

# 15 Game Theory

*Varian: Chapters 28-29.*

The key novelty compared to the competitive (Walrasian) equilibrium analysis is that game theoretic analysis **allows for the possibility that utility/profit/payoffs depend directly on what other people do**. Moreover, all *players* are “intelligent” in the sense that they **understand the structure of the economy**. One can think of the difference between standard competitive price theory and game theory as follows.

In terms of the mathematical model we will now let  $s_i$  denote some choice variable of agents  $i = 1, \dots, n$  and write the **payoff function for agent  $i$**  (that is the profit function in applications where participants are firms, utility function if participants are individual agents) as

$$u_i(s_1, s_2, \dots, s_n).$$

The advantage of this compared to a competitive type of model is simply that it allows us to analyze situations that can't be handled with the tools we've used so far. Examples include:

- **Oligopoly markets:** The price or quantity decisions by one firm has to directly influence profits by other firms if firms are competing for customers.
- **Auctions:** If a player gets the object depends on own bid and bids by others.
- **Public Goods:** To the extent that the programming depends on contributions to the local public radio station, the utility for a person who listens to public radio depends on the level of contributions from other people in the area.
- **Patent Races:** The likelihood of getting a patent depends on own research effort, but also on research effort by others.

## 15.1 Basic Concepts

A *game* refers to any “social situation” that involves two or more participants. Specifically, a game always has three components;

1. When specifying a game we need to be explicit about who the participants are. These are called *players in the lingo of game theory*. Depending on the application, a “player” may be a firm in a market, a consumer, a general at war, or even a participant in a game of poker. Usually, we label the players by  $i = 1, \dots, n$ .
2. We also need to be explicit about what every player can conceivably do. The choice variable of a player is referred to as a *strategy*, and often times we use the notation  $s_i$  for a generic choice (=strategy) for player  $i$ . In an oligopoly example  $s_i$  could be a price or a quantity, or (say when there is a time element) a contingent pricing plan. In a simple game of military conflict  $s_i$  could be “storm that hill” or “don’t storm that hill”
3. The final (and crucial) component of a game is a *payoff function* for each agent that specifies how each player evaluates every *strategy profile* (=a strategy for each agent). I.e., to figure out what each agent wants to do and what they should expect others to do we need to know the utility for every choice of the agent as well as every choice of everyone else. Hence, the payoff function is on form  $u_i(s_1, \dots, s_i, \dots, s_n)$ .

## 15.2 Equilibrium in Games

The point with a model is to be able to say something about what will likely happen in a given situation, so we need somehow to be able to predict what will happen. That is we need some *solution concept* or a concept of *equilibrium* that applies to a game. It turns out however that we need different notions of equilibrium for different sorts of games, so we will actually need a few different equilibrium concepts. We will focus on the three most commonly used: *Dominant Strategy equilibrium*, *Nash equilibrium* and *Subgame perfect Nash equilibrium* (or “*Backwards induction equilibrium*”).

## 15.3 Payoff Matrices and Dominant Strategies

The ideas are best illustrated through examples. Consider first a simple example where,

- There are two players, Axel and Birgitta
- Two possible strategies for each player. Call them “top” ( $T$ ) and “bottom” ( $B$ ) for Axel and “left” ( $L$ ) and “right” ( $R$ ) for Birgitta.

A payoff function for Axel then gives the utility for each combination of his own choice and Birgittas choice, that is a “happiness index” for each of the four possible pairs ( $T, L$ ), ( $T, R$ ), ( $B, L$ ) and ( $B, R$ ). Similarly, Birgittas payoff function also gives a happiness index for each of the four strategy pairs. Suppose that the payoff functions are

$$\begin{array}{rcl}
 & u_A(T, L) = 1 & u_B(T, L) = 2 \\
 \text{Axel:} & u_A(T, R) = 0 & \text{Birgitta: } u_B(T, R) = 1 \\
 & u_A(B, L) = 2 & u_B(B, L) = 1 \\
 & u_A(B, R) = 1 & u_B(B, R) = 0
 \end{array}$$

We can then represent this simple two player game in a *payoff matrix* as in Figure 1. Observe that all information in the specification above is in the matrix. We know who the players are (Axel & Birgitta), what the players can do (“top”, “bottom” and “left”, “right”.) and have they evaluate the options. Note that the convention that we follow is to give the payoff for the “row-player” (Axel) first and then the payoff for the “column-player” (Birgitta).

