Chapter 3: Nash Equilibrium

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

Nash Equilibrium builds upon the notion that every player finds the strategy that maximizes her payoff against each of his rivals' strategies, which we refer as her "best response".

- Discrete Setting:

  If player i has only two available strategies, "maximizing" his payoff means choosing the strategy that yields the highest payoff, taking her rivals' strategies as given.

- Continuous Setting:

  Player i chooses her strategy from a continuum of available strategies, her best response is found by maximizing this player's payoff, literally:

  Solving a utility or profit maximization problem, where we will take the strategies of her rivals as given (as if they were a parameter).

Best Response

Definition: Player i regards strategy s_i as a best response to strategy profile s_{-i}, if

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$$

for every available strategy $s_i' \neq s_i$.

- Intuitively, in a game with two players, i and j, strategy s_i is player i's best response to player j's strategy, s_j, if s_i yields a weakly higher payoff than any other strategy s_i' against s_j.
- In other words, when player j chooses s_j, player i maximizes her payoff by responding with s_i than with any other available strategies.

Best Response

Definition: Player i regards strategy s_i as a best response to strategy profile s_{-i}, if

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$$

for every available strategy $s_i' \neq s_i$.

- If player i finds that the same strategy s_i is a best response against *every strategy profile* of her rivals, all s_{-i}, then s_i must also be strictly dominant.
- Therefore:
  - A strictly dominant strategy is a BR, but…
  - not every BR is a dominant strategy.

Tool 3.1. How to find best responses (BR) in matrix games

1. Focus on the row player by fixing your attention on one strategy of the column player (i.e., one specific column).
   - Cover with your hand (or with a piece of paper) all columns that you are not considering.
   - Find the highest payoff for the row player by comparing the first component of every pair.
   - For future reference, underline this payoff. This is the row player's best response payoff to the column that you considered from the column player.
2. Repeat step 1, but now fix your attention on a different column.
3. For the column player, the method is analogous, but now direct your attention on one strategy of the row player (i.e., one specific row). Cover with your hand all the rows that you are not considering, and compare the payoffs of the column player (i.e., second component of every pair).

Example BR with discrete strategy spaces

		Player 2	
		l	*r*
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

Example BR with discrete strategy spaces

We can start finding Player 1's best responses.

		Player 2	
		l	r
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

$BR_1(l) = U$; $BR_1(r) = D$, or simply,

$$BR_1(s_2) = \begin{cases} U\ if\ s_2 = l, and \\ D\ if\ s_2 = r \end{cases}$$

Example BR with discrete strategy spaces

We can start finding Player 1's best responses.

		Player 2	
		l	*r*
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

		Player 2	
		l	*r*
Player 1	U	$\underline{5}$,3	2,1
	D	3,6	$\underline{4}$,7

Matrix 3.1b. Finding best responses in a 2x2 matrix – Player 1

Example: BR with discrete strategy spaces

We can start finding Player 2's best responses.

		Player 2	
		l	r
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

$BR_2(U) = l\,;\, BR_2(D) = r$, or simply,

$$BR_2(s_1) = \begin{cases} l \;\; if\; s_1 = U, and \\ r\; if\; s_1 = D \end{cases}$$

Example BR with discrete strategy spaces

We can start finding Player 2's best responses.

		Player 2	
		l	*r*
Player 1	U	5,3	2,1
	D	3,6	4,7

Matrix 3.1a. Finding best responses in a 2x2 matrix

		Player 2	
		l	*r*
Player 1	U	$5,\underline{3}$	2,1
	D	3,6	$4,\underline{7}$

Matrix 3.1c. Finding best responses in a 2x2 matrix – Player 2

Finding BRs with Continuous Strategy Spaces

If player i faces a payoff function $u_i(s_i, s_{-i})$,

- We only need to fix her rival's strategies at a generic s_{-i}, and
- Differentiate with respect to s_i to find the value of the strategy s_i that maximizes player i's payoff.

Example 3.1: BR under Cournot Quantity Competition

- Consider an industry with two firms, 1 and 2,
- Every firm simultaneously and independently choosing its output level.
- Every firm i faces a symmetric cost function $C(q_i) = cq_i$,
  - where $c > 0$ denotes the marginal cost of additional units of output,
- Firms face inverse demand function $p(Q) = 1 - Q$,
  - where $Q = q_1 + q_2$ represents aggregate output.

