proof of Proposition 1 for this modified example are shown in Figure 2.

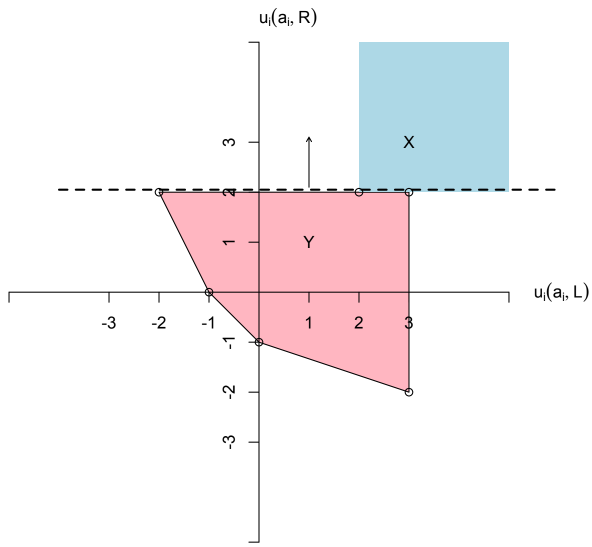

Figure 2

The hyperplane separating X and Y is the dashed line in Figure 2. The orthogonal vector that we have drawn has coordinates 0 and 1, which correspond to the belief of player 1 that puts probability 1 on player 2 playing R. The construction in the proof of Proposition 1 allows thus beliefs that put probability zero on some actions. We demonstrate below how a modified construction shows that full support beliefs can be obtained whenever an action is not weakly dominated.

Proof. The argument that proves that a best response to a full support belief is not weakly dominated is essentially the same as STEP 1 of the proof of Proposition 1, and we omit it.[3]

[3] All that changes is the reasoning that leads to inequality (4), and which we did not even spell out in the proof of Proposition 1.

The proof of the converse is a less obvious modification of the argument in STEP 2 of the proof of Proposition 1.[4] As in that proof, we shall construct a the belief to which a not weakly dominated action is a best reply, and that belief will be the orthogonal vector of a hyperplane separating two sets X and Y. To ensure that this belief has full support, we shall use a stronger separating hyperplane theorem than we used in the proof of Proposition 1. Specifically, it will be useful to appeal to a separating hyperplane theorem that allows us to replace the weak inequality in (10) above by a strict inequality. Such a theorem is the "strong separating hyperplane theorem " which is Theorem 6 in Border [2]. To apply this theorem we need, in particular, that the set X is compact. To achieve this, we have to change the definition of X substantially. This will be done below. This change in the definition of X will then also necessitate a small change in the definition of Y. Hopefully, the remainder of the proof will clarify how these changes in the definitions of X and Y are used in the proof.

We define X and Y as follows:

$$X = \{x \in \mathbb{R}^{|A_{-i}|} | x = u_i(a_i^*, \langle a_{-i} \rangle) + y \text{ for some } y \in \Delta\}. \tag{23}$$

(where Δ is the unit simplex in $\mathbb{R}^{|A_{-i}|}$) and

$$\begin{aligned} Y \;=\; \Big\{ y \in \mathbb{R}^{|A_{-i}|} | \quad & y = u_i(a_i^*, \langle a_{-i} \rangle) \\ & + \sum_{a_i \in A_i} \lambda(a_i) \left(u_i(a_i, \langle a_{-i} \rangle) - u_i(a_i^*, \langle a_{-i} \rangle) \right) \\ & \text{for some function } \lambda : A_i \to \mathbb{R}_+ \Big\}. \end{aligned} \tag{24}$$

Thus, X is now the set of all payoff vectors that result if we add an element of the unit simplex to $u_i(a_i^*, \langle a_{-i} \rangle)$, whereas before it was the set of all payoff vectors that are strictly greater in every component than $u_i(a_i^*, \langle a_{-i} \rangle)$. Note that the modified definition of X leaves X a convex set. Obviously, X is also compact. The set Y is the smallest cone with vertex $u_i(a_i^*, \langle a_{-i} \rangle)$ that includes the set Y as previously defined. Observe that with this new definition Y is convex. By Lemma 6 in Border [3] it is also closed because it is a finitely generated convex cone. It is generated by the finite collection of vectors $u_i(a_i, \langle a_{-i} \rangle) - u_i(a_i^*, \langle a_{-i} \rangle)$ (where $a_i \in A_i$).