		Birgitta	
		Left	Right
Axel	Top	1,2	0,1
	Bottom	2,1	1,0

Figure 1: Representing a two player game in a payoff matrix

Now, the question is **what would happen in this situation?** That is, **what is a natural notion of equilibrium?** Inspecting the payoff matrix we see that:

1. If Birgitta plays “Left”, then Axel gets 2 if playing “Bottom” and 1 if choosing “Top”.  
Hence “Bottom” is better than “Top” if Birgitta would play “Left”
2. If Birgitta plays “Right”, then Axel gets 1 if playing “Bottom” and 0 if choosing “Top”.  
Hence “Bottom” is better than “Top” if Birgitta would play “Right”
3.  $\Rightarrow$  “Bottom” is better than “Top” **no matter what Birgitta is doing**. Hence it seems like a minimal requirement of rationality to predict that Axel would indeed play “Bottom” since this is the best **independently of what Axel thinks Birgitta will do**.
4. Similarly, if Axel plays “Top”, Birgitta gets 2 if playing “Left” and 1 if playing “Right”, so “Left” is better than “Right” if Axel plays “Top”.
5. If Axel plays “Bottom”, Birgitta gets 1 if playing “Left” and 0 if playing “Right”, so “Left” is better than “Right” if Axel plays “Bottom”.
6.  $\Rightarrow$  “Left” is better than “Right” **no matter what Birgitta is doing**. Hence, we’d think that Birgitta would indeed play “Left”.
7. We can then conclude that the only natural prediction in this example is (“Bottom”, “Left”)

We have now “solved” the game in the sense that given the game summarized in Figure 1 our prediction is (“Bottom”, “Left”). The procedure can be applied in many other (although not all) situations to get predictions about what will happen. In general:

- If a strategy is better than all other strategies independently of the strategies chosen by the other players we call the strategy a *dominant strategy*.
- If all players have a dominant strategy we say that the strategy profile where all players play its dominant strategy is a *dominant strategy equilibrium*

### 15.3.1 Formal Definition

In a general (2 player) game we can define dominance as follows:

**Definition 1**  $s_1^*$  is called a (weakly) dominant strategy for player  $i$  if  $u_1(s_1^*, s_2) \geq u_1(s_1, s_2)$  for all  $s_1 \neq s_1^*$  and all  $s_2$  and  $u_1(s_1^*, s_2) > u_1(s_1, s_2)$  for all  $s_1 \neq s_1^*$  and some  $s_2$ .

The definition for player 2 is symmetric-just reverse the roles of the players. Note that the second part of the definition rules out calling something a dominant strategy if  $u_1(s_1, s_2) = k$  for all  $s_1$  and  $s_2$ , while still allowing indifference given some choices by the other player.

**Definition 2**  $(s_1^*, s_2^*)$  is a dominant strategy equilibrium if  $s_1^*$  and  $s_2^*$  are both dominant strategies.

It is trivial to generalize the notion to  $n > 2$  player games, but then the payoff is  $u_i(s_1, \dots, s_n)$  so we write the inequality for player 1 as  $u_1(s_1^*, s_2, \dots, s_n) \geq u_1(s_1, s_2, \dots, s_n)$  for all  $s_1 \neq s_1^*$  and  $s_2, \dots, s_n$ . For an arbitrary  $i$  we would write  $u_i(s_1, \dots, s_{i-1}, s_i^*, s_{i+1}, \dots, s_n) \geq u_i(s_1, \dots, s_{i-1}, s_i, s_{i+1}, \dots, s_n)$ . The idea is perfectly the same: if for each conceivable strategy profile (a strategy for each player)  $s_i^*$  gives a higher payoff than any other  $s_i$  we refer to it as a dominant strategy.

## 15.4 The Prisoner's Dilemma

**Story:** Imagine that Axel and Birgitta have been out on a crime-spree and got caught. They are questioned in separate cells and have to decide whether to confess or not. They are told (and/or know that this is true);

1. if one of them confesses (and agrees to be a witness against the partner), he/she will go free in exchange for being a key witness. The person not confessing is in for a long time in jail.
2. if nobody confesses they know that they can only be nailed on some minor offense.

3. if both confess they get some time off relative to the case when only the partner in crime confesses (because they are cooperative), but not much.

		BIRGITTA	
		Confess	Deny
AXEL	Confess	-30,-30	0,-40
	Deny	-40,0	-1,-1

Figure 2: A Prisoners Dilemma

We can capture the strategic nature of this story as a game with payoff matrix gives by the one in Figure 2. We may think of the entries as the negative of the number of years in prison associated with each possibility. Now we see that **confess is a dominant strategy** ( $-30 > -40$  and  $0 > -1$ ), so the prediction of this game, the **dominant strategy equilibrium**, is that both agents will confess, resulting in 30 years of prison for them each.