Example 3.1: BR under Cournot Quantity Competition

To find BR for firm 1,

$$\max_{q_1 \geq 0} \; (1 - q_1 - q_2)\, q_1 - c q_1$$

Differentiating with respect to q_1, we find:

$$1 - 2q_1 - q_2 = c$$

Solving for q_1:

$$q_1 = \frac{1-c}{2} - \frac{1}{2} q_2$$

Firm 1's best response function

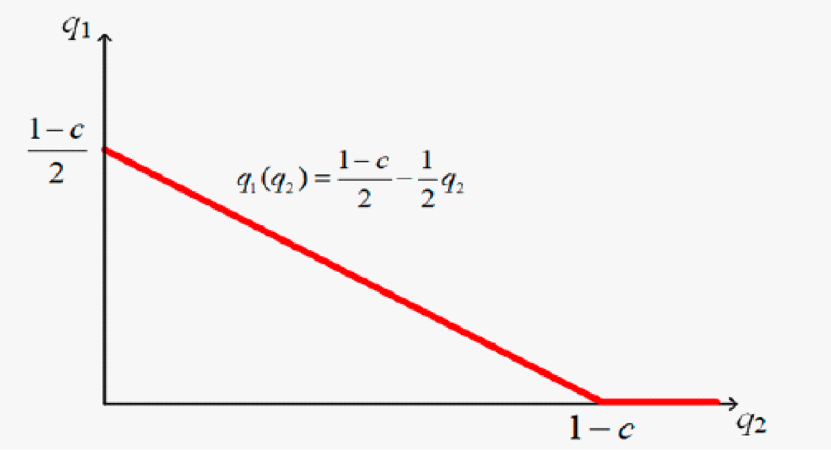

Figure 3.1. Firm 1's best response function.

Firm 1's best response function:

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 \; if \; q_2 < 1 - c \\ 0 \;\; \text{otherwise} \end{cases}$$

Because firm 1 responds with a lower output q_1 when its rival increases its output, q_2, we say that firms' output decisions in this context are "strategic substitutes."

A symmetric best response function applies for firm 2.

Deleting Strategies that are Never a Best Response (NBR)

- Before we use BRs to find the Nash equilibrium of a game…

- We can apply NBR in a similar fashion as in IDSDS.
- Recall that, in IDSDS, we deleted a player’s strictly dominated strategies as never being used.
- Now we use a similar approach:
  - If a player never uses a strategy as a BR, we can label it to be “Never a best response” (NBR), and…
  - Delete it from her strategy set.

Deleting Strategies that are Never a Best Response (NBR)

Definition. Never a Best Response (NBR):
Strategy s_i is never a best response if:

$$u_i(s_i, s_{-i}) \ > (s_i', s_{-i}) \text{ for every } s_i' \ \neq s_i.$$

does not hold for any strategy profile of her rivals, s_{-i}.

Alternatively, if a strategy s_i is NBR, there are

no beliefs that player i can sustain about how her opponents behave

that would lead her to use strategy s_i as a best response.

Deleting Strategies that are Never a Best Response (NBR)

- If player i finds strategy s_i is strictly dominated by s_i', then:
  - s_i yields a strictly lower payoff than s_i'
  - *regardless* of the strategy profile her rivals choose.
- As a consequence, s_i cannot be a BR against any strategy profile of i's rivals.
  - That is, s_i is NBR.
- In summary,

$$s_i \text{ strictly dominated} \Longrightarrow s_i \text{ is NBR.}$$

Rationalizability

- We can go through an iterative process –analogous to IDSDS– but…
  - Identifying strategies that are NBR for either player,
  - rather than strategies that are strictly dominated.
- This iterative process is known as “rationalizability,”
  - Because it finds strategies that player i can rationalize (explain) to be BRs to at least one of her opponents’ strategies.

Tool 3.2. Applying Rationalizability

1. Starting with player i, delete every strategy that is NBR, obtaining the reduced strategy set $S_i' \subset S_i$.
2. Using common knowledge of rationality, we continue the above reasoning, arguing that player $j \neq i$ can anticipate player $i's$ best responses and, as a consequence, the strategies that are NBR for player i, deleting them from S_i.

Given this reduced strategy space $S_i' \subset S_i$, player j can now examine her own best responses to player i, seeking to identify if one or more are never used, and further restricting her strategy space to $S_j' \subset S_j$.

Applying Rationalizability

3. We obtain the Cartesian product $S_i' \times S_j'$, representing the remaining rationalizable strategies after deleting NBRs for two steps. Player i then finds if some of her strategies in this reduced game are NBR, deleting them from her strategy set S_i', obtaining $S_i'' \subset S_i$.

