Game Theory

Lecture 3: Relations between strategies - dominance and best replies

Frieder Neunhoeffer

Lisbon School of Economics & Management
Universidade de Lisboa

ULISBOA | UNIVERSIDADE DE LISBOA

Example: Prisoner's dilemma

- The payoff (years in prison) matrices of this game are:

$$A = \begin{pmatrix} 3 & 0 \\ 5 & 1 \end{pmatrix}, B = \begin{pmatrix} 3 & 5 \\ 0 & 1 \end{pmatrix}, \qquad (A, B) = \begin{pmatrix} 3,3 & 0,5 \\ 5,0 & 1,1 \end{pmatrix}$$

- where A,B stand for the payoff matrices of the *Row* and *Column* players, respectively. For each player, the first pure strategy ("Confess") strictly dominates the second one ("Deny"). Since there are only two pure strategies, the first pure strategy strictly dominates all the mixed strategies different from it. These mixed strategies are convex combinations of the two pure strategies.
- **Remark:** It is possible that a pure strategy is dominated by a mixed strategy without being dominated by any pure strategy.

Dominance: weak, strict and iterated strict

- Weak dominance: Strategy x_i weakly dominates strategy y_i (both are strategies available to player i) if $u_i(x_i, z_{-i}) \geq u_i(y_i, z_{-i})$ for any profile of strategies selected by the other players (i.e., for any $z \in \Theta$), where a strict inequality holds for at least one profile $z \in \Theta$.

- Non-dominated: Strategy x_i is non-dominated if there is no strategy that weakly dominates it.

- Strict dominance: Strategy x_i strictly dominates strategy y_i if $u_i(x_i, z_{-i}) > u_i(y_i, z_{-i})$ for any profile of strategies chosen by the other players (i.e., for any $z \in \Theta$).

Dominance relations

Dominance relations introduce a *partial order* between strategies. It does not matter if we consider the set of mixed strategies of a player i (i.e., the whole unit simplex Δ_i) or, only its vertices (i.e., the set of pure strategies S_i): the relevant definitions and propositions remain the same.

"Partial order" means that dominance allows us to rank any two strategies *sometimes* but not *always*.

Another example

- Let us consider a game with two players where the payoff matrix of player 1 (whose strategies are represented by the matrix rows) is given by,

$$A = \begin{pmatrix} 3 & 0 \\ 0 & 3 \\ 1 & 1 \end{pmatrix}$$

- Player 1 has three pure strategies, whereas player 2 (whose strategies are represented by the columns of the matrix) has two pure strategies. The third pure strategy of player 1, $x_1 = e_1^3$, is not weakly dominated by any of the other two pure strategies. However, it can yield a lower payoff than a mixed strategy for any strategy choice made by the opponent.
- Formally, let the mixed strategy be $y_1 = \left(\frac{1}{2}, \frac{1}{2}, 0\right)$. Then, we have the relation

  $1 = u_1(x_1, z_2) < u_1(y_1, z_2) = \frac{3}{2}$ for any strategy (pure or mixed) selected by player 2.

Why is identifying dominated strategies important from the viewpoint of finding a solution (or "equilibrium") of the game?

- It is clear that a "rational" player does not use a **strictly dominated** strategy.
- Moreover, some authors (such as Kohlberg and Mertens) argue that a "rational" player does not even use a **weakly dominated** strategy.
- Hence, strictly dominated strategies can be eliminated from the game without affecting its outcome.
- Let $S_D \subset S$, be the set of profiles of pure strategies that survive the process of elimination of strictly dominated strategies. If each player is left with a unique pure strategy, i.e., if S_D owns only one strategy profile, then the game is said to be *solvable by dominance*.

Example 1

- The *Prisoner's Dilemma* is a game that is solvable by dominance.
- Elimination is *iterated*. A pure strategy in a game G is not iteratively strictly dominated if:

1. It is not strictly dominated in the original game G.
2. Or it is not strictly dominated in the reduced game G_1 that is obtained from G, through the elimination of some or all the strategies that are strictly dominated in this game.
3. Or it is not strictly dominated in the further reduced game G_2, which is obtained from G_1 through elimination of strictly dominated strategies in G_2.
4. And so on ...
5. Until it becomes impossible to eliminate more strategies, i.e., until $G_{t+1} = G_t$ for some positive integer t.