One important lesson from this example is that in game theory, **equilibria need not be Pareto optimal**. If both would deny, they would both be better off than in the only equilibrium of the game. However, both denying is not a self-enforcing situation since each player then would have an incentive to deviate by confessing.

The payoff structure in the prisoner's dilemma arises naturally in many contexts with more economic relevance. One example is when agents need to cooperate for some common good. For a simple example of this, suppose that Axel and Birgitta runs a coop (say, making organic fair trade shade grown cappuccinos). Assume that sales depend on whether they are lazing or working hard (if the line becomes too long then customers run away). To make it simple, suppose that effort can only take on two values,  $L$  and  $H$  and denote by  $e^A$  the choice of effort by Axel and  $e^B$  the choice of effort by Birgitta. We assume that sales,  $y$ , is an increasing function of effort and that it is symmetric in the effort by Axel and Birgitta

so that

$$y(L, L) < y(L, H) = y(H, L) < y(H, H).$$

Next, we assume that Axel and Birgitta are good egalitarians, so that they've decided to split income equally. Moreover, they are also lazy, meaning that for a given level of income they prefer low effort. We formalize this by assuming that there is a strictly positive utility cost  $c_H$  associated with high effort, whereas  $L$  costs nothing (by normalization). Now, low effort is a (strictly) dominant strategy if

$$\begin{aligned} \frac{y(L, L)}{2} &> \frac{y(L, H)}{2} - c_H \text{ and} \\ \frac{y(L, H)}{2} &> \frac{y(H, H)}{2} - c_H. \end{aligned}$$

Low effort being a dominant strategy does not make the game into a prisoners dilemma in itself. We call it a Prisoners dilemma only if the efficient thing to do would be for both agents to expand high effort. That is, if

$$\begin{aligned} y(H, H) - 2c_H &> y(H, L) - c_H = y(L, H) - c_H \text{ and} \\ y(H, H) - 2c_H &> y(L, L) \end{aligned}$$

For example, let

$$\begin{aligned} y(L, L) &= 0 \\ y(L, H) &= y(H, L) = 1 \\ y(H, H) &= 2 \\ c_H &> \frac{1}{2} \end{aligned}$$

Then,

$$\frac{y(L, L)}{2} = 0 > \frac{1}{2} - c_H = \frac{y(L, H)}{2} - c_H$$

and

$$\frac{y(L, H)}{2} = \frac{1}{2} > 1 - c_H = \frac{y(H, H)}{2} - c_H,$$

so low effort is a dominant strategy, implying that  $(e^A, e^B) = (L, L)$  is the dominant strategy equilibrium. Moreover, if  $c_H < 1$  then

$$y(H, H) - 2c_H = 2(1 - c_H) > 1 - c_H = y(H, L) - c_H = y(L, H) - c_H$$

and

$$y(H, H) - 2c_H = 2(1 - c_H) > 0 = y(L, L),$$

so  $(e^A, e^B) = (H, H)$  is the only efficient outcome. We conclude that the game is a prisoners dilemma if  $\frac{1}{2} < c_H < 1$  (and the given relation  $y(e^A, e^B)$  between effort and output postulated above).

Similar logic can be shown that the *tragedy of the commons* may be interpreted as a prisoners dilemma situation. Unilateral determination of tariffs in international trade is another important example.

## 15.5 Nash Equilibrium

Now consider the game in Figure 3. Observe that:

- If Birgitta plays “Left”, then “Middle” is the best option for Axel
- If Birgitta plays “Center”, then “Top” is the best option for Axel
- If Birgitta plays “Right”, then “Bottom” is the best option for Axel
- We conclude that there is no dominant strategy for Axel and you can verify that the same is true for Birgitta. Neither is there a *dominated* strategy (something that is worse than some other strategy no matter what the opponent is doing).