The properties of X and Y listed in the previous paragraph are sufficient to allow us to apply the strong separating hyperplane theorem provided that X and Y are disjoint. We prove this indirectly. Suppose $z \in X \cap Y$. Observe that $z \in X$ implies:

$$y \geq u_i(a_i^*, \langle a_{-i} \rangle). \tag{25}$$

[4] An example and a graph that illustrate the modification of STEP 2 follow after the end of the proof.

Because $z \in Y$ there is therefore some function $\lambda : A_i \to \mathbb{R}_+$ such that:

$$u_i(a_i^*, \langle a_{-i}\rangle) + \sum_{a_i \in A_i} \lambda(a_i)\left(u_i(a_i, \langle a_{-i}\rangle) - u_i(a_i^*, \langle a_{-i}\rangle)\right) \geq u_i(a_i^*, \langle a_{-i}\rangle) \tag{26}$$

which is equivalent to:

$$\sum_{a_i \in A_i} \lambda(a_i) u_i(a_i, \langle a_{-i}\rangle) \geq \sum_{a_i \in A_i} \lambda(a_i) u_i(a_i^*, \langle a_{-i}\rangle). \tag{27}$$

Because the two sides of (27) are not equal, we have to have $\lambda(a_i) \neq 0$ for at least some $a_i \in A_i$. We can therefore divide both sides of (27) by $\sum_{a_i \in A_i} \lambda(a_i)$ and obtain:

$$\sum_{a_i \in A_i} \frac{\lambda(a_i)}{\sum\limits_{a_i' \in A_i} \lambda(a_i')} u_i(a_i, \langle a_{-i}\rangle) \geq u_i(a_i^*, \langle a_{-i}\rangle) \tag{28}$$

Inequality (28) says that a_i^* is weakly dominated by the mixed action which gives each action in A_i the probability $\lambda(a_i)/\sum_{a_i' \in A_i} \lambda(a_i')$. This contradicts our assumption that a_i^* is not weakly dominated.

We can now conclude using the strong separating hyperplane theorem (Theorem 6 in Border [2] that X and Y can be strongly separated, that is, there are some $\pi \in \mathbb{R}^{|A_{-i}|}$ which is not equal to zero in every component, and some $\delta \in \mathbb{R}$, such that:

$$\pi \cdot x > \delta \quad \forall x \in X \tag{29}$$

and

$$\pi \cdot y \leq \delta \quad \forall y \in Y. \tag{30}$$

We have two observations about π and δ. The first is:

$$\delta \geq \pi \cdot u_i(a_i^*, \langle a_{-i}\rangle). \tag{31}$$

This follows from (30) and $u_i(a_i^*, \langle a_{-i}\rangle) \in Y$. The second observation is:

$$\pi \gg 0. \tag{32}$$

where "$\gg$" means "strictly greater in every component." Suppose (32) were not true, that is, one component of π, say the k-th component, were less than or equal to zero. Denote by ι_k the unit vector in Δ that equals 1 in the k-th component and zero in all other components. Thus $u_i(a_i^*, \langle a_{-i}\rangle) + \iota_k \in X$. Moreover, by our assumption about π, we would have:

$$\pi \cdot (u_i(a_i^*, \langle a_{-i}\rangle) + \iota_k) \leq \pi \cdot u_i(a_i^*, \langle a_{-i}\rangle) \leq \delta \tag{33}$$

where the last inequality is (31). But (33) contradicts (29).

Inequality (32) allows us to normalize π so that it becomes a probability vector. We denote this probability vector by μ_i.

$$\mu_i = \frac{\pi}{\|\pi\|} \tag{34}$$

Note that $\mu_i \gg 0$, and therefore we can complete the proof by showing that a_i^* maximizes expected utility in A_i when beliefs are μ_i.

