MINISTRY OF EDUCATION AND SCIENCE OF THE RUSSIAN FEDERATION
State autonomous educational institution of higher education
«Lobachevsky State University of Nizhni Novgorod»

Markina M.V.

MATRIX GAMES. PURE NASH EQUILLIBRIUM

Teaching aid
For students studying in English

Recommended by the Methodological Commission of the IITMM
for students of UNN, students in the field of training
"Fundamental informatics and information technology"

Nizhni Novgorod 2018

МИНИСТЕРСТВО ОБРАЗОВАНИЯ И НАУКИ
РОССИЙСКОЙ ФЕДЕРАЦИИ

Федеральное государственное автономное образовательное
учреждение высшего образования
«Национальный исследовательский Нижегородский
государственный университет им. Н.И. Лобачевского»

М.В. Маркина

МАТРИЧНЫЕ ИГРЫ. РАВНОВЕСИЕ ПО НЭШУ В ЧИСТЫХ СТРАТЕГИЯХ

Учебно-методическое пособие
Для студентов, обучающихся на английском языке

Рекомендовано методической комиссией ИИТММ
для студентов ННГУ, обучающихся по направлению подготовки
“ фундаментальная информатика и информационные технологии”

Нижний Новгород
2018

УДК 519.85
ББК В22.18

Маркина М.В: МАТРИЧНЫЕ ИГРЫ. РАВНОВЕСИЕ ПО НЭШУ В ЧИСТЫХ СТРАТЕГИЯХ**.** Методическое пособие. – [электронный ресурс] Нижний Новгород: Нижегородский госуниверситет, 2018. – 24 с.

Рецензент: к.ф.м.н., доцент **Круглов Е.В.**

Настоящее пособие содержит материалы по курсу «Исследование операций», читаемого бакалаврам направления подготовки «Фундаментальная информатика и информационные системы» ИИТММ. Рассматриваются разделы: построение математической модели стратегической ситуации, определение концепции равновесия по Нэшу, редукция матричной игры. В каждом разделе приведен необходимый минимум теоретических сведений, разобраны практические примеры.

. Пособие предназначено для студентов и магистров, специализирующихся в области теории математического моделирования, оптимизации и математического программирования.

УДК 519.85
ББК 22.18

Contents

1. Mathematical model of strategic situation

A STRATEGIC GAME is a model of interacting decision-makers. In recognition of the interaction, we refer to the decision-makers as players. Each player has a set of possible actions. The model captures interaction between the players by allowing each player to be affected by the actions of all players, not only her own action. Specifically, each player has preferences about the action profile—the list of all the players' actions.

More precisely, a strategic game is defined as follows.

_ DEFINITION Mathematical model of a strategic game (with ordinal preferences) consists of

• a set of players
• for each player, a set of actions
• for each player, preferences over the set of action profiles.

A very wide range of situations may be modeled as strategic games. For example, the players may be firms, the actions prices, and the preferences a reflection of the firms' profits. Or the players may be candidates for political office, the actions campaign expenditures, and the preferences a reflection of the candidates' probabilities of winning. Or the players may be animals fighting over some prey, the actions concession times, and the preferences a reflection of whether an animal wins or loses.

As in the model of rational choice by a single decision-maker, it is frequently convenient to specify the players' preferences by giving payoff functions that represent them. Bear in mind that these payoffs have only ordinal significance. If a player's payoffs to the action profiles a, b , and c are 1, 2, and 10, for example, the only conclusion we can draw is that the player prefers c to b and b to a ; the numbers do not imply that the player's preference between c and b is stronger than her preference between a and b .

Time is absent from the model. The idea is that each player chooses her action once and for all, and the players choose their actions "simultaneously" in the sense that no player is informed, when she chooses her action, of the action chosen by any other player. (For this reason, a strategic game is sometimes referred to as a "simultaneous move game".) Nevertheless, an action may involve activities that extend over time, and may take into account an unlimited number of contingencies. An action might specify, for example, "if company X 's stock falls below $10, buy 100 shares; otherwise, do not buy any shares". (For this reason, an action is sometimes called a "strategy".) However, the fact that time is absent from the model means that when analyzing a situation as a strategic game, we abstract from the complications that may arise if a player is allowed to change her plan as events unfold: we assume that actions are chosen once and for all.

Example 1: the Prisoner's Dilemma

One of the most well-known strategic games is the Prisoner's Dilemma. Its name comes from a story involving suspects in a crime; its importance comes from the huge variety of situations in which the participants face incentives similar to those faced by the suspects in the story.

Two suspects in a major crime are held in separate cells. There is enough evidence to convict each of them of a minor offense, but not enough evidence to convict either of them of the major crime unless one of them acts as an informer against the other (finks). If they both stay quiet, each will be convicted of the minor offense and spend one year in prison. If one and only one of them finks, she will be freed and used as a witness against the other, who will spend four years in prison. If they both fink, each will spend three years in prison.

This situation may be modeled as a strategic game:

Players: The two suspects.

