MAS348 brief notes
(material shown on slides during the lectures in compact, printer-friendly way.)
This is not a textbook– annotate these with the content of lectures.

Moty Katzman

Contents

Chapter 1. Inroduction to the course 7
What is a Game 7
Example: Prisoners' Dilemma 7
Utility and preferences 8
An Example 8

Chapter 2. Cooperative games– pure strategies 9
Games in strategic form 9
The strategic form of a Prisoners' Dilemma 9
Dominant strategies 9
Example: iterative elimination of dominated strategies 9
Example: The Median Voter Theorem 10
Weakly dominant strategies 11
Example: Sealed-bid second-price auctions 11
Pareto optimality 12
Stabiliy. Example: the investment game 12
Stability: best responses and Nash equilibria 12
Example: no stable strategy profile 13
Best response 13
Nash equilibrium 13
Example: A coordination game with multiple Nash equilibria 13
Example: Best response functions 13
Elimination of dominated strategies does not loose Nash equilibria 14
Example: weak domination and Nash equilibria 14
Weird Example: Braess's Paradox 14

Chapter 3. Nash equilibria in Economics: monopolies, duopolies and oligopolies 15
The Cournot model 15
The Bertrand model 16

Chapter 4. Cooperative games– mixed strategies 19
Example: matching pennies 19
Expected utility 19
Mixed strategies and mixed-strategy Nash equilibria 19
Example: matching pennies again 20
Example: The battle of the sexes 20
Example: Rock-Scissors-Paper 21
Example: Alice and Bob play tennis 21

Example: The attrition game 21
Dominating mixed strategies 22
Existence of Nash equilibria 22

Chapter 5. Two-person zero-sum games 23
“Maxminization” 23
“Maxminimizing” with mixed strategies 23
Zero-sum games 23
Pure strategies: looking for saddle points 23
Another example 24
minmax $\geq$ maxmin 25
Mixed strategies 26
Two-person zero-sum games have mixed-strategy Nash-equilibria! 27
Zero sum games in matrix form 28
Example: symmetric games 28
Example: a symmetric game 28
Example: finding the value of a game 29
Example: finding the value of a game 29
A very different example: A duel 30

Chapter 6. Sequential games 31
An example 31
Example: Solving games using backward induction 31
Directed rooted trees 31
Example: “The hungry lions” game 32
Cournot duopoly revisited: the Stackelberg model 32
Solving strictly competitive games with no chance moves 32
Subgames and subtrees 33
Example: 2-pile Nim 33
Zermelo’s theorem 33
Strictly competitive, finite, sequential games have a value 33
Chess-like games 34
Chomp! 34
Sequential games in strategic form 34
Example 35
Example: A bargaining process 35
Monopolist vs. new entrant 35
Imperfect information and information sets 36
Example: Alice and Bob compete, Part I 36
Example: Alice and Bob compete, Part II 37
An example 38
Subgame perfect Nash equilibria 39
Example: Market monopolist and new entrant, revisited. 39
Example 40

Chapter 7. Repeated Games 41

Example: Repeated Prisoners' Dilemma 41
Payoffs of infinitely repeated games 41
A finitely repeated game 41
Nash equilibria of repeated games 42
GRIM vs. GRIM 42
The "folk theorem"– first version (indefinite games) 43
Notation: G^{∞} 44
Minimax values 44
Convex sets 44
The cooperative payoff region of games 44
The "folk theorem"– second version (infinite games) 45

Chapter 8. Bayesian Games 47
Example: Alice and Bob face a showdown 47
Strategies in Bayesian Games 48
Bayes-Nash equilibrium 48
Example: Alice and Bob have a dispute 48
Example: More information may hurt 49

CHAPTER 1

Inroduction to the course

What is a Game

A game consists of

(a) a set of players, each having
(b) a set of actions or strategies; a choice of strategies by all players determines
(c) outcomes (or payoffs) for each player.