Example 1

It can be easily demonstrated that:

- The process of elimination of strictly dominated strategies stops after a finite number of steps.
- The final outcome of the process does not depend on the order of dominated strategy elimination.

Example 2

- Another game that is solvable by dominance is given by the payoff matrices A (player 1) and B (player 2).

$$A = \begin{pmatrix} 3 & 1 & 6 \\ 0 & 0 & 4 \\ 1 & 2 & 5 \end{pmatrix}, B = \begin{pmatrix} 3 & 0 & 1 \\ 1 & 0 & 2 \\ 6 & 4 & 5 \end{pmatrix}$$

- Player 1 inspects matrix B and concludes that, for player 2, the second pure strategy is strictly dominated by both pure strategies 1 and 3. Consequently, she assumes that player 2 will never use the pure strategy 2, so that matrix B can be reduced to

$$B^1 = \begin{pmatrix} 3 & 1 \\ 1 & 2 \\ 6 & 5 \end{pmatrix}$$

Example 2

- Likewise, as player 2 never uses strategy 2, player 1 can reduce her matrix A by suppressing the second column, thus becoming,

$$A^1 = \begin{pmatrix} 3 & 6 \\ 0 & 4 \\ 1 & 5 \end{pmatrix}$$

- In the reduced matrix A^1, pure strategies 2 and 3 are strictly dominated by pure strategy 1, which is the only survivor to the elimination process. As player 2 knows that player 1 must use the pure strategy 1, she can further reduce her matrix by suppressing the second and third rows to get

$$B^2 = \begin{pmatrix} 3 & 1 \end{pmatrix}$$

- Under these conditions, player 2 eliminates the third strategy. Pure strategies 1 are for both players the remaining ones after a process of iterated elimination of strictly dominated pure strategies

Iterated dominance vs. one-round dominance

The assumptions on iterated dominance are far stricter than on one-round dominance:

- In one-round dominance, each player is assumed to know her payoff function. Thus, dominated strategies in the player's own payoff function can be deleted.
- In iterated dominance, she has to know the payoff functions of all the players, since each player eliminates pure strategies in the matrix of her opponent.
- In one-round dominance, each player should be rational.
- In iterated dominance, the rationality of each player should be *common knowledge*: each player is rational; each one knows that the other is also rational; each knows that the other knows that she is rational, and so on ...

Best replies

We introduce the difference between the concepts of "function" and "correspondence":

- A *function* associates to each point x a **unique** point $f(x)$.
- A *correspondence* associates to each point x a **non-empty set** $\phi(x)$.