Actions: Each player's set of actions is {Quiet, Fink}.

Preferences Suspect 1's ordering of the action profiles, from best to worst, is (Fink, Quiet) (she finks and suspect 2 remains quiet, so she is freed), (Quiet, Quiet) (she gets one year in prison), (Fink, Fink) (she gets three years in prison), (Quiet, Fink) (she gets four years in prison). Suspect 2's ordering is (Quiet, Fink), (Quiet, Quiet), (Fink, Fink), (Fink, Quiet).

We can represent the game compactly in a table. First choose payoff functions that represent the suspects' preference orderings. For suspect 1 we need a function u1 for which u_1(Fink, Quiet) > u_1(Quiet, Quiet) > u_1(Fink, Fink) > u_1(Quiet, Fink).

A simple specification is u_1(Fink, Quiet) = 3, u1(Quiet, Quiet) = 2, u_1(Fink, Fink) =1, and u_1(Quiet, Fink) = 0. For suspect 2 we can similarly choose the function u_2 for which u_2(Quiet, Fink) = 3, u_2(Quiet, Quiet) = 2, u_2(Fink, Fink) = 1, and u_2(Fink, Quiet) = 0. Using these representations, the game is illustrated in Figure 1. In this figure the two rows correspond to the two possible actions of player 1, the two columns correspond to the two possible actions of player 2, and the numbers in each box are the players' payoffs to the action profile to which the box corresponds, with player 1's payoff listed first.

		Suspect 1	
		Quiet	Fink
Suspect 2	Quiet	2, 2	0, 3
	Fink	3, 0	1, 1

Figure 1. The Prisoner's Dilemma

The Prisoner's Dilemma models a situation in which there are gains from cooperation (each player prefers that both players choose Quiet than they both choose Fink) but each player has an incentive to "free ride" (choose Fink) whatever the other player does. The game is important not because we are interested in understanding the incentives for prisoners to confess, but because many other situations have similar structures. Whenever each of two players has two actions, say C (corresponding to Quiet) and D (corresponding to Fink), player 1 prefers (D, C) to (C, C) to (D, D) to (C, D), and player 2 prefers (C, D) to (C, C) to (D, D) to (D, C), the Prisoner's Dilemma models the situation that the players face. Some examples follow.

Example 2: Working on a joint project

You are working with a friend on a joint project. Each of you can either work hard or goof off. If your friend works hard then you prefer to goof off (the outcome of the project would be better if you worked hard too, but the increment in its value to you is not worth the extra effort). You prefer the outcome of your both working hard to the outcome of your both goofing off (in which case nothing gets accomplished), and the worst outcome for you is that you work hard and your friend goofs off (you hate to be "exploited"). If your friend has the same preferences then the game that models the situation you face is given in Figure 2, which, as you can see, differs from the Prisoner's Dilemma only in the names of the actions.

	Work hard	Goof off
Work hard	2, 2	0, 3
Goof off	3, 0	1, 1

Figure 2 Working on a joint project

Example 3: Duopoly

In a simple model of a duopoly, two firms produce the same good, for which each firm charges either a low price or a high price. Each firm wants to achieve the highest possible profit. If both firms choose High then each earns a profit of \$1000.
If one firm chooses High and the other chooses Low then the firm choosing High obtains no customers and makes a loss of \$200, whereas the firm choosing Low earns a profit of \$1200 (its unit profit is low, but its volume is high). If both firms choose Low then each earns a profit of \$600. Each firm cares only about its profit, so we can represent its preferences by the profit it obtains, yielding the game in Figure 3.

	High	Low
High	1000, 1000	−200, 1200
Low	1200, −200	600, 600

Figure 3 A simple model of a price-setting duopoly

Bearing in mind that what matters are the players' preferences, not the particular payoff functions that we use to represent them, we see that this game, like the previous one, differs from the Prisoner's Dilemma only in the names of the actions.
The action High plays the role of Quiet, and the action Low plays the role of Fink; firm1 prefers (Low, High) to (High, High) to (Low, Low) to (High, Low) , and firm2 prefers (High, Low) to (High, High) to (Low, Low) to (Low, High) .
As in the previous example, I do not claim that the incentives in a duopoly are necessarily those in the Prisoner's Dilemma; different assumptions about the relative sizes of the profits in the four cases generate a different game. Further, in this case one of the abstractions incorporated into the model—that each firm has only two prices to choose between—may not be harmless; if the firms may choose among many prices then the structure of the interaction may change.

<u>Example 4: Bach or Stravinsky?</u>

Two people wish to go out together. Two concerts are available: one of music by Bach, and one of music by Stravinsky. One person prefers Bach and the other prefers Stravinsky. If they go to different concerts, each of them is equally unhappy listening to the music of either composer.
We can model this situation as the two-player strategic game in Figure 4, in which the person who prefers Bach chooses a row and the person who prefers Stravinsky chooses a column.