4…

k. The process continues until we cannot find more strategies that are NBR for either player.

Applying Rationalizability

- The strategy profile (or set of strategy profiles) that survives this iterative process are referred as “rationalizable” strategy profiles because every player can sustain beliefs about her rival’s behavior.
- In games with two players:
  - Both solution concepts yield identical equilibrium results, that is, IDSDS=Rationalizability.
- In games with three or more players:
  - Equilibrium outcomes do not necessarily coincide.
  - Rationalizability produces more precise equilibrium outcomes than IDSDS.
  - [Technical, see Fudenberg and Tirole, 1995]
- That is, for a given strategy profile s,

$$s \text{ is rationalizable} \Longrightarrow s \text{ survives IDSDS}$$
$$\not\Leftarrow$$

Example 3.2: Rationalizability and IDSDS

		Firm 2	
		h	l
	H	4,4	0,2
Firm 1	M	1,4	2,0
	L	0,2	0,0

Matrix 3.2a. Applying Rationalizability – First Step

Starting from firm 1, row L is NBR:

- When firm 2 chooses h, firm 1's BR is H
- When firm 2 chooses l, firm 1's BR is M

In other words, firm 1 does not have incentives to respond with L regardless of the believes it sustains on firm 2's behavior (i.e., the beliefs on firm 2 chooses column h or l).

Example 3.2: Rationalizability and IDSDS

		Firm 2	
		h	l
Firm 1	H	4,4	0,2
	M	1,4	2,0

Matrix 3.2b. Applying Rationalizability – Second Step

Moving to firm 2, we find that:

- When firm 1 chooses H, firm 2’s BR is h.
- When firm 1 chooses L, firm 2’s BR is h.

Therefore, column l is NBR for firm 2.

Example 3.2: Rationalizability and IDSDS

		Firm 2
		h
Firm 1	H	4,4
	M	1,4

Matrix 3.2c. Applying Rationalizability –Third Step

Moving back to firm 1, we find that M is its BR to firm 2 choosing l.

Therefore, M is NBR in this step of applying rationalizability.

	h
H	4,4

In conclusion, (H,h) survives rationalizability.

As expected, this coincides with the unique strategy profile surviving IDSDS.

Evaluating Rationalizability as a solution concept

Spoiler: It exhibits the same properties as IDSDS.

1. Existence? Yes
   - Rationalizability satisfies existence, meaning that at least one strategy profile in every game must be rationalizable.
2. Uniqueness? No.
   - One or more strategy profiles may survive rationalizability
   - Does not yield a unique equilibrium prediction in all games with three or more players
3. Robust to small payoff perturbations? Yes.
   - Rationalizability is robust to small perturbations because rationalizability does not change equilibrium outcomes if we alter the payoff of one of the players by a small amount.
4. Socially Optimal? No.
   - Rationalizability does not necessarily yield social optimal outcomes, as illustrated by the Prisoner's Dilemma game, where the only strategy profile surviving rationalizability is (Confess, Confess), which is not socially optimal.

Evaluating Rationalizability as a solution concept

	IDSDS	IDWDS	SDE	Rationalizability
Existence	Yes	Yes	No	Yes
Uniqueness	No	No	Yes	No
Robustness of payoff changes	Yes	Yes	Yes	Yes
Pareto optimal	No	No	No	No

Application of Rationalizability: Finding NBRs in the Beauty Contest

- Consider that your instructor of the Game Theory course shows up in the classroom proposing to play the following “Guess the Average” game (also known as “Beauty Contest”).
- She distributes small pieces of paper and asks each student to write a number (an integer) between 0 and 100.
- There are $N > 2$ students and each of them simultaneously and independently writes her number, taking into account the following rules of the game:
  - The instructor will collect all numbers, write them on the board,
  - Compute the average, and then multiply that average by $\frac{1}{2}$ (half of the average).
  - The instructor, then, declares a winner the student who submitted a number closest to half the average.

Finding NBRs in the Beauty Contest Contd.

Which strategies survive rationalizability?

- In the first step, we can try to eliminate strategies that are NBR for any student.
- Starting with the highest number, $s_i = 100$, we can see that, even if all other students submit $s_j = 100$ for all $j \neq i$, and the number of students is large enough to yield an average of $\bar{s} = 100$, we would obtain that half of the average approaches 50.
- Therefore, student $i's$ best response would be 50, i.e., $BR_i(s_j) = 50$ where $s_j = 100$ for all $j \neq i$, entailing that i cannot rationalize $s_i = 100$ being a best response, regardless of which numbers she believes her rivals submit.
- In summary, $s_i = 100$ is NBR, helping us restrict her strategy space to $S_i' = \{0,99\}$, in the first step of rationalizability.