Hence, if we are only going to rule out dominated strategies, then our prediction is that in the strategic situation described in Figure 3 “something will happen”, which is not very precise or instructive. However, a natural (really clever) idea is to say that an “equilibrium” is a situation where **each agent behaves optimally given the strategy chosen by the other(s)**, which is the concept of *Nash equilibrium*. I.e.,

		Birgitta		
		Left	Center	Right
Axel	Top	0,4	4,0	5,3
	Middle	4,0	0,4	5,3
	Bottom	3,5	3,5	6,6

Figure 3: A Game with no Dominant Strategy

**Definition 3** A strategy profile (a strategy for each player) is a Nash equilibrium if each player plays a strategy that is optimal given the strategies played by the other players in the game. That is (in a two player game)  $(s_1^*, s_2^*)$  is a Nash equilibrium if;

1.  $u_1(s_1^*, s_2^*) \geq u_1(s_1, s_2^*)$  for all  $s_1$  in the set of available strategies for player 1
2.  $u_2(s_1^*, s_2^*) \geq u_2(s_1^*, s_2)$  for all  $s_2$  in the set of available strategies for player 2

A useful interpretation is that a Nash equilibrium requires that every player maximizes his/her utility given some beliefs about what other players will do and that the beliefs are correct. Hence, unlike dominant strategy equilibrium the concept of Nash equilibrium requires **rational expectations-every agent knows exactly what the other players are going to do in a Nash equilibrium**, which may seem weird, but is the solution to the problem that otherwise players would have to form conjectures about other players play...and conjectures about other players conjectures about their own play...and conjectures about conjectures about conjectures about....

Consider Figure 3. In a matrix form game like this, the fool-proof algorithm for finding all (pure strategy) Nash equilibria is to first assume that Birgitta plays “Left” and then simply check the first column of the payoff matrix to see where Axel gets the highest utility. In this example, it turns out to be “Middle” , which we call the **best response** to “Left” . To keep track, underline the relevant entry in the payoff matrix (i.e., the payoff for Axel from (Middle,Left)). Next, assume that Birgitta plays “Center”. Then, “Top” is the best response

		Birgitta		
		Left	Center	Right
Axel	Top	0, <u>4</u>	<u>4</u> , 0	5, 3
	Middle	<u>4</u> , 0	0, <u>4</u>	5, 3
	Bottom	3, 5	3, 5	<u>6</u> , <u>6</u>

Figure 4: "Bottom", "Right" is the Unique Nash Equilibrium

and we underline the corresponding entry in the payoff matrix, and in the final column we see that "Bottom" is the best response to "Right", which leaves 6 underlined. Next, do the same thing for Birgitta (who picks the column). In a Nash equilibrium each player must behave optimally given the choice of the other player, implying that every cell where we have underlined both the payoff for Axel and the payoff for Birgitta is a Nash equilibrium. In this example, the only Nash equilibrium is (Bottom, Right).

### 15.5.1 The Battle of the Sexes

		Birgitta	
		Hockeyl	Opera
Axel	Hockey	2, 1	0, 0
	Opera	0, 0	1, 2

Figure 5: The Battle of the Sexes

**Story:** Now imagine Axel and Birgitta are discussing where to meet for a date (obviously this discussion takes place using cell phones). They have eliminated all other options but a hockey game and the opera. Suddenly, Axels' cell phone dies and the connection is broken. There is no way for them to communicate with each other before the actual time for the date, so they both have to guess what the other will do when deciding on whether to go to

the hockey game or the opera. Suppose that Axel prefers hockey to opera, and that Birgitta prefers opera to hockey. However, what they REALLY care about is going out for a date. Hence, if Axel goes to the hockey game and Birgitta doesn't show, he is so miserable that he may as well be (alone) at the opera, whereas opera with Birgitta is better than being alone (at either place). The situation for Birgitta is symmetric. We can represent this strategic situation as in Figure 5.

Proceeding as with the last game we see that:

- “Hockey”, “Hockey” is a Nash equilibrium.
- “Opera”, “Opera” is also a Nash equilibrium

The point with the example is twofold:

1. There may be several Nash equilibria (there is two “pure strategy” equilibria and also a mixed strategy equilibrium where both agents randomize).
2. When there are several Nash equilibria, then it is not obvious that we would expect real world agents to always coordinate their play on Nash equilibrium behavior. There is one equilibrium where both go to the game and another where both go to the opera, so it doesn't seem obvious that we would never see outcomes where Axel goes to the game and Birgitta goes to the opera (indeed, there is a mixed strategy equilibrium where there is mis-coordination with positive probability).

## 15.6 Mixed Strategies

Figure 6 depicts a game (probably familiar to many) called “Rock, Scissors, Paper”. The idea is that Rock beats Scissors, Scissors beats Paper and Paper beats Rock, whereas there is a tie if both players pick the same strategy. The payoff is 1 for the winner and  $-1$  for the loser, and both agents get 0 in case of a tie.

- It is easy to check that there is no *pure strategy* equilibrium. The reason is simply that if you know what your opponent is doing, then you can win for sure.