	Bach	Stravinsky
Bach	2, 1	0, 0
Stravinsky	0, 0	1, 2

Figure 4 Bach or Stravinsky

This game is also referred to as the "Battle of the Sexes" (though the conflict it models surely occurs no more frequently between people of the opposite sex than it does between people of the same sex). I refer to the games as BoS, an acronym that fits both names. (I assume that each player is indifferent between listening to Bach and listening to Stravinsky when she is alone only for consistency with the standard specification of the game. As we shall see, the analysis of the game remains the same in the absence of this assumption.)

Like the Prisoner's Dilemma, BoS models a wide variety of situations. Consider, for example, two officials of a political party deciding the stand to take on an issue.

Suppose that they disagree about the best stand, but are both better off if they take the same stand than if they take different stands; both cases in which they take different stands, in which case voters do not know what to think, are equally bad.

Then BoS captures the situation they face. Or consider two merging firms that currently use different computer technologies. As two divisions of a single firm they will both be better off if they both use the same technology; each firm prefers that the common technology be the one it used in the past. BoS models the choices the firms face.

<u>Example 5: Matching Pennies</u>

Two people choose, simultaneously, whether to show the Head or the Tail of a coin. If they show the same side, person 2 pays person 1 a dollar; if they show different sides, person 1 pays person 2 a dollar. Each person cares only about the amount of money she receives, and (naturally!) prefers to receive more than less. A strategic game that models this situation is shown in Figure 5. (In this representation of the players' preferences, the payoffs are equal to the amounts of money involved. We could equally well work with another representation—for example, 2 could replace each 1, and 1 could replace each − 1.)

	Head	Tail
Head	1, − 1	− 1, 1
Tail	− 1, 1	1, − 1

Figure 5 Matching Pennies

In this game the players' interests are diametrically opposed (such a game is called "strictly competitive"): player 1 wants to take the same action as the other player, whereas player 2 wants to take the opposite action. This game may, for example, model the choices of appearances for new products by an established producer and a new firm in a market of fixed size. Suppose that each firm can choose one of two different appearances for the product. The established producer prefers the newcomer's product to look different from its own (so that its customers will not be tempted to buy the newcomer's product), whereas the newcomer prefers that the products look alike. Or the game could model a relationship between two people in which one person wants to be like the other, whereas the other wants to be different.

2. Nash equilibrium

What actions will be chosen by the players in a strategic game? We wish to assume, as in the theory of a rational decision-maker, that each player chooses the best available action. In a game, the best action for any given player depends, in general, on the other players' actions. So when choosing an action a player must have in mind the actions the other players will choose. That is, she must form a belief about the other players' actions.

On what basis can such a belief be formed? The assumption underlying the analysis in this chapter and the next two chapters is that each player's belief is derived from her past experience playing the game, and that this experience is sufficiently extensive that she knows how her opponents will behave. No one tells her the actions her opponents will choose, but her previous involvement in the game leads her to be sure of these actions.

Although we assume that each player has experience playing the game, we assume that she views each play of the game in isolation. She does not become familiar with the behavior of specific opponents and consequently does not condition her action on the opponent she faces; nor does she expect her current action to affect the other players' future behavior.

It is helpful to think of the following idealized circumstances. For each player in the game there is a population of many decision-makers who may, on any occasion, take that player's role. In each play of the game, players are selected randomly, one from each population. Thus each player engages in the game repeatedly, against ever-varying opponents. Her experience leads her to beliefs about the actions of "typical" opponents, not any specific set of opponents.

As an example, think of the interaction between buyers and sellers.

Buyers and sellers repeatedly interact, but to a first approximation many of the pairings may be modeled as random. In many cases a buyer transacts only once with any given seller, or interacts repeatedly but anonymously (when the seller is a large store, for example).
In summary, the solution theory we study has two components. First, each player chooses her action according to the model of rational choice, given her belief about the other players' actions. Second, every player's belief about the other players' actions is correct. These two components are embodied in the following definition.
A Nash equilibrium is an action profile $\mathbf{a*}$ with the property that no player $\mathbf{i}$ can do better by choosing an action different from $\mathbf{a*}_i$, given that every other player j adheres to $\mathbf{a*}_j$.
In the idealized setting in which the players in any given play of the game are drawn randomly from a collection of populations, a Nash equilibrium corresponds to a steady state. If, whenever the game is played, the action profile is the same Nash equilibrium $\mathbf{a*}$, then no player has a reason to choose any action different from her component of $\mathbf{a*}$; there is no pressure on the action profile to change. Expressed differently, a Nash equilibrium embodies a stable "social norm": if everyone else adheres to it, no individual wishes to deviate from it.
The second component of the theory of Nash equilibrium—that the players' beliefs about each other's actions are correct—implies, in particular, that two players' beliefs about a third player's action are the same. For this reason, the condition is sometimes said to be that the players' "expectations are coordinated".
The situations to which we wish to apply the theory of Nash equilibrium do not in general correspond exactly to the idealized setting described above. For example, in some cases the players do not have much experience with the game; in others they do not view each play of the game in isolation. Whether or not the notion of Nash equilibrium is appropriate in any given situation is a matter of judgment. In some cases, a poor fit with the idealized setting may be mitigated by other considerations. For example, inexperienced players may be able to draw conclusions about their opponents' likely actions from their experience in other situations, or from other sources. Ultimately, the test of the appropriateness of the notion of Nash equilibrium is whether it gives us insights into the problem at hand. With the aid of an additional piece of notation, we can state the definition of a Nash equilibrium precisely. Let $\mathbf{a}$ be an action profile, in which the action of each player $\mathbf{i}$ is $\mathbf{a}_i$. Let $\mathbf{a'}_i$ be any action of player i (either equal to a_i, or different from it). Then $(\mathbf{a'}_i, \mathbf{a}_{-i})$ denotes the action profile in which every player j except i chooses her action $\mathbf{a}_j$ as specified by $\mathbf{a}$, where as player i chooses $\mathbf{a'}_i$. (The –i subscript on $\mathbf{a}$ stands for "except i ".) That is, $(\mathbf{a'}_i, \mathbf{a}_{-i})$ is the