Finding NBRs in the Beauty Contest Contd.

Which strategies survive rationalizability?

- In the second step, even if all other students submit the highest number (in the reduced strategy space S_i'), $s_j = 99$ for all $j \neq i$, the average would be $\bar{s} = 99$, entailing that half of the average is 49.5. As a result, every student $i$finds that submitting $s_i = 99$ is NBR either, so we can further restrict her strategy space to $S_i'' = \{0{,}98\}$ in the second step of rationalizability.
- …
- After 97 more rounds, we are left with a restricted strategy space of $\{0{,}1\}$ which cannot be further reduced. Indeed, submitting $s_i = 0$ $(s_i = 1)$ is a best response if student i believes that her classmates will all submit $s_j = 0$ $\left(s_j = 1\right)$, so we cannot rule out strategies as being NBR. Therefore, all symmetric strategy profiles, where every student i submits the same number (all submit $s_i = 0$ or all submit $s_i = 1$) survive rationalizability; and so do all asymmetric strategy profiles, where at least on student submits $s_i = 0$ and all other $j \neq i$ students submit $s_j = 1$.

Experimental Tests

- This game has been tested in experimental labs quite often.
  - See, for instance, Nagel (1995) and Stahl and Wilson (1995), which initiated these studies, and for a literature review, see Camerer (2003) and Crawford et al. (2011).
- If you participate in one of these experiments and believe that all other participants randomly submit a number uniformly distributed between 0 and 100, often referred by the literature as “Level-0” players, then you can participate that the average will be around 50 (assuming a sufficiently large number of participants), entailing that half of the average would be 25.
- Therefore, if you think that your rivals are “Level-0” players, you submit $s_i = 25$.
- However, if you think that your rivals must have gone through that same thought process, and they are submitting $s_j = 25$, so they are “Level-1” players…
  - You could outsmart them anticipating that the average will be 25, and half of it will be 12.5, inducing you to submit $s_i = 12.5$.

Experimental Tests

- If you think that your rivals went through that thought process too, and are submitting $s_j = 12.5$, thus being "Level-2" players.
  - Then half the average will be 6.25, implying that, for you to outsmart them, you must submit $s_i = 3.125$.
- The argument extends to further steps, so that if a player thinks that her rivals are "Level-k" players, she must submit

$$s_i = \frac{50}{2^{k+1}}.$$

- Interestingly, this argument implies that, as a participant, it is not necessarily optimal for you to submit some of the strategies that survive rationalizability $s_i = 0 \text{ or } s_i = 1$, as shown above).
  - Instead, you submit a different number depending on the type of Level- k rivals that you deal with, submitting lower numbers as your rivals become more sophisticated (higher k, so they can go through more iterations).

Experimental Tests

- Nagel (1995), for instance, shows that:
  - Many subjects submit $s_i = 25$, suggesting they are Level-1 players; but…
  - Undergraduate students who took Game Theory courses tend to submit $s_i = 12.5$ or $s_i = 6.25$ (thus being Level-2 or -3 players),
    - and so do Caltech students in the Economics major (who generally have a strong Math background)
    - and usual readers of financial newspapers (*Financial Times* in the UK and *Expansion* in Spain).

Finding NBRs in the Cournot Duopoly

From earlier

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 & if\ q_2 < 1-c \\ 0 & \text{otherwise} \end{cases}$$

and similarly for firm 2.

- Firm 1's output ranges between $q_1 = \frac{1-c}{2}$ and $q_1 = 0$, implying we can delete all output levels $q_1 > \frac{1-c}{2}$ as NBR for firm 1. So $S_1' = \left[0, \frac{1-c}{2}\right]$
- Following a similar argument with firm 2's best response function (which is symmetric to firm 1's), we claim that $q_2 > \frac{1-c}{2}$ is NBR for firm 2, yielding a reduced strategy space $S_2' = \left[0, \frac{1-c}{2}\right]$.
- Therefore, the first step of rationalizability entails that $S_i' = \left[0, \frac{1-c}{2}\right]$ for every firm i.

Finding NBRs in the Cournot Duopoly contd.