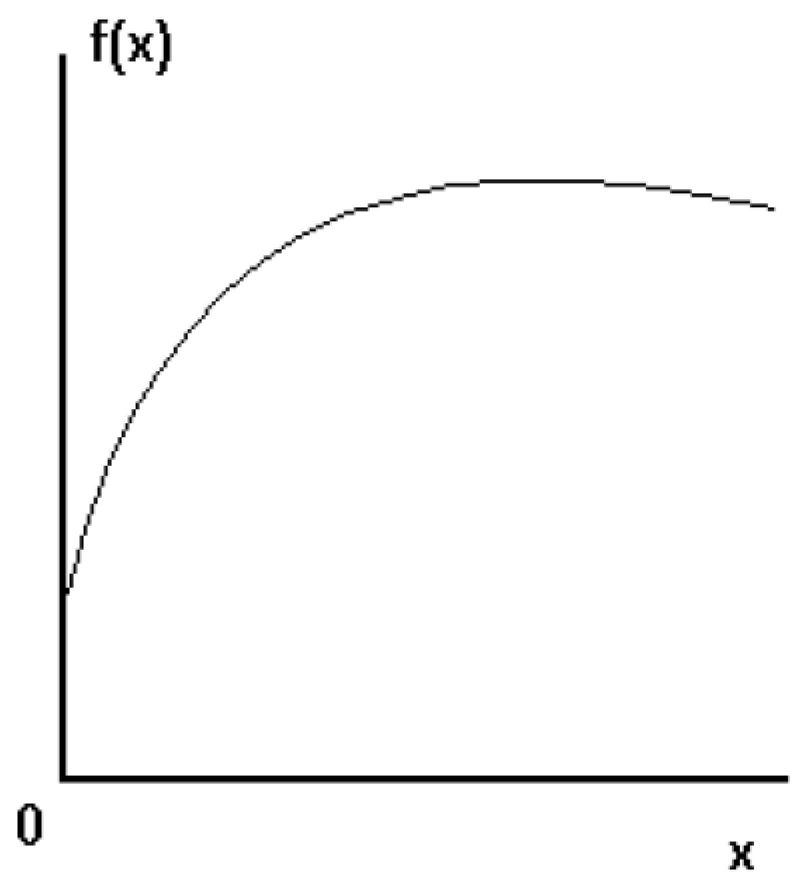

Function

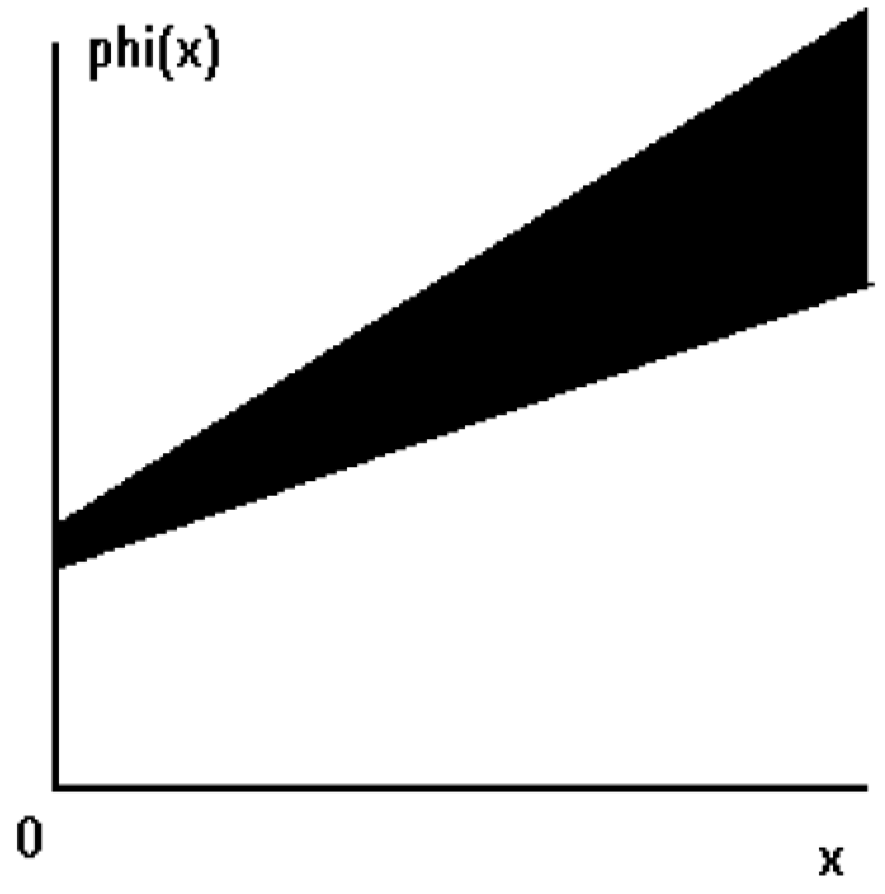

Correspondence

Pure best reply

A **pure** best reply for player i to the profile of strategies $y \in \Theta$ is a *pure strategy* such that no other pure strategy available to the player gives her a higher payoff against the profile y.