Example.

- Rock-paper-scissors
- Sealed-bid auctions (many players, partial information)
- M.A.D. (Mutually Assured Destruction) (scary game from my childhood)
- Duopolies, Oligopolies,
- Chess (sequential)
- Chomp!
- backgammon (partial information, potentially infinite)

Example: Prisoners' Dilemma

Alice and Bob are arrested for murder and theft, but while there is enough evidence to convict them of theft, there is not enough evidence to convict them of murder, unless one or both prisoners confess. Alice and Bob are both told:

> "If you confess and your partner does not, you go free and your partner will be jailed 20 years. If neither of you confess, you will go to jail for a year. If both of you confess, you will both go to jail for 10 years."

What's a prisoner to do?

		Bob	
		Confess	Don't Confess
Alice	Confess	Alice: 10 years, Bob: 10 years	Alice: free, Bob: 20 years
	Don't Confess	Alice: 20 years, Bob: free	Alice: 1 year, Bob: 1 year

Preferences:
going free is best,
1 year in jail preferable to 10 years,
10 years in jail preferable to 20 years,
(1 year in jail preferable to 20 years.)

Optimal strategy: *Confess*.

Don't Confess is dominated by *Confess*, i.e., no matter what choice the other player makes, *Confess* does worse than *Don't Confess*!

Utility and preferences

We want to compare attractiveness of outcomes of games.
Utility: a real number which is the measure of attractiveness of an outcome.
High utility preferable to low utility.
Game players aim to maximize utility (or expected utility later in the course).

An Example

Alice's utilities of 0, 1, 10 and 20 years in jail are 100, 3, 0, -1. Bob's utilities of 0, 1, 10 and 20 years in jail are 10, -5, -10, -20.

	Confess	Don't Confess
Confess	(0,-10)	(100,-20)
Don't Confess	(-1,10)	(3,-5)

If we change, say, 100, 3, 0, -1, with any $a > b > c > d$ the analysis doesn't change.

CHAPTER 2

Cooperative games– pure strategies

Games in strategic form

Games consist of players, available actions, and utilities of outcomes.

DEFINITION. A *strategic form* of an n-person game with players $1, 2, \ldots, n$ consists of sets $S_1, S_2, \ldots, S_n$ and functions $u_i : S_1 \times S_2 \times \cdots \times S_n \to \mathbb{R}$ $(1 \leq i \leq n)$, S_i being the set of *actions* (or *strategies*) for player i and u_i the utility function for player i .

Elements in $S_1 \times S_2 \times \cdots \times S_n$ are called *strategy profiles.*

The strategic form of a Prisoners' Dilemma

Here $n = 2$, $S_1 = S_2 = \{\text{confess}, \text{do not confess}\}$, and u_1, u_2 take values
$u_1(\{\text{confess}, \text{confess}\}) = u_2(\{\text{confess}, \text{confess}\}) = -5$,
$u_1(\{\text{do not confess}, \text{do not confess}\}) = u_2(\{\text{do not confess}, \text{do not confess}\}) = -1$,
$u_1(\{\text{confess}, \text{do not confess}\}) = 10$, $u_2(\{\text{confess}, \text{do not confess}\}) = -10$,
$u_1(\{\text{do not confess}, \text{confess}\}) = -10$, $u_2(\{\text{do not confess}, \text{confess}\}) = 10$.

Dominant strategies

DEFINITION. Consider a game in strategic form $(S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$.

A strategy $s_i \in S_i$ *dominates* $s'_i \in S_i$ if

$$u_i(s_1, \ldots, s_i, \ldots, s_n) > u_i(s_1, \ldots, s'_i, \ldots, s_n)$$

for all

$$s_1 \in S_1, \ldots, s_{i-1} \in S_{i-1}, s_{i+1} \in S_{i+1}, \ldots, s_n \in S_n.$$

We denote this $s_i \gg s'_i$.