action profile in which all the players other than i adhere to a while i "deviates" to $\mathbf{a'_i}$. (If $\mathbf{a'_i} = \mathbf{a_i}$ then of course $(\mathbf{a'_i}, \mathbf{a_{-i}}) = (\mathbf{a_i}, \mathbf{a_{-i}}) = \mathbf{a}$.) If there are three players, for example, then $(\mathbf{a'_2}, \mathbf{a_{-2}})$ is the action profile in which players 1 and 3 adhere to **a** (player 1 chooses $\mathbf{a_1}$, player 3 chooses $\mathbf{a_3}$) and player 2 deviates to $\mathbf{a'_2}$.
Using this notation, we can restate the condition for an action profile $a\star$ to be a Nash equilibrium: no player i has any action a_i for which she prefers $(\mathbf{a_i}, \mathbf{a\star_{-i}})$ to $a\star$.
Equivalently, for every player i and every action $\mathbf{a_i}$ of player i , the action profile $a\star$ is at least as good for player i as the action profile $(\mathbf{a_i}, \mathbf{a\star_{-i}})$.

DEFINITION (Nash equilibrium of strategic game with ordinal preferences)

The action profile $a\star$ in a strategic game with ordinal preferences is a Nash equilibrium if, for every player **i** and every action $\mathbf{a_i}$ of player **i**, $\mathbf{a\star}$ is at least as good according to player i 's preferences as the action profile $(\mathbf{a_i}, \mathbf{a\star_{-i}})$ in which player i chooses $\mathbf{a_i}$ while every other player j chooses $a\star_j$. Equivalently, for every player i ,

$$u_i(a\star) \geq u_i(\mathbf{a_i}, \mathbf{a\star_{-i}}) \text{ for every action } \mathbf{a_i} \text{ of player i }, \qquad (\mathbf{a})$$

where $\mathbf{u_i}$ is a payoff function that represents player i 's preferences. This definition implies neither that a strategic game necessarily has a Nash equilibrium, nor that it has at most one. Examples in the next section show that some games have a single Nash equilibrium, some possess no Nash equilibrium, and others have many Nash equilibria.
The definition of a Nash equilibrium is designed to model a steady state among experienced players. An alternative approach to understanding players' actions in strategic games assumes that the players know each others' preferences, and considers what each player can deduce about the other players' actions from their rationality and their knowledge of each other's rationality.

Examples of Nash equilibrium

By examining the four possible pairs of actions in the Prisoner's Dilemma we see that (Fink, Fink) is the unique Nash equilibrium.
The action pair (Fink, Fink) is a Nash equilibrium because (i) given that player 2 chooses Fink , player 1 is better off choosing Fink than Quiet (looking at the right column of the table we see that Fink yields player 1 a payoff of 1 whereas Quiet yields her a payoff of 0), and (ii) given that player 1 chooses Fink , player 2 is better off choosing Fink than Quiet (looking at the bottom row of the table we see that Fink yields player 2 a payoff of 1 whereas Quiet yields her a payoff of 0).
No other action profile is a Nash equilibrium:

• (Quiet, Quiet) does not satisfy (**a**) because when player 2 chooses Quiet , player 1's payoff to Fink exceeds her payoff to Quiet (look at the first components of the entries in the left column of the table). (Further, when player 1 chooses Quiet, player 2's payoff to Fink exceeds her payoff to Quiet : player 2, as well as player 1, wants to deviate. To show that a pair of actions is not a Nash equilibrium, however, it is not necessary to study player 2's decision once we have established that player 1 wants to deviate: it is enough to show that one player wishes to deviate to show that a pair of actions is not a Nash equilibrium.)

• (Fink Quiet) does not satisfy (**a**) because when player 1 chooses Fink , player 2's payoff to Fink exceeds her payoff to Quiet (look at the second components of the entries in the bottom row of the table).