Obviously, inequality (30) holds if we replace π by μ_i and divide the right hand side by $\|\pi\|$.

$$\mu_i \cdot y \leq \frac{\delta}{\|\pi\|} \quad \forall y \in Y. \tag{35}$$

Now suppose that a_i^* did not maximize expected utility in A_i:

$$\mu_i \cdot (u_i(a_i, \langle a_{-i}\rangle) - u_i(a_i^*, \langle a_{-i}\rangle)) > 0 \tag{36}$$

for some $a_i \in A_i$. We can find some large enough $\hat{\lambda} > 0$ such that:

$$\mu_i \cdot \hat{\lambda}\,(u_i(a_i, \langle a_{-i}\rangle) - u_i(a_i^*, \langle a_{-i}\rangle)) > \mu_i \cdot u_i(a_i^*, \langle a_{-i}\rangle) + \frac{\delta}{\|\pi\|} \tag{37}$$

This is equivalent to:

$$\mu_i \cdot \Big(\hat{\lambda}\,(u_i(a_i, \langle a_{-i}\rangle) - u_i(a_i^*, \langle a_{-i}\rangle)) - u_i(a_i^*, \langle a_{-i}\rangle)\Big) > \frac{\delta}{\|\pi\|} \tag{38}$$

Now note that $\Big(\hat{\lambda}\,(u_i(a_i, \langle a_{-i}\rangle) - u_i(a_i^*, \langle a_{-i}\rangle)) - u_i(a_i^*, \langle a_{-i}\rangle)\Big) \in Y$ (set $\lambda(a_i) = \hat{\lambda}$ and $\lambda(a_i') = 0$ for all $a_i' \neq a_i$). Thus, we have found an element y of Y such that $\mu_i \cdot y > \delta/\|\pi\|$, which contradicts (35). □

We illustrate Step 2 of the proof of Proposition 2 with the same game that we used above to illustrate the proof of Proposition 1. In that game, player 1's action A is not only not strictly dominated, but also not weakly dominated. We illustrate in Figure 3 the construction of the sets X and Y to which the proof of Proposition 2 refers.

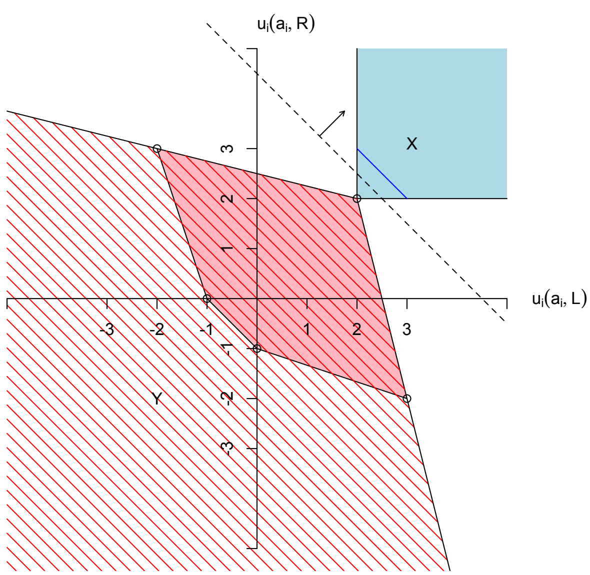

Figure 3

The set Y is now the cone generated by the set that we had called Y in Figure 1. In Figure 3 this cone is indicated by parallel lines from the top left to the bottom right. The set X is the unbroken line in the bottom left corner of the shaded rectangle that represented X in Figure 1. The hyperplane separating X and Y is the dashed line in Figure 3. Figure 3 also shows the orthogonal vector for this hyperplane.

4. Iterated Elimination of Strictly Dominated Actions

If players know each others' utility functions, and know that other players are rational, and know that other players know that players are rational, etc.,[5] it seems plausible to iterate the elimination of strictly dominated actions.