This definition leads to a *correspondence of best replies of player i in pure strategies* $\beta_i : \Theta \to S_i$, that maps each profile of mixed strategies $y \in \Theta$ into the non-empty (and finite) set of pure best replies of player i to y, which is formally defined by:

$$\beta_i(y) = \left\{h \in S_i : u_i\left(e_i^h, y_{-i}\right) \geq u_i\left(e_i^k, y_{-i}\right), \forall k \in S_i\right\}$$

It is clear that if two pure strategies h and k of player i are best replies to the profile of strategies $y \in \Theta$, then they lead to the same expected payoff for player i, i.e.,

$$u_i\left(e_i^h, y_{-i}\right) = u_i\left(e_i^k, y_{-i}\right)$$

Example

Let us recall the example of lecture 2. The payoff matrix A of player 1, whose pure strategies are represented by the lines of the matrix, is,

$$A = \begin{pmatrix} 3 & 0 \\ 0 & 3 \\ 1 & 1 \end{pmatrix}$$

In this example, the third pure strategy is not a best reply to no strategy profile.

- For any y_{-i}, the third pure strategy yields a payoff of 1 to player 1.
- If $y_{-i} = \left(\frac{1}{2}, \frac{1}{2}\right)$, the first and second pure strategies yield both $\frac{3}{2}$.
- If $y_{-i} = (1,0)$, the first pure strategy yields 3.
- If $y_{-i} = (0,1)$, the second pure strategy yields 3.

Best reply correspondence in pure strategies

Since

1. any mixed strategy x_i of player i is a convex combination of pure strategies $e_i^h, h = 1, \dots, m_i$,
2. the expected payoff function $u_i(x_i, y_{-i})$ of player i is linear in x_i.

Then we can conclude that no mixed strategy x_i can give player i a higher payoff against the profile of mixed strategies $y \in \Theta$ than any of the pure best replies to y.

Hence in this case, we can define the best reply correspondence in pure strategies *in relation to a mixed strategy*:

$$\beta_i(y) \; = \left\{h \in S_i : u_i(e_i^h, y_{-i}) \geq u_i(x_i, y_{-i}), \forall x_i \in \Delta_i\right\}$$

Best reply correspondence in mixed strategies

Generalizing, we can define a best reply correspondence in *mixed strategies* $\tilde{\beta}_i$. A mixed strategy x_i is a best reply for player i to the profile of strategies $y \in \Theta$ if no other mixed strategy z_i gives player i a higher pay payoff against y than x_i does. Formally,

$$\tilde{\beta}_i : \Theta \to \Delta_i$$
$$\tilde{\beta}_i(y) = \{x_i \in \Delta_i : u_i(x_i, y_{-i}) \geq u_i(z_i, y_{-i}), \forall z_i \in \Delta_i\}$$

Strategy x_i, that is **never a best reply**, is a strategy for which there is no strategy profile $y \in \Theta$ such that x_i is a **best reply** to y.

In practice, x_i is a **best reply** to the profile of strategies y chosen by all players if and only if it is **a best choice** for player i given that she holds a conjecture y_{-i} on the strategies chosen by her rivals.

Interpretation

A mixed strategy x_i is a best reply for player i against the profile of strategies $y \in \Theta$, if each pure strategy that is assigned a positive probability by x_i is also a pure best reply against y.

Since two pure best replies against the same strategy profile y must have an equal expected payoff, the mixed best reply x_i has the property that the pure strategies endowed with positive probabilities are *equalized* in payoff terms.

We can also write the *combined payoff correspondences of best replies*, as the Cartesian product of the respective correspondences of each player i, in pure and mixed strategies:

$$\beta(y) = \times_{i \in I}\, \beta_i(y) \subset S$$
$$\tilde{\beta}(y) = \times_{i \in I}\, \tilde{\beta}_i(y) \subset \Theta$$

Meaning of $\tilde{\beta}(y)$: This correspondence associates with each profile of mixed strategies $y = (y_1, y_2, \dots, y_n)$ for n players a non-empty set of points of the form $x = (x_1, x_2, \dots, x_n)$ where x_i is a mixed best reply to y.