		Birgitta			
			Rock	Scissors	Paper
Axel	Rock	0,0	1,-1	-1,1	
	Scissors	1,-1	0,0	1,-1	
	Paper	1,-1	-1,1	0,0	

Figure 6: A Game with no Pure Strategy Nash Equilibrium

- The solution to this problem is to *randomize*. Therefore we allow the agents to pick a probability distribution over the three options. We call such a randomization a *mixed strategy*. In this game of “Rock, Scissors, Paper”, a mixed strategy for Axel is a triple  $(a_R, a_S, a_P)$ , where  $a_R \geq 0, a_S \geq 0$  and  $a_P \geq 0$  and  $a_R + a_S + a_P = 1$ . Similarly, a mixed strategy for Birgitta is a triple  $(b_R, b_S, b_P)$ , where  $b_R \geq 0, b_S \geq 0$  and  $b_P \geq 0$  and  $b_R + b_S + b_P = 1$ .
- Suppose that Birgitta plays mixed strategy  $(b_R^*, b_S^*, b_P^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ . The expected payoff for Axel from playing Rock is then

$$b_R^*0 + b_S^*1 + b_P^*(-1) = \frac{1}{3}0 + \frac{1}{3}1 + \frac{1}{3}(-1) = 0.$$

If Axel plays paper we conclude that the expected payoff is

$$b_R^*1 + b_S^*(-1) + b_P^*0 = \frac{1}{3}1 + \frac{1}{3}(-1) + \frac{1}{3}0 = 0,$$

and if Axel plays scissors his payoff is

$$b_R^*(-1) + b_S^*0 + b_P^*1 = \frac{1}{3}(-1) + \frac{1}{3}0 + \frac{1}{3}1 = 0.$$

- We conclude that Axel is *indifferent* between all is available pure strategies. Hence, it is optimal for Axel to pick any probability distribution over the pure strategies “Rock, Scissors, Paper”. In particular,  $(a_R^*, a_S^*, a_P^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$  is optimal when Birgitta randomizes in accordance with  $(b_R^*, b_S^*, b_P^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ .

- But, assuming that Axel picks the mixed strategy  $(a_R^*, a_S^*, a_P^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ , then we can appeal to the same calculations as above to conclude that it is optimal for Birgitta to pick any probability distribution over the pure strategies “Rock, Scissors, Paper”. In particular,  $(b_R^*, b_S^*, b_P^*) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$  is optimal when Axel randomizes in accordance with  $(a_R^*, a_S^*, a) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ .
- We conclude that  $\{(a_R^*, a_S^*, a_P^*) (b_R^*, b_S^*, b_P^*)\} = \{(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})\}$  is a Nash equilibrium (in mixed strategies).
- Indeed, this is the only Nash equilibrium in the game. It is a good exercise to verify this. To do so, you need to rely on the fact that the only way to make it optimal for a player to randomize is that the player is indifferent between the pure strategies that are played with positive probability.

### 15.6.1 The Mixed Strategy Equilibrium in the Battle of the Sexes

For games with only two actions for each player it is easy to compute mixed strategy equilibria. Take the battle of the sexes game discussed above and assume that there is a mixed strategy equilibrium (we will verify this by the calculation below). Let  $p$  be the probability that Birgitta plays “Hockey” and  $q$  be the probability that Axel plays “Hockey”. Then, the expected payoff for Axel if he plays “Hockey” is

$$p2 + (1 - p)0 = 2p,$$

whereas “Opera” gives

$$p0 + (1 - p)1 = 1 - p$$

Now, if  $2p > 1 - p$  the optimal response for Axel is “Hockey”, while if  $2p < 1 - p$ , then the optimal response is “Opera”. This means that for it to be optimal for Axel to pick a nondegenerate randomization (some  $0 < q < 1$ ) it must be that Axel is indifferent, which he is if and only if

$$2p = 1 - p \Leftrightarrow p = \frac{1}{3}.$$

A similar calculation for Birgitta yields the expected payoffs

$$q1 + (1 - q)0 = q \text{ if she plays "Hockey"}$$

$$q0 + (1 - q)2 \text{ if she plays "Opera",}$$

Indifference requires that  $q = (1 - q)2 \Leftrightarrow q = \frac{2}{3}$ . We conclude that  $(p, q) = (\frac{1}{3}, \frac{2}{3})$  is the only Nash equilibrium in mixed strategies in this game. Hence, there are three Nash equilibria in total in the Battle of the sexes; (Hockey,Hockey), (Opera, Opera) and a mixed equilibrium where Axel goes to the hockey game with probability  $\frac{2}{3}$  and Sheila goes to the Hockey game with probability  $\frac{1}{3}$ . In this equilibrium the probability that they actually meet is

$$\frac{1}{3} \frac{2}{3} + \frac{2}{3} \frac{1}{3} = \frac{4}{9},$$

which indeed is lower than if they simply randomized with equal probabilities.