- In the second step of rationalizability, we consider firm 1's best response function again, but taking into account that firm 2's rationalizable output levels are those in $S_2' = \left[0, \frac{1-c}{2}\right]$, that is

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 \text{ if } q_2 < \frac{1-c}{2} \\ 0 \text{ otherwise} \end{cases}$$

- As a consequence, firm 1's output now ranges between:
  - $q_1 = \frac{1-c}{2}$, which occurs when $q_2 = 0$, and
  - $q_1 = \frac{1-c}{4}$, which happens when firm 2 produces $q_2 = \frac{1-c}{2}$.
- This helps us further reduce firm 1's strategy space to $S_1'' = \left[\frac{1-c}{4}, \frac{1-c}{2}\right]$.
  - A similar argument applies to firm 2, yielding $S_2'' = \left[\frac{1-c}{4}, \frac{1-c}{2}\right]$.

Finding NBRs in the Cournot Duopoly contd.

- In the third step of rationalizability, firm 1's best response function becomes:

$$q_1(q_2) = \begin{cases} \frac{1-c}{2} - \frac{1}{2}q_2 & \text{if } \frac{1-c}{4} < q_2 < \frac{1-c}{2} \\ 0 & \text{otherwise} \end{cases}$$

- Therefore, firm 1's output now ranges between:
  - $q_1 = \frac{1-c}{4}$, which occurs when $q_2 = \frac{1-c}{2}$, and
  - $q_1 = \frac{3(1-c)}{8}$, which happens when firm 2 produces $q_2 = \frac{1-c}{4}$.
- Hence, firm 1's strategy space shrinks to $S_1''' = \left[\frac{1-c}{4}, \frac{3(1-c)}{8}\right]$.
  - Similarly for firm 2's strategy space.

Finding NBRs in the Cournot Duopoly contd.

- Repeating the process again, we see that firm 2's output levels that induce firm 1 to remain active (upper part of firm 1's best response function) keeps shrinking, approaching the 45-degree line, until the point where it does not change in further iterations of rationalizability.
- For this to occur, we need that $q_1 = q_2$. Inserting this property in the best response function yields $q_1 = \frac{1-c}{2} - \frac{1}{2} q_1$. Solving for q_1 , we obtain $q_1 = \frac{1-c}{3}$.
- Later, we will confirm that this output level coincides with the Nash equilibrium of the Cournot game, where $q_1 = q_2 = \frac{1-c}{3}$.

Nash Equilibrium

Definition. **Nash Equilibrium (NE).**

A strategy profile $s^* = (s_i^*, s_{-i}^*)$ is a Nash Equilibrium if every player chooses a best response given her rivals' strategies.

- In a two-player game, the above definition says that:
  - Strategy s_i^* is player i's best response to s_j^* and, similarly,
  - Strategy s_j^* is player j's best response to s_i^*.
- For compactness, we often write player i's best response to s_j^* as:
  - $BR_i(s_j^*) = s_i^*$,
  - and player j's as $BR_j(s_i^*) = s_j^*$.
- Therefore, a strategy profile is a NE if it is a *mutual* best response.

Tool 3.3 How to find NEs in matrix games

1. Find the best responses to all players, underlining best response payoffs.
2. Identify which cell or cells in the matrix has every payoff underlined, meaning that all players have a best response payoff. These cells are the NEs of the game.

Example 3.3: Finding NEs in a two-player game

		Firm 2	
		h	m
Firm 1	M	<u>5</u>,<u>3</u>	2,1
	L	3,6	<u>4</u>,<u>7</u>

Matrix 3.3. Finding Best Responses and NEs

Also, show that no player has strictly dominated strategies, implying that IDSDS has no bite:

$$NE = \{(M, h), (L, m)\}$$

NE in common games: Prisoner's Dilemma game

		Player 2	
		Confess	*Not Confess*
Player 1	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 3.4a. The Prisoner's Dilemma Game

NE in common games: Prisoner's Dilemma game

Player 2

Player 1		*Confess*	*Not Confess*
	Confess	-4,-4	0,-8
	Not Confess	-8,0	-2,-2

Matrix 3.4a. The Prisoner's Dilemma Game

Player 2

Player 1		*Confess*	*Not Confess*
	Confess	<u>-4</u>,<u>-4</u>	<u>0</u>,-8
	Not Confess	-8,<u>0</u>	-2,-2

Matrix 3.4b. The Prisoner's Dilemma Game – Underlining Best Response payoffs

$$NE = \{(Confess, Confess)\}$$

NE in common games: Battle of the Sexes game

		Wife	
		Football	Opera
Husband	Football	10,8	6,6
	Opera	4,4	8,10

Matrix 3.5a. The Battle of the Sexes Game

NE in common games: Battle of the Sexes game

Wife

		Football	Opera
Husband	Football	10,8	6,6
	Opera	4,4	8,10

Matrix 3.5a. The Battle of the Sexes Game

Wife

		Football	Opera
Husband	Football	<u>10</u>,<u>8</u>	6,6
	Opera	4,4	<u>8</u>,<u>10</u>