Definition 6. *A finite sequence* $((X_i^t)_{i \in N})_{t=0}^T$ *(where* $T \in \mathbb{N}$*) is a* process of iterated elimination of strictly dominated actions *(IESDA) if for all* $i \in N$*:*

(i) $X_i^0 = A_i$

and for all $i \in N$ *and* $t = 0, 1, \ldots, T-1$*:*

(ii) $X_i^{t+1} \subseteq X_i^t$

(iii) $a_i \in X_i^t \setminus X_i^{t+1}$ *only if* a_i *is strictly dominated in the strategic game with player set* N*, action sets* X_j^t *for every* $j \in N$*, and utility functions equal the restrictions of* u_k *to* $\times_{j \in N} X_j^t$*.*

(iv) *if* $a_i \in X_i^T$ *then* a_i *is not strictly dominated in the strategic game with player set* N*, action sets* X_j^T *for every* $j \in N$*, and utility functions equal to the restrictions of* u_k *to* $\times_{j \in N} X_j^T$*.*

Note that there are multiple processes of IESDA. However, as we show below, they all terminate with the same sets of actions, and therefore the phrase "actions that survive IESDA" is unambiguously defined. We first need a preliminary lemma.

Lemma 1. *Let* $((X_j^t)_{j \in N})_{t=0}^T$ *be a process of IESDA. Suppose* $a_i \in X_i^T$ *is a best reply to some belief* $\mu_i \in \Delta(X_{-i}^T)$ *among all actions in* X_i^T*. Then* a_i *is also a best response to* μ_i *among all actions in* A_i*.*

Proof. Let $\tilde{a}_i$ be any best response to $\mu_i \in \Delta(X_{-i}^T)$ among all actions in A_i. Then $\tilde{a}_i$ can not be eliminated by IESDA in any step t, because of the equivalence of not being strictly dominated and being a best response to some belief, because μ_i having support in $\Delta(X_{-i}^T)$ implies that μ_i has support in $\Delta(X_{-i}^t)$, and because $\tilde{a}_i$ is a best response to μ_i among all actions in X_i^t. Therefore, $\tilde{a}_i \in X_i^T$. If a_i is a best response to μ_i in X_i^T, it must therefore yield at least as high expected utility, and hence the same expected utility, as $\tilde{a}_i$. Therefore, it is also a best response to μ_i among all actions in A_i. $\square$

[5] I.e. utility functions and rationality are common knowledge.

Proposition 3. *Let* $((X_i^t)_{j\in N})_{t=0}^T$ *and* $((\widehat{X}_i^t)_{j\in N})_{t=0}^{\widehat{T}}$ *both be processes of IESDA. Then* $X_i^T = \widehat{X}_i^{\widehat{T}}$ *for all* $i \in N$.

Proof. It suffices to prove:

$$\widehat{X}_i^{\widehat{T}} \subseteq X_i^t \quad \forall t = 0, 1, \ldots, T. \tag{39}$$

This is sufficient because it implies $\widehat{X}_i^{\widehat{T}} \subseteq X_i^T$, and, by the symmetric argument, we can also infer $X_i^T \subseteq \widehat{X}_i^{\widehat{T}}$, so that the assertion follows. We prove (39) by induction over t. The assertion is trivially true for $t = 0$. For the induction step we have to prove that $\widehat{X}_i^{\widehat{T}} \subseteq X_i^t$ for all $i \in N$ implies $\widehat{X}_i^{\widehat{T}} \subseteq X_i^{t+1}$ for all $i \in N$. Consider some $i \in N$ and some $a_i \in \widehat{X}_i^{\widehat{T}}$. We have to show $a_i \in X_i^{t+1}$. Because a_i is not strictly dominated in the game left over at the end of the process $((\widehat{X}_i^t)_{i\in N})_{t=0}^{\widehat{T}}$, it is a best response among all actions in A_i to a belief $\mu_i \in \Delta(\widehat{X}_{-i}^{\widehat{T}})$ (using the above Lemma). By the inductive assumption, μ_i's support is included in X_{-i}^t. Because a_i is a best response to μ_i among all actions in A_i, it is also a best response to μ_i among the actions in in X_i^t. Therefore, it cannot be eliminated in step t of the process $((X_i^t)_{i\in N})_{t=0}^T$, and therefore $a_i \in X_i^{t+1}$. □

Adapting Definition 6 appropriately, one might also define the iterated elimination of weakly dominated actions. The analog of Proposition 3 is, however, not correct when iterated elimination of weakly dominated actions is considered. For this and other reasons iterated elimination of weakly dominated actions is a problematic solution concept that we do not consider any further in these notes.