• (Quiet, Fink) does not satisfy (**a**) because when player 2 chooses Fink , player 1's payoff to Fink exceeds her payoff to Quiet (look at the first components of the entries in the right column of the table).

<u>Strict and nonstrict equilibria</u>

In all the Nash equilibria of the games we have studied so far a deviation by a player leads to an outcome worse for that player than the equilibrium outcome.

The definition of Nash equilibrium, however, requires only that the outcome of a deviation be no better for the deviant than the equilibrium outcome. And, indeed, some games have equilibria in which a player is indifferent between her equilibrium action and some other action, given the other players' actions.

Consider the game in Figure 6. This game has a unique Nash equilibrium, namely (T, L). (For every other pair of actions, one of the players is better off changing her action.) When player 2 chooses L, as she does in this equilibrium, player 1 is equally happy choosing T or B; if she deviates to B then she is no worse off than she is in the equilibrium. We say that the Nash equilibrium (T, L) is not a strict equilibrium.

	L	M	R
T	1 1,	1,0,	0, 1
B	1,0,	0 1,	1 0,

Figure 6 A game with a unique Nash equilibrium, which is not a strict equilibrium

For a general game, an equilibrium is strict if each player's equilibrium action is better than all her other actions, given the other players' actions. Precisely, an action profile **a∗** is a strict Nash equilibrium if for every player i we have $\mathbf{u_i(a*) > u_i(a_i, a*_{-i})}$ for every action $\mathbf{a_i = a*_i}$ of player **i.** (Contrast the strict inequality in this definition with the weak inequality in (**a**).)

3. Best response functions

We can find the Nash equilibria of a game in which each player has only a few actions by examining each action profile in turn to see if it satisfies the conditions for equilibrium. In more complicated games, it is often better to work with the players' "best response functions".

Consider a player, say player i. For any given actions of the players other than i, player i's actions yield her various payoffs. We are interested in the best actions— those that yield her the highest payoff. For example, Bach is the best action for player 1 if player 2 chooses Bach; Stravinsky is the best action for player 1 if player 2 chooses Stravinsky. In particular, in BoS, player 1 has a single best action for each action of player 2. By contrast, in the game in Figure 31.1, both T and B are best actions for player 1 if player 2 chooses L: they both yield the payoff of 1, and player 1 has no action that yields a higher payoff (in fact, she has no other action)

We denote the set of player i 's best actions when the list of the other players' actions is a_{-i} by $Bi(a_{-i})$. Thus in BoS we have B_1(Bach) = {Bach} and B_1(Stravinsky) = {Stravinsky} ; in the game in Figure 6 we have $B_1(L) = \{T , B\}$.

Precisely, we define the function Bi by

$$\mathbf{B_i(a_{-i}) = \{a_i \text{ in } Ai : u_i(a_i , a_{-i}) \geq u_i(a'_i , a_{-i}) \text{ for all } a'_i \text{ in } A_i\}} :$$

any action in $\mathbf{B_i(a_{-i})}$ is at least as good for player **i** as every other action of player i when the other players' actions are given by $\mathbf{a_{-i}}$. We call $\mathbf{B_i}$ the best response function of player **i** .

The function B_i is set-valued : it associates a set of actions with any list of the other players' actions. Every member of the set $B_i(\mathbf{a_{-i}})$ is a best response of player **i** to $\mathbf{a_{-i}}$: if each of the other players adheres to $\mathbf{a_{-i}}$ then player **i** can do no better than choose a member of $\mathbf{B_i(a_{-i})}$. In some games, the set $\mathbf{B_i(a-i)}$ consists of a single action for every list $\mathbf{a_{-i}}$ of actions of the other players: no matter what the other players do, player **i** has a single optimal action. In other games, $\mathbf{B_i(a_{-i})}$ can contains more than one action for some lists $\mathbf{a_{-i}}$ of actions of the other players.

Using best response functions to define Nash equilibrium

A Nash equilibrium is an action profile with the property that no

player can do better by changing her action, given the other players' actions. Using the terminology just developed, we can alternatively define a Nash equilibrium to be an action profile for which every player's action is a best response to the other players' actions.
That is, we have the following result.
PROPOSITION
The action profile **a⋆** is a Nash equilibrium of a strategic game with ordinal preferences if and only if every player's action is a best response to the other players' actions:

$$a\star_i \text{ is in } Bi(a\star_{-i}) \text{ for every player } i \,. \qquad (b)$$

If each player i has a single best response to each list $\mathbf{a_{-i}}$ of the other players' actions, we can write the conditions in (b) as equations. In this case, for each player **i** and each list a_{-i} of the other players' actions, denote the single member of $\mathbf{B_i(a_{-i})}$ by $\mathbf{b_i(a_{-i})}$ (that is, $\mathbf{B_i(a_{-i}) = \{b_i(a_{-i})\}}$).
Then (b) is equivalent to

$$\mathbf{a\star_i = b_i(a\star_{-i})} \text{ for every player } i \,, \qquad (c)$$

a collection of n equations in the n unknowns $a\star_i$, where n is the number of players in the game. For example, in a game with two players, say 1 and 2, these equations are

$$a\star_1 = b_1(a\star_2) \qquad a\star_2 = b_2(a\star_1) \,.$$

That is, in a two-player game in which each player has a single best response to every action of the other player, $(a\star_1, a\star_2)$ is a Nash equilibrium if and only if player 1'saction $\mathbf{a\star_1}$ is her best response to player 2's action $\mathbf{a\star_2}$, and player 2's action $\mathbf{a\star_2}$ is her best response to player 1's action $\mathbf{a\star_1}$.