Relations between “dominance” and “best reply”

Working independently, BERNHEIM and PEARCE proposed in 1984 the two following relations between dominance and best replies.

- **Theorem 1**: If a pure strategy is a best reply to some profile of mixed strategies, then it can not be strictly dominated (but it can be weakly dominated).
- **Theorem 2**: If a pure strategy is a best reply to some profile of *completely* mixed strategies, then it is not dominated.

Rationalizability

Hence, an important concept is the set of *rationalizable* strategies: the set of strategies that survive an iterative process of eliminating strategies that are not a best reply to any profile of strategies by the set of players. By definition, these strategies are **neither** strategies that are **never** best replies, **nor** strategies that are a best reply to a strategy that, in turn, is **never** a best reply, and so on ...

By definition, *rationalizable* strategies are the only strategies that can be used by rational players in a game where both players' rationality and the game structure are **common knowledge**.

It can be proved that each player has, at least, one rationalizable strategy and that the set of rationalizable strategies does not depend upon the specific order followed by the elimination process.

If a unique pure strategy for each player survives this kind of elimination process, then the set of *rationalizable strategies* is a solution concept of the game.

Example

To explain better, the meaning of *rationalizability*, we use the finite game drawn from BERNHEIM:

	b_1	b_2	b_3	b_4
a_1	$0,7$	$2,5$	$7,0$	$0,1$
a_2	$5,2$	$3,3$	$5,2$	$0,1$
a_3	$7,0$	$2,5$	$0,7$	$0,1$
a_4	$0,0$	$0,-2$	$0,0$	$10,-1$

Let us find the sets of pure rationalizable strategies for both players:

1. *Round*: We remove b_4, which is never a best reply because it is strictly dominated by a mixed strategy, which selects the pure strategies b_1 and b_3, with the same probability $\frac{1}{2}$.

Example

2. *Round*: Once b_4 is removed, strategy a_4, can be eliminated because it is strictly dominated by a_2.

3. *Round*: Henceforth, no further strategy can be removed:
   - a_1 is a best reply to b_3
   - a_2 is a best reply to b_2
   - a_3 is a best reply to b_1
   - b_1 is a best reply to a_1
   - b_2 is a best reply to a_2
   - b_3 is a best reply to a_3

Justification chain

For each *rationalizable strategy*, a player can build a **justification chain** for her choice which is assumes that no player believes another player selects a strategy that is never a best reply.

For instance, in the above game, Player 1 can justify the choice of a_2 by the belief that Player 2 chooses b_2, which in turn Player 1 can justify for herself believing that Player 2 thinks that Player 1 believes that Player 2 selects b_2, and so on. Hence, Player 1 can build an infinite justification chain $(a_2, b_2, a_2, b_2, \dots)$.

Similarly, Player 1 can rationalize the use of strategy a_1 with the justification chain $(a_1, b_3, a_3, b_1, a_1, b_3, a_3, b_1, a_1, \dots)$. In this context, Player 1:

1. justifies playing a_1, by believing that Player 2 will play b_3.
2. justifies the belief that Player 2 will play b_3, by believing that Player 2 believes that Player 1 will choose a_3.

Justification chain

3. justifies the latter belief (i.e., Player 2 believes that Player 1, will select a_3), by thinking that Player 2 thinks Player 1 believes that Player 2 will play b_1.
4. and so on ...

By contrast, let us assume that Player 1 tries to justify the choice of a_4. The only way to do so is to hold a belief that Player 2 will play b_4. The contradiction lies in that there is **no belief** which Player 2 could possibly hold which would allow to justify b_4. Consequently, Player 1 is not able to justify playing the **non-rationalizable** strategy a_4.

References

- MAS-COLLEL, A. et AL. (1995), Microeconomic Theory, chapter 8, sections 8.B, 8.C, pages: 235-245.
- WEIBULL, Jörgen W. (1997), Evolutionary Game Theory, Cambridge Mass./London, England, chapter 1: 10-14.