5. Nash Equilibrium

We now come to the most classic solution concept for strategic games.

Definition 7. *A list of mixed actions* $\alpha \in \times_{i\in N}\Delta(A_i)$ *is a* Nash equilibrium *if for every player* $i \in N$*:*

$$\alpha_i \in \arg\max_{\alpha_i' \in \Delta(a_i)} U_i(\alpha_i', \alpha_{-i}) \tag{40}$$

It is useful to re-write the definition of Nash equilibrium in the following way. For every player i and for every $\alpha_{-i} \in \times_{j\in N, j\neq i}\Delta(a_j)$ define the set of best replies of player i:

$$B_i(\alpha_{-i}) = \arg\max_{\alpha_i' \in \Delta(A_i)} U_i(\alpha_i', \alpha_{-i}). \tag{41}$$

For every $\alpha \in \times_{j\in N}\Delta(A_j)$ define:

$$B(\alpha) = \times_{i\in N} B_i(\alpha_{-i}). \tag{42}$$

A Nash equilibrium is then a fixed point of the correspondence B, that is, an α such that $\alpha \in B(\alpha)$.

The following result gives a characterization of best responses that is useful when determining Nash equilibria.

Lemma 2. *For every player i and every α_{-i} we have $\alpha_i \in B_i(\alpha_{-i})$ if and only if $\alpha_i(a_i) > 0$ implies $a_i \in \arg\max_{a'_i \in A_i} U_i(a'_i, \alpha_{-i})$.*

Proof. The same argument that leads to (5) in STEP 1 of the proof of Proposition 1 shows that for every $\alpha_i \in \Delta(A_i)$:

$$U_i(\alpha_i, \alpha_{-i}) = \sum_{a_i \in A_i} \alpha_i(a_i) \left(\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \prod_{j \in N; j \neq i} \alpha_j(a_j) \right). \tag{43}$$

This implies:

$$U_i(\alpha) \leq \max_{a_i \in A_i : \alpha_i(a_i) > 0} U_i(a_i, \alpha_{-i}). \tag{44}$$

which in turn implies:

$$U_i(\alpha) \leq \max_{a_i \in A_i} U_i(a_i, \alpha_{-i}). \tag{45}$$

Therefore:

$$\max_{\alpha_i \in \Delta_i(A_i)} U_i(\alpha_i, \alpha_{-i}) \leq \max_{a_i \in A_i} U_i(a_i, \alpha_{-i}). \tag{46}$$

On the other hand, because $a_i \in \Delta(A_i)$ for all $a_i \in A_i$:

$$\max_{\alpha_i \in \Delta(a_i)} U_i(\alpha_i, \alpha_{-i}) \geq \max_{a_i \in A_i} U_i(a_i, \alpha_{-i}). \tag{47}$$

Therefore,

$$\max_{\alpha_i \in \Delta(a_i)} U_i(a_i, \alpha_{-i}) = \max_{a_i \in A_i} U_i(a_i, \alpha_{-i}). \tag{48}$$

which means that α_i is a best response to α_{-i} if and only if

$$U_i(\alpha) = \max_{a_i \in A_i} U_i(a_i, \alpha_{-i}). \tag{49}$$

The proof is completed by observing that (43) shows that whenever $a_i \in \arg\max_{a'_i \in A_i} U_i(a'_i, \alpha_{-i})$ for all $\alpha_i(a_i) > 0$ equation (49) holds. By contrast, if there is some $a_i \notin \arg\max_{a'_i \in A_i} U_i(a'_i, \alpha_{-i})$ with $\alpha_i(a_i) > 0$ then $U_i(\alpha_i, \alpha_{-i}) < \max_{a_i \in A_i; \alpha_i(a_i) > 0} U_i(a_i, \alpha_{-i})$. □

We now prove a preliminary result that then allows us to prove the existence theorem for Nash equilibria.