The definition of a Nash equilibrium in terms of best response functions suggests a method for finding Nash equilibria:
- find the best response function of each player
- find the action profiles that satisfy (b) (which reduces to (c) if each player has a single best response to each list of the other players' actions).

To illustrate this method, consider the game in Figure 7. First find the best response of player 1 to each action of player 2. If player 2 chooses L, then player 1's best response is M (2 is the highest payoff for player 1 in this column); indicate the best response by attaching a star to player 1's payoff to (M, L) . If player 2 chooses C, then player 1's best response is T, indicated by the star attached to player 1's payoff to (T, C). And if player 2 chooses R, then both T and B are best responses for player 1; both are indicated by stars. Second, find the best response of player 2 to

each action of player 1 (for each row, find highest payoff of player 2); these best responses are indicated by attaching stars to player 2's payoffs. Finally, find the boxes in which both players' payoffs are starred. Each such box is a Nash equilibrium: the star on player 1's payoff means that player 1's action is a best response to player 2's action, and the star on player 2's payoff means that player 2's action is a best response to player 1's action. Thus we conclude that the game has two Nash equilibria:
(M, L) and (B, R .

	L	C	R
T	1 ,2*	2* , 1	1* , 0
M	2* , 1*	0 ,1*	0 ,0
B	0 ,1	0 ,0	1* , 2*

Figure 7 Using best response functions to find Nash equilibria in a two-player game in which each player has three actions.

4. Dominated actions

Strict domination

You drive up to a red traffic light. The left lane is free; in the right lane there is a car that may turn right when the light changes to green, in which case it will have to wait for a pedestrian to cross the side street. Assuming you wish to progress as quickly as possible, the action of pulling up in the left lane "strictly dominates" that of pulling up in the right lane. If the car in the right lane turns right then you are much better off in the left lane, where your progress will not be impeded; and even if the car in the right lane does not turn right, you are still better off in the left lane, rather than behind the other car.
In any game, a player's action "strictly dominates" another action if it is superior, no matter what the other players do.
DEFINITION (Strict domination)
In a strategic game with ordinal preferences, player i's action $\mathbf{a''_i}$ strictly dominates her action $\mathbf{a'_i}$ if $\mathbf{u_i(a''_i , a_{-i}) > u_i(a'i , a_{-i})}$ for every list $\mathbf{a_{-i}}$ of the other players' actions, where $\mathbf{u_i}$ is a payoff function that represents player i's preferences.
In the Prisoner's Dilemma, for example, the action Fink strictly dominates the action Quiet: regardless of her opponent's action, a player prefers the outcome when she chooses Fink to the outcome when she chooses Quiet . In BoS , on the other hand, neither action strictly dominates the other: Bach is better than Stravinsky if the other player chooses Bach , but is worse than Stravinsky if the other player chooses

Stravinsky. If an action strictly dominates the action $\mathbf{a_i}$, we say that $\mathbf{a_i}$ is strictly dominated . A strictly dominated action is not a best response to any actions of the other players: whatever the other players do, some other action is better. Since a player's Nash equilibrium action is a best response to the other players' Nash equilibrium actions, a strictly dominated action is not used in any Nash equilibrium.

When looking for the Nash equilibria of a game, we can thus eliminate from consideration all strictly dominated actions. For example we can eliminate Quiet for each player in the Prisoner's Dilemma , leaving (Fink, Fink) as the only candidate for a Nash equilibrium. (As we know, this action pair is indeed a Nash equilibrium.) The fact that the action $\mathbf{a''_i}$ strictly dominates the action $\mathbf{a'_i}$ of course does not imply that $\mathbf{a''}_i$ strictly dominates all actions. Indeed, $\mathbf{a''}_i$ may itself be strictly dominated.

In the left-hand game in Figure 8, for example, M strictly dominates T, but B is better than M if player 2 chooses R. (I give only the payoffs of player 1 in the figure, because those of player 2 are not relevant.) Since T is strictly dominated, the game has no Nash equilibrium in which player 1 uses it; but the game may also not have any equilibrium in which player 1 uses M . In the right-hand game, M strictly dominates T but itself strictly dominated by B. In this case, in any Nash equilibrium player 1's action is B (her only action that is not strictly dominated).