Matrix 3.5b. The Battle of the Sexes Game – Underlining Best Response payoffs

$$NE = \{(F, F), (O, O)\}$$

NE in common games: Stag Hunt game

		Player 2	
		Stag, S	Hare, H
Player 1	Stag, S	6,6	1,4
	Hare, H	4,1	2,2

Matrix 3.6a. The Stag Hunt Game

NE in common games: Stag Hunt game

Player 2

Player 1		Stag, S	Hare, H
	Stag, S	6,6	1,4
	Hare, H	4,1	2,2

Matrix 3.6a. The Stag Hunt Game

Player 2

Player 1		Stag, S	Hare, H
	Stag, S	<u>6,6</u>	1,4
	Hare, H	4,1	<u>2,2</u>

Matrix 3.6b. The Stag Hunt Game – Underlining Best Response payoffs

$$NE = \{(S,S),(H,H)\}$$

NE in common games: The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	-8,10
	Stay	12,-8	-30,-30

Matrix 3.7a. Anticoordination Game

NE in common games: The Game of Chicken

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	-8,10
	Stay	12,-8	-30,-30

Matrix 3.7a. Anticoordination Game

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	$\underline{-8},\underline{10}$
	Stay	$\underline{12},\underline{-8}$	-30,-30

Matrix 3.7b. Anticoordination Game – Underlining Best Response payoffs

$$NE = \{(Swerve, Stay), (Stay, Swerve)\}$$

Multiple Nash Equilibria

- In some games, such as the Battle of the Sexes, the Stag Hunt game, and the Game of Chicken, we found more than one NE.
- The literature often relied on some equilibrium refinement tools:
    - Pareto dominance or Risk dominance, which we will cover at the end of this chapter.
    - More technical tools, such as Trembling Hand Perfect Equilibrium or Proper Equilibrium, which we will cover at the end of Chapter 5 (we need to learn mixed strategy NE first).
- An alternative often used are the so-called “focal points” from Schelling (1960), which players tend to choose in the absence of pre-play communication.
- However, it is unclear *how a specific focal point arises* (i.e., what makes a NE more salient than another):
    - it may vary depending on the time or place where players interact,
    - leading many researchers to experimentally test the emergence of focal points in simultaneous-move games with more than one NE.
    - For experimental tests, see Camerer et al. (2004).

Relationship between NE and IDSDS

If a strategy profile $s = (s_i, s_{-i})$ is a NE, then it must survive IDSDS, but the opposite does not necessarily hold:

$$s \text{ is a NE} \Rightarrow s \text{ survives IDSDS}$$
$$\not\Leftarrow$$

In other words, the NEs in a game is a subset of the strategy profiles surviving IDSDS.

- For the first line of implication, think about (C,C) in the PD game (it is a NE and also survives IDSDS).
- For the second line of implication, think about (F,F) in the BoS game (it survives IDSDS, but it's not a NE).