Lemma 3. *For every player i and every α_{-i} the set $B_i(\alpha_{-i})$ is non-empty and convex.*

Proof. $B_i(\alpha_{-i})$ is non-empty because among the finitely many elements of A_i at least one must be a best response among all elements of A_i to α_{-i} in A_i, and by Lemma 2 this action is contained in $B_i(\alpha_{-i})$. To see that $B_i(\alpha_{-i})$ is convex, assume $\alpha_i, \alpha_i' \in B_i(\alpha_{-i})$ and $0 < \lambda < 1$. We have to show: $\alpha_i^\lambda = \lambda\alpha_i + (1-\lambda)\alpha_i' \in B_i(\alpha_{-i})$. By Lemma 2 this means that we have to show $\alpha_i^\lambda(a_i) > 0$ implies that $a_i \in \arg\max_{a_i' \in A_i} U_i(a_i', \alpha_{-i})$. But $\alpha_i^\lambda(a_i) > 0$ implies that either $\alpha_i(a_i) > 0$, or $\alpha_i'(a_i) > 0$, or both. Because α_i and α_{-i} are best responses to α_{-i}, we can apply Lemma 2 to at least one of these mixed actions, and obtain directly what we have to show. □

Proposition 4. *Every strategic game has at least one Nash equilibrium.*

Proof. It suffices to show that B satisfies the conditions of Kakutani's fixed point theorem (Lemma 20.1 in Osborne and Rubinstein [1]), which guarantees the existence of at least one fixed point. For this we have to verify that the domain of B is compact and convex, that $B(\alpha)$ is non-empty and convex for every α, and that B has a closed graph. That the conditions for the domain are satisfied is obvious. Non-emptiness and convexity of $B(\alpha)$ for every α follow from Lemma 3. That B has a closed graph follows from the maximum theorem (Theorem 9.14 in Sundaram [4]). □

Finally, we relate Nash equilibrium actions to actions that survive IESDA.

Proposition 5. *Suppose that α is a Nash equilibrium, and that for some $i \in N$ and $a_i \in A_i$ we have: $\alpha_i(a_i) > 0$. Then a_i survives IESDA.*

Proof. Let $((X_i^t)_{j\in N})_{t=0}^T$ be a process of IESDA. We shall show that $a_i \in X_i^t$ for all $t = 0, 1, \ldots, T$. We prove this by induction over t. The assertion is obvious for $t = 0$. Now suppose we had shown the assertion for t. We want to prove it for $t+1$. By Proposition 1 it is sufficient to show that a_i is a best response to some belief μ_i over $\times_{j\in N, j\neq i} X_j^t$. Define this belief by:

$$\mu_i(a_{-i}) = \prod_{j\in N, n\neq i} \alpha_j(a_j). \tag{50}$$

This is a probability measure over $\times_{j\in N, j\neq i} X_j^t$ because, by the inductive assumption, $\alpha_j(a_j) > 0$ implies $a_j \in X_j^t$. Moreover, a_i is a best response because by assumption α is a Nash equilibrium, and therefore, by Lemma 2 every pure action that is contained in the support of α_i is a best response to α_{-i}. □

References

[1] Martin J. Osborne and Ariel Rubinstein, *A Course in Game Theory*, Cambridge and London: MIT Press, 1994.

[2] Kim C. Border, *Separating Hyperplane Theorems*, course notes available at: http://www.hss.caltech.edu/~kcb/Notes/SeparatingHyperplane.pdf.

[3] Kim C. Border, *Alternative Linear Inequalities*, course notes available at: http://www.hss.caltech.edu/~kcb/Notes/Alternative.pdf

[4] Rangarajan K. Sundaram, *A First Course in Optimization Theory*, Cambridge et. al.: Cambridge University Press, 1996.