	L	R
T	1,1	0,0
M	2,2	1,1
B	1,3	3,2

Figure 8 Two games in which player 1's action T is strictly dominated by M

In the left-hand game, B is better than M if player 2 chooses R; in the right-hand game, M itself is strictly dominated, by B.

A strictly dominated action is incompatible not only with a steady state, but also with rational behavior by a player who confronts a game for the first time.

Weak domination

As you approach the red light in the situation at the start of the previous section, there is a car in each lane. The car in the right lane may, or may not, be turning right; if it is, it may be delayed by a pedestrian crossing the side street. The car in the left lane cannot turn right. In this case your pulling up in the left lane "weakly dominates", though does not strictly dominate, your pulling up in the right lane.

If the car in the right lane does not turn right, then both lanes are equally

good; if it does, then the left lane is better.
In any game, a player's action "weakly dominates" another action if the first action is at least as good as the second action, no matter what the other players do, and is better than the second action for some actions of the other players.
DEFINITION (Weak domination)
In a strategic game with ordinal preferences, player i 's action $\mathbf{a''_i}$ weakly dominates her action $\mathbf{a'_i}$ if $\mathbf{u_i(a''_i , a_{-i}) \geq u_i(a'_i , a_{-i})}$ for every list $\mathbf{a_{-i}}$ of the other players' actions and $\mathbf{u_i(a''_i , a_{-i}) > u_i(a'_i , a_{-i})}$ for some list $\mathbf{a_{-i}}$ of the other players' actions, where $\mathbf{u_i}$ is a payoff function that represents player i 's preferences.
For example, in the game in Figure 9 (in which, once again, only player 1's payoffs are given), M weakly dominates T , and B weakly dominates M ; B strictly dominates T .

	L	R
T	1	0
M	2	0
B	2	1

Figure 9. A game illustrating weak domination. (Only player 1's payoffs are given.)

The action M weakly dominates T; B weakly dominates M. The action B strictly dominates T equilibrium action.
Can an action be weakly dominated in a nonstrict Nash equilibrium? Definitely.

5. **Zero-sum game**

A situation in which one person's gain is equivalent to another's loss, so the net change in wealth or benefit is zero.
Zero-sum game: a game in which the sum of all players' payoffs equals zero for every outcome. So essentially you have a game that game can produce some number of outcomes and if we are looking at a zero-sum game then it must be the case that you take any outcome from the game and you sum all of the players pay off together and that some must equal exactly zero.
Minmax (maxmin) approach
Minmax is a decision rule used in game theory for *mini*mizing the possible loss or for maximizing the possible win for a worst case scenario.

	Y1	y2
X1	3	5
X2	1	-2

What will happen with Y if Y choose y1?
If X chooses x1 Y can lose 3 .
If X chooses x2 Y can lose 1.
So maximum possible loss of player Y is going to be 3.
Similarly maximum possible lost for Y if Y chooses y2 is 5.
Y tries to minimize the possible losses. The idea is very simple. So Y must choose y1
The criteria being used by Y in this situation is minimax criteria which is minimizing the maximum lost. This 3 is called UPON VALUE of the game .
Similarly if you consider strategies used by the player X. X tries to maximize possible win.
In first row minimum possible win is 3.
In second row is -2
So X must choose x1. The criteria being used by X in this situation is maxmin criteria which is maximizing the minimum win. This 3 is called LOWER VALUE of the game .
In this game
UPPER=LOWER=VALUE=3
This 3 is average expected game outcome of this game if play many times. So both players X and Y will use strategies in such a way that they end up with this value 3.
This solution is SADDLE POINT for zero-sum game and is Pure Nash Equilibrium.

Definition of saddle point

A saddle point of a matrix is the position of such an element in the payoff matrix, which is minimum in its row and the maximum in its column.

Procedure to find the saddle point

- • Select the minimum element of each row of the payoff matrix and mark them with circles.
- • Select the maximum element of each column of the payoff matrix and mark them with squares.
- • If their appears an element in the payoff matrix with a circle and a square together then that position is called saddle point and the element is the value of the game.

Solution of games with saddle point

To obtain a solution of a game with a saddle point, it is feasible to find out

- Best strategy for player A
- Best strategy for player B
- The value of the game