## 15.7 Extensive Form Games




Figure 7: The Battle of the Sexes in Extensive Form

So far we have described games by using what is called the “normal form representation”, which turns out to be convenient for games where players move simultaneously (as in all examples so far). However, there is an alternative way to depict games using a “game tree”, which is referred to as **the extensive form**. Figure 7 depicts the “battle of the sexes” using this alternative (but equivalent) way. Every place where some player makes a decision is a node in the graph and as you see the lines fully describes all possibilities in the game. Observe that being a simultaneous game, Birgitta doesn’t know what Axel played when

making her decision and we have shown this in the graph by the “ring” around the nodes where Birgitta is choosing. This ring means that Birgitta must make the same decision no matter whether Axel picked opera or football. In the graph I’ve depicted one of the equilibria (“Opera”, “Opera”).

Notice that once we guess an equilibrium, then we can check that the suggested profile is an equilibrium by following the path to the endnodes following any unilateral deviation. In this example, we already know what the equilibria are, but in more complicated games, this guess and verify approach is often useful.

## 15.8 Sequential (Dynamic) Games




Figure 8: The Battle of the Sexes with Axel Moving First

**Story:** Suppose now that everything is as in the “battle of the sexes” except that Axel makes his choice before Birgitta, **where Birgitta is assumed to be able to observe what Axel did before making her choice.** To make sense of this you may imagine that Axel is the guy with the cell phones, whereas Birgitta is at home. Axel can therefore use a pay phone once he made his choice and tell Birgitta where he is (at a point where it is too late for Axel to reconsider, whereas Birgittas’ apartment is located somewhere in between the stadium and the opera house so that she can make it in time for either place). Also, Sheila has a caller ID telling the location of the pay phone, so that it is not possible for Bruce to lie.

This game we would depict as in Figure ??, where we have taken away the “ring” around the decision nodes for Birgitta to indicate that she knows where Axel is when making her

decision. It is very important to understand that we now have to rethink what a strategy is:

- For Axel, a strategy is still where to go: Opera or Hockey.
- For Birgitta a strategy is now much richer than in the simultaneous setting. It is now a *contingent plan that specifies an action (Opera or Hockey) after any possible choice that Bruce made.*
- Hence the possible strategies for Birgitta are
  1. Opera after Opera and Opera after Hockey
  2. Opera after Opera and Hockey after Hockey
  3. Hockey after Opera and Opera after Hockey
  4. Hockey after Hockey and Hockey after Hockey

For brevity, we write the available strategies as  $\{oO, oH, hO, hH\}$ , where the lowercase letters indicate the action chosen if Axel picks Opera and the capitals indicate the action if Axel picks Hockey”.

In Figure 8 we have depicted the strategy profile  $(O, oO)$ , that is where Axel picks Opera and Birgitta picks Opera no matter what Axel did. Let’s check if this is a Nash equilibrium.

1. Suppose that Axel would deviate, then given the play specified they would end up in the end node corresponding to (Hockey,Opera). Axel would then get a payoff of 0 rather than 1. Axel can therefore not gain by a unilateral deviation.
2. If Birgitta would deviate to  $fO$  or  $fF$  she would get 0 rather than 2, while if she deviated to  $oF$  she’d still get 2. Hence, Birgitta can’t gain from a deviation. The conclusion is that  $(O, oO)$  is a Nash equilibrium.

### 15.8.1 Backwards Induction

Arguably  $(O, oO)$  is not what we think would happen. Imagine yourself being Axel. You’d know that *if* you’d go to the game, then Birgitta has the choice of either going to the game

with you, or going to the opera alone. Hence, a natural line of reasoning from Axels' point of view is that going to the game would guarantee that Birgitta would also show up for the game. The same reasoning would imply that going to the opera would guarantee that also Sheila goes to the opera. Hence **taking Birgittas rational responses to each action into consideration** Axel would choose between getting 2 (going to the game with Birgitta) or getting 1 (going to the opera with Birgitta). Hence, Bruce would choose Hockey.