Relationship between NE and IDSDS

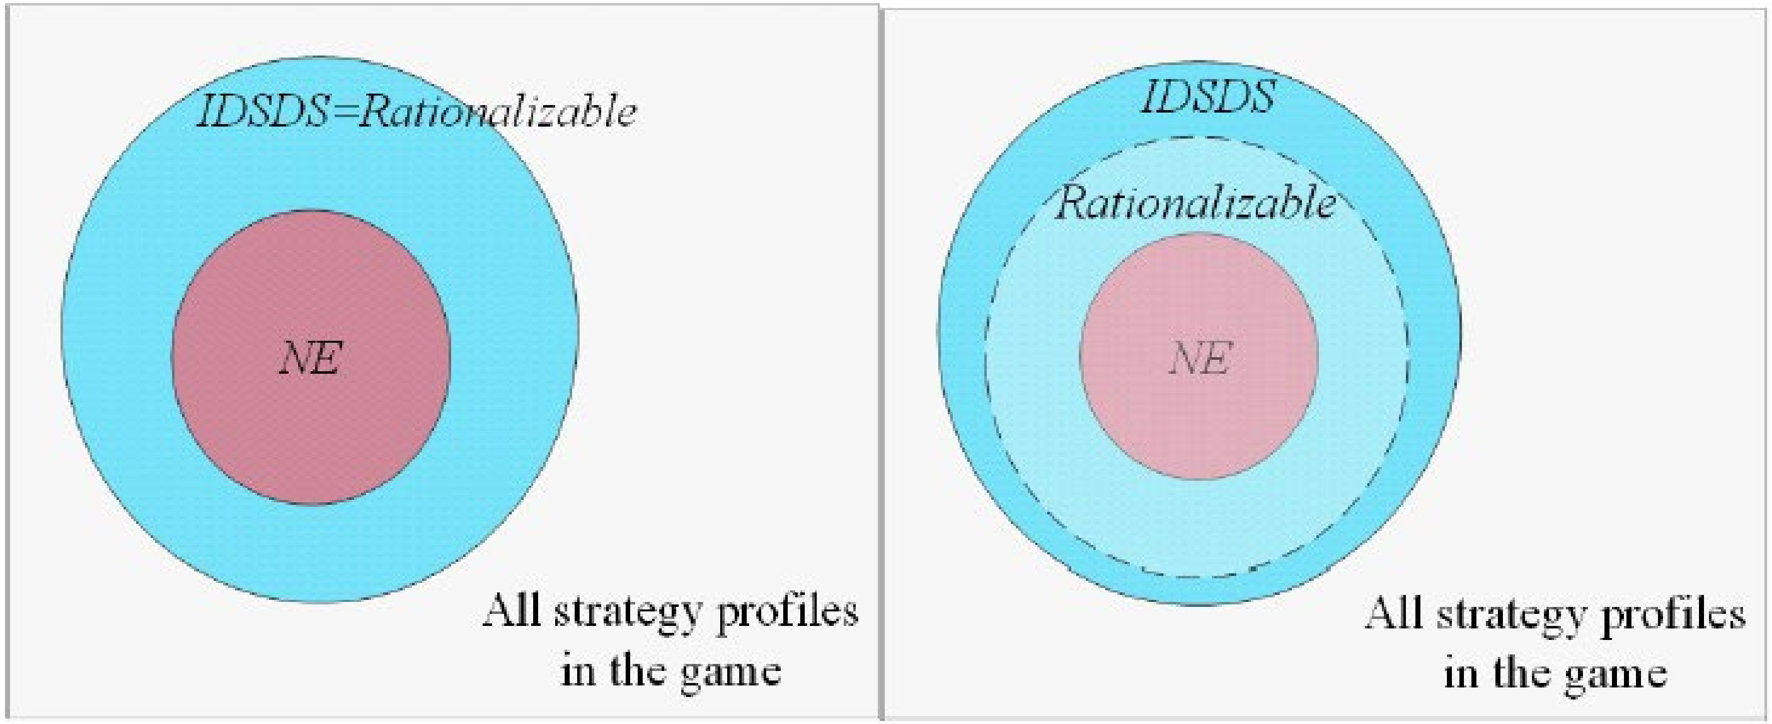

Figure 3.2a. Comparing strategy profiles - Two players.

Games with no NE

Do all games have at least one NE in pure strategies?

No, let’s see an example.

Games with no NE

		Player 2	
		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Matrix 3.8. Matching Pennies Game

Games with no NE

Player 2

		Heads	Tails
Player 1	Heads	1,-1	-1,1
	Tails	-1,1	1,-1

Matrix 3.8. Matching Pennies Game

Player 2

		Heads	Tails
Player 1	Heads	$\underline{1}$,-1	-1,$\underline{1}$
	Tails	-1,$\underline{1}$	$\underline{1}$,-1

Matrix 3.9. Matching Pennies Game – Underlining Best Response payoffs

Evaluating NE as a solution concept

1. Existence? Yes.
   - A NE may not exist in pure strategies, but then there must be one when we allow for mixed strategies.
   - This result is often known as "Nash's existence theorem."
   - Requires the use of mixed strategies where players randomize their choices (Chapter 5)
2. Uniqueness? No.
   - While the above discussion helps us rest assured that a NE will arise in most games, we cannot claim that every game will have a unique NE.
   - As an example, recall BoS or Chicken games where more than one NE exists.
3. Robust to small payoff perturbations? Yes.
   - If we alter the payoff of one of the players by a small amount, ε, we would find the same NE than before the payoff change.

Evaluating NE as a solution concept

4. Socially optimal? No.
   - As discussed, in the Prisoner's Dilemma game, the NE of a game does not need to be socially optimal.
   - A similar argument applies to the following "Modified BoS," where (F,F) and (O,O) are both NEs, but (O,O) Pareto dominates (F,F).