JOHN F. NASH , JR .
A few of the ideas of John F. Nash Jr., developed while he was a graduate student at Princeton from1948 to 1950, transformed game theory. Nash was born in 1928 in Bluefield, West Virginia, USA, where he grew up. He was an undergraduate mathematics major at Carnegie Institute of Technology from 1945 to 1948. In 1948 he obtained both a B.S. and an M.S., and began graduate work in the Department of Mathematics at Princeton University. (One of his letters of recommendation, from a professor at Carnegie Institute of Technology, was a single sentence: "This man is a genius" (Kuhn et al. 1995, 282).) A paper containing the main result of his thesis was submitted to the Proceedings of the National Academy of Sciences in November 1949, fourteen months after he started his graduate work. ("A fine goal to set . . .graduate students", to quote Kuhn! (See Kuhn et al. 1995, 282.)) He completed his PhD the following year, graduating on his 22nd birthday. His thesis, 28 pages in length, introduces the equilibrium notion now known as "Nash equilibrium" and delineates a class of strategic games that have Nash equilibrium (Proposition 116.1 in this book). The notion of Nash equilibrium vastly expanded the scope of game theory, which had previously focused on two-player "strictly competitive" games (in which the players' interests are directly opposed). While a graduate student at Princeton, Nash also wrote the seminal paper in bargaining theory, Nash (1950b)
(the ideas of which originated in an elective class in international economics he took as an undergraduate). He went on to take an academic position in the Department (Milnor 1995, 15); he has been described as "one of the most original mathematical minds of [the twentieth] century" (Kuhn 1996). He shared the 1994 Nobel prize in economics with the game theorists John C. Harsanyi and Reinhard Selten.

6. EXERCISES

1. Formulate this situation as a strategic game and find its Nash equilibrium

Task 1.1
Players I and II simultaneously call out one of the numbers one or two. Player I's name is Odd; he wins if the sum of the numbers is odd. Player II's name is Even; he wins if the sum of the numbers is even.

Task 1.2
Both players simultaneously choose an integer from 0 to 3 and they both

win the smaller of the two numbers in points. In addition, if one player chooses a larger number than the other, then he have to give up two points to the other. For example: if Player1 chooses '1' and Player 2 chooses '2', we have outcome (3,-1).

Task 1.3

The previous game (Task 2) is modified so that the two players win the named amount if they both choose the same number, and otherwise win nothing.

Task 1.4

There are two bars in a city, with owners labelled *A* and *B*, who can charge $2, $4, or $5 per drink. Each day, there are 6,000 tourists and 4,000 locals who decide which bar to visit. (Each person can only go to one bar and each must go to at least one bar, where each person has exactly one drink.) Since tourists have no idea about the bars, they randomize without reference to the pricing. The locals, however, always go to the cheapest bar (and randomize if the prices are the same). The question is: what prices should the owners set if they choose simultaneously? (Use iterated elimination of dominated strategies)

Examples of the outcomes:

f1 (2,2)= (3000+2000)*2=10000

f1(4,2)=3000*4=12000

f1(2,4)=(3000+4000)*2=14000

f2(2,4)=3000*4=12000

Task 1.5

Let there are 6 possible trade points evenly spaced along the street. They are equidistance from each other. The first point is at the beginning of the block and the latest point is at the end of the block.
Both vendors can choose any trade point for the sale of ice cream. What is the winning strategy for one of these two vendors? ? (Use iterated elimination of dominated strategies)

Task 1.6

Both players simultaneously choose an integer from 0 to 3 and they both win the smaller of the two numbers in points. In addition, if one player chooses a larger number than the other, then they have to give up two points to the other.

Task 1.7

Each of the two firms, loses $ 2 million per period if they both sell Internet browsers. When a firm does not have a rival, then, becoming a monopolist, it will earn $ 10 per period. Firms can withdraw from the market in 2016 (with income 0) and in 2017 -2018 (having lost 2 million for each previous year), or remain until the end of 2018.

Task 1.8

Two firms lease adjacent land over a 100 million-tonne oil tank. The cost of one ton is $ 200. Each of the firms should decide whether to drill a well and, if drilled, what size? Drilling and servicing a narrower well costs $ 100 million, and a broad $ 300 million. But at the same time, three times more oil will be pumped through a wide well a day.

2. Use iterated elimination of dominated strategies and find the Nash Equalibrium in pure strategies

2.1

	x1	x2	x3	x4	x5
y1	28;1	63;1	2;0	2;45	3;19
y2	2;2	32;1	2;5	33;0	2;3
y3	95;1	54;1	0;2	4;1	0;4
y4	3;43	1;33	1;39	1;12	1;17
y5	1;13	22;0	1;88	2;57	3;72
y6	0;12	15;0	0;0	0;99	1;88

2.2

	x1	x2	x3
y1	1,1	-2,0	4,-1
y2	0,3	3,1	5,4
y3	1,5	4,2	6,2

2.3

	x1	x2	x3
y1	1;23	2;55	0;33
y2	22;0	1;13	1;88
y3	0;2	5;6	2;14

2.4

	x1	x2	x3
y1	2;6	33;0	2;3
y2	0;2	4;1	0;4
y3	1;39	1;12	1;17

3. Find LOWER and UPPER VALUE of the game and check about Pure Nash Equilibrium in zero-sum game

3.1

12	22	65	67
33	11	90	12
88	25	67	28

3.2

12	11	11	13
55	10	11	15
44	12	11	16

3.3

10	11	12	13
14	15	16	17
18	19	20	21

References:

1. Martin Osborne “Gamebook”
http://pioneer.netserv.chula.ac.th/~ptanapo1/gamebook.pdf
2. Thomas S. Ferguson “Game Theory”
http://www.math.ucla.edu/~tom/Game_Theory/Contents.html