The reason why  $(O, oO)$  seems like an “unlikely outcome” is that it relies on a **non-credible treat**. The strategy  $oO$  means that Sheila would go to the opera even if put in a situation where she facing *fait accompli* would not want to do it. For this reason we will solve dynamic games “backwards” in order to guarantee that **every player behaves optimally after any “history”** and doing this we get only Nash equilibria that are **credible** (i.e., every player would always want to do what the strategy specifies no matter where in the game we are. This procedure is called “backwards induction” and for the example we see that the only remaining equilibrium is  $(H, oH)$ . This kind of equilibrium is called a *backwards induction equilibrium* or *subgame perfect equilibrium* (where the reason for the second name has to do with that the play is optimal for every agent in every part of the game that could be thought of as a game in itself).

### 15.8.2 Repeated Prisoners Dilemma

One way one may think of as sustaining good outcomes in situations where agents have incentives to defect is by a construction where if there is a deviation, then the agent is punished by the other agent(s) by switching to a bad outcome. This will sometimes work, but not always.

To think about this, let's consider a prisoners dilemma situation. To minimize the arithmetic involved, simplify the static payoff matrix to have defection by each player generate a payoff of 0, cooperation generate a payoff of 1 for each player, and the off-diagonal elements be  $(-1, 2)$  and  $(2, -1)$ . For simplicity, suppose the game is played twice, and assume that the players can see the outcome of the “first round” before playing it a second time. I have

		Birgitta	
		Defect	Cooperate
Axel	Defect	0,0	2-1
	Cooperate	-1,2	1,1

Figure 9: A Simpler Prisoners Dilemma

also changed the labeling of the strategies. I now label the strategy that is best for the agents collectively “cooperate” and the other strategy “defect”. There are several reasons for this; (i) the original story seems a bit of a stretch if it is repeated, (ii) the basic point with looking at repetitions is to see whether this can help to sustain collectively good outcomes when there is a short term incentive to cheat, (iii) to make the notes consistent with the vast literature on the subject.

A strategy is a *full contingent plan*. In this context that means an “action” (either  $C$  or  $D$ ) in the first period and an action (either  $C$  or  $D$ ) for each of the possible outcomes in the first period. To write down the set of all possible strategies, let lowercase letters denote the action following  $C$  by the opponent in period 1 (it is redundant to condition on the own first period action), so that we may write

$$C_1cD$$

for the strategy that has the player cooperate in the first period, that cooperates in the second period if the other player cooperated in the first, and defects if the other player defected in the first period. In total, there are now 8 strategies for each player;

$$\{C_1cC, C_1cD, C_1dC, C_1dD, D_1cC, D_1cD, D_1dC, D_1dD\}.$$

**Backwards Induction** Let’s first consider backwards induction equilibria. Then, in the second period, the players play a static prisoners dilemma. Hence,  $(D, D)$  must be played in the second period after any history of play. But then, in the first period both players

know that, regardless of what happens in the first period,  $(D, D)$  is played in the second period. Hence, also the first period is a static prisoners dilemma. We conclude that the unique backwards induction equilibrium is for both players to play  $D_1dD$ , that is, always defect, regardless of the history of play.

This logic extends to an arbitrary (finite) rounds of play. In the last period, we will have defect for any history. Hence, in the second to last period, we will have defect after any history, and so on.

**Nash Equilibria** We saw above that it is sometimes possible to support non-credible Nash equilibria, by threats of punishments that would be non-optimal if they actually need to be carried out. However, in the Prisoners dilemma this turns out to be impossible. To see this, observe that

1. In the static game,  $D$  is a strictly dominant strategy. Hence;

- $C_1dD$  is strictly better than either  $C_1cC$  or  $C_1cD$  if the other player cooperates in the second period.
- $D_1dD$  is strictly better than either  $D_1cC$  or  $D_1cD$  if the other player cooperates in the second period.
- $C_1dD$  is strictly better than either  $C_1dC$  or  $C_1cC$  if the other player defects in the second period.
- $D_1dD$  is strictly better than either  $D_1dC$  or  $D_1cC$  if the other player defects in the second period.

$\Rightarrow$ neither  $C_1cC$  or  $D_1cC$  can be part of a Nash equilibrium.

2. Suppose that player 1 cooperates in the first period. That is, she plays  $C_1cD$ ,  $C_1dC$  or  $C_1dD$ . Then,

- Suppose she plays  $C_1dD$ . Then, the opponent knows that player 1 will defect in the second period regardless of history. Hence the opponent must play either

$D_1dD$  or  $D_1dC$ , in which case player 1 gets a payoff of  $-1 + 0 = -1$ . Clearly  $D_1dD$  is strictly better since it gives at least 0 (either 0 or 2)

- Suppose she plays  $C_1cD$ . Then, if the opponent plays any strategy involving cooperation in the first period (in which case player 1 cooperates in the second period) then  $C_1dD$  is strictly better. Suppose instead that the opponent defects in period 1. Then the payoff for player 1 is

$$-1 + 0$$

if the opponent defects in the second period and

$$-1 + 2$$

if the opponent cooperates in the second period. In the first case, a profitable deviation for 1 is  $D_1dD$  (worst that can happen is a payoff of zero). In the second case the other player would be better off deviating so that he plays defect in the second period.