		Wife	
		Football	Opera
Husband	Football	8,8	6,6
	Opera	4,4	10,10

Matrix 3.10. Modified Battle of the Sexes Game

Evaluating NE as a solution concept

	IDSDS	IDWDS	SDE	Rationalizability	NE
Existence	Yes	Yes	No	Yes	Yes
Uniqueness	No	No	Yes	No	No
Robustness of payoff changes	Yes	Yes	Yes	Yes	Yes
Pareto optimal	No	No	No	No	No

Then, it’s similar to IDSDS or Rationalizability in its properties, but has more “bite” than any of them.

Appendix

Appendix – Equilibrium Selection

- Coordination and anticoordination games give rise to two NEs.
  - In games with two players but N discrete strategies, we can have N strategy profiles sustained as NEs (see examples in the EconS 424 website).
- In games with more than one NE, a typical question is whether we can select one of them as being more "natural" or "reasonable" to occur according to some criteria.
- The literature offers several *equilibrium selection criteria* and we present the most popular below.
- Throughout our presentation, let $s^* = (s_i^*, s_{-i}^*)$ denote a NE strategy profile.

Appendix – Pareto Dominance

Definition. **Pareto Dominated NE**.

A NE strategy profile s^* Pareto dominates another NE strategy profile $s' \neq s^*$ if

$$u_i(s^*) > u_i(s') \text{ for every player } i.$$

Example:

		Firm 2	
		Tech. A	Tech. B
Firm 1	Tech. A	<u>10</u>,<u>10</u>	0,2
	Tech. B	2,0	<u>5</u>,<u>5</u>

Matrix 3.11. Technology Coordination Game

$$NE = \{(A, A), (B, B)\}$$

Here (A, A) Pareto dominates (B, B), since every firm earns a higher profit at (A, A) than at (B, B), i.e., $10 > 5$.

Appendix – Pareto Dominance

- Pareto dominance, however, does not help us rank NEs in *all* games.

Example

		Player 2	
		Swerve	Stay
Player 1	Swerve	-1,-1	-8,10
	Stay	12,-8	-30,-30

Matrix 3.7b. Anticoordination Game – Underlining Best Response payoffs

$$NE = \{(Swerve, Stay), (Stay, Swerve)\}$$

In this case, we cannot select $(Stay, Swerve)$ over $(Swerve, Stay)$ or vice versa, since player 1 prefers $(Stay, Swerve)\ to\ (Swerve, Stay)$, while player 2 prefers otherwise.

Appendix – Risk Dominance

- Pareto dominance helps us compare NEs where, for instance, both firms simultaneously change their technology decision from B to A, yielding (A, A).
- But what if each firm unilaterally deviates from NE?
  - If (B, B) is played, every firm's payoff loss when unilaterally switching to A is $5 - 0 = 5$.
  - But if (A, A) is being played, every firm loses $10 - 2 = 8$ when unilaterally deviating to B.
- Informally, at (A, A), unilateral deviations are "riskier" than at (B, B).
  - Note that, as opposed to Pareto dominance, we now seek to understand the payoff loss that players can suffer in the path when each of them unilaterally moves from one NE to another.
- How to measure the "riskiness" of deviations in a more formal way?

Appendix – Risk Dominance

- Definition. **Risk Dominated NE.**
- A NE strategy profile $s^* = (s_1^*, s_2^*)$ risk dominates another NE strategy profile $s' = (s_1', s_2')$ if the total payoff loss of moving from s' to s^*,

$$[u_1(s_1^*, s_2^*) - u_1(s_1', s_2^*)] \times [u_2(s_1^*, s_2^*) - u_2(s_1^*, s_2')]$$

exceeds that of moving from s^* to s',

$$[u_1(s_1', s_2') - u_1(s_1^*, s_2')] \times [u_2(s_1', s_2') - u_2(s_1', s_2^*)]$$

Appendix – Risk Dominance

Example:

		Firm 2	
		Tech. A	Tech. B
Firm 1	Tech. A	10,10	0,2
	Tech. B	2,0	5,5

Matrix 3.11. Technology Coordination Game

We can say that (B, B) risk dominates (A, A) because the total payoff loss of moving from (B, B) to (A, A),

$$[u_1(A, A) - u_1(B, A)] \times [u_2(A, A) - u_2(A, B)]$$
$$= (10 - 2) \times (10 - 2) = 64$$

exceeds that of moving from (A, A) to (B, B)

$$[u_1(B, B) - u_1(A, B)] \times [u_2(B, B) - u_2(B, A)]$$
$$= (5 - 0) \times (5 - 0) = 25.$$