If $s_i \gg s'_i$, then it cannot possibly be advantageous ever for player i to choose strategy s'_i, because choosing s_i would always fare better, no matter what other players do!

Example: iterative elimination of dominated strategies

Consider the following game given in tabular form

	l	r
U	0, 0	-1,-1
D	-3, 3	1, 1

(here there are two, players: 1 and 2 (aka Alice and Bob), $S_1 = \{U, D\}$, $S_2 = \{l, r\}$.)

$0 > -1$, $3 > 1$ so $l \gg r$, and we can delete r!

	l
U	0, 0
D	-3, 3

Now $0 > -3$, so $U \gg D$, and so Alice plays U.

This is an example of a *dominance-solvable game*, i.e., a game for which iterative elimination of dominated strategies results in a game with one action for each player.

Another example: I ask all students in this class to choose an integer between 1 and 100. The person whose number is closest to 2/3 of the average will receive £5. Which number do you choose?

Let x denote 2/3 of the average of the class choices.

$x < 67$, would anyone choose ≥ 68?
Now we know that $x < 45$, would anyone choose ≥ 45?
Now we know that $x \leq 30$, would anyone chose > 30?

We can continue this process until we are left with the single strategy of choosing 1.

(Why is this argument likely to fail in real life?)

Example: The Median Voter Theorem

Consider two candidates running for election, and assume that the electors care only about one issue which can be quantified with one number X (e.g., tax rates, immigration quotas, etc.)

The candidates, being politicians, will adopt the position which is most likely to get them elected. The Median Voter Theorem states that the winning strategy is to adopt the position of the median voter.

To see this assume that voters vote for the candidate whose position is closest to theirs (and if there is a tie they vote to either candidate with probability 50%), and assume, for simplicity, that the candidates can adopt the positions of the 10% centile, 20% centile, ... up to the 90% centile.

If we also assume that each candidate gets half the voters in between two positions we get the following game:

	10%	20%	30%	40%	50%	60%	70%	80%	90%
10%	50%	15%	20%	25%	30%	35%	40%	45%	50%
20%	85%	50%	25%	30%	35%	40%	45%	50%	55%
30%	80%	75%	50%	35%	40%	45%	50%	55%	60%
40%	75%	70%	65%	50%	45%	50%	55%	60%	65%
50%	70%	65%	60%	55%	50%	55%	60%	65%	70%
60%	65%	60%	55%	50%	45%	50%	65%	70%	75%
70%	60%	55%	50%	45%	40%	35%	50%	75%	80%
80%	55%	50%	45%	40%	35%	30%	25%	50%	85%
90%	50%	45%	40%	35%	30%	25%	20%	15%	50%

(entries in the table denote Alice's share of the vote.) Alice's 20% strategy dominates her 10% strategy, her 80% strategy dominates her 90% strategy. Eliminate these and similarly Bob's 10% and 90% strategies.

Now, (but not before!) the 30% strategies dominate the 20% strategies, and the 70% strategies dominate the 80% strategies–

	20%	30%	40%	50%	60%	70%	80%
20%	50%	25%	30%	35%	40%	45%	50%
30%	75%	50%	35%	40%	45%	50%	55%
40%	70%	65%	50%	45%	50%	55%	60%
50%	65%	60%	55%	50%	55%	60%	65%
60%	60%	55%	50%	45%	50%	65%	70%
70%	55%	50%	45%	40%	35%	50%	75%
80%	50%	45%	40%	35%	30%	25%	50%

Now, (but not before!) the 30% strategies dominate the 20% strategies, and the 70% strategies dominate the 80% strategies– eliminate the dominated strategies. We eventually end up with a single 50% strategy for both candidates.

Weakly dominant strategies

Definition. Consider a game in strategic form $(S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$.

A strategy $s_i \in S_i$ *weakly dominates* $s'_i \in S_i$ if

$$u_i(s_1, \ldots, s_i, \ldots, s_