- Finally, suppose that she plays  $C_1dC$ . If the other player defects in period 1,  $C_1dD$  is strictly better. If the other player cooperates in period 1 then her payoff is

$$1 + 0$$

if the other player defects in period 2. In this case, a defection in the first period would yield at least

$$2 + 0$$

since 0 is the worst case scenario from a defection in period 2. If the other player defects in period 1, that player must defect also in the second period on the equilibrium path (following  $C_1$  by player 1), so the payoff for player 1 must be

$$-1 + 0$$

Again,  $D_1dD$  is a profitable deviation.

3. Hence, we have ruled out any strategy with cooperate in the first period being part of a Nash equilibrium. The game is symmetric, so the same is true for the other player as well.
4. Given that both players defect in the first period, any strategy that calls for a cooperation after first period defection can also be ruled out due to defect being a dominant strategy in the static game. We conclude that, in any Nash equilibrium, both players must defect in both periods.
5. Again, this argument can be extended to an arbitrary finite number of periods (but this is kind of hard).

**Infinite Repetitions** In the prisoners dilemma game, no finite repetition is of any help at all, not even when allowing the players to use non-credible threats. However, we **can** get more desirable outcomes if we allow for an infinite number of rounds (where the infinity can be reinterpreted as a situation where a coin biased is tossed after every round and the game continues if it comes up heads...this would almost certainly end in finite time).

The simplest construction is as follows. Let Axel follow the strategy:

1. Play  $C$  in every period if  $(C, C)$  has been played in every previous period.
2. If anything different has ever been played, then play  $D$  in every period following the first defection.

Now, assume that players discount the payoffs at rate  $\delta < 1$ . Then, the strategy profile described will lead to  $(C, C)$  played in every period, so the discounted utility is

$$S = 1 + \delta + \delta^2 + \delta^3 + \dots$$

But

$$\delta S = \delta + \delta^2 + \delta^3 + \delta^4 + \dots$$

so

$$S - \delta S = 1 \Rightarrow S = \frac{1}{1 - \delta}.$$

We need to check that there is no incentive to deviate (in any period). Since the future looks the same at any point in the game with these strategies this can be done in a single calculation. If a player deviates at time  $t$ , then the payoff in the period is 2 rather than 1, but the downside is that the future is now worth 0 instead of  $\delta \frac{1}{1-\delta}$ . Hence, there is no reason to defect if

$$\frac{1}{1-\delta} \geq 2 \text{ or } \delta \geq \frac{1}{2}.$$

For subgame perfection we also have to check that  $(D, D)$  forever is an equilibrium in the continuation of the game since this is used after a defection, but this is obvious since it is repeating the static equilibrium in every period.

## 15.9 Three Player Games

We end this section by noting that we have a somewhat useful matrix representation also of games with three players. I will illustrate this with a voting problem with three voters (Axel, Birgitta and Carl), who are voting on a referendum. Assume that the referendum passes if at least two votes are cast in favor, and suppose that the voters only care about whether the referendum passes or not. Let  $a$  (which could be positive or negative) be the relative gain or loss from the referendum being passed for Axel,  $b$  be the gain/loss for Birgitta and  $c$  being the gain/loss for Carl. The strategy for each player is simply to vote “yes” or ”no”, and we may represent this game as in Figure 10

We note;

- For Axel and Birgitta, we can go on and consider the best responses just like before. I.e., Axel picks the row and Birgitta picks the column.
- The only difference is that Carl picks the matrix.
- Going through the best responses, we observe that (Yes, Yes, Yes) and (No, No, No) are both Nash equilibria, regardless of preferences.

		Carl votes Yes		Carl Votes No		
		Birgitta		Birgitta		
		Yes	No	Yes	No	
Axel	Yes	$a, b, c$	$a, b, c$	Yes	$a, b, c$	$0, 0, 0$
	No	$a, b, c$	$0, 0, 0$	No	$0, 0, 0$	$0, 0, 0$

Figure 10: A Three Player Voting Game

- There is also another potential equilibrium where each player votes in favor if and only if they prefer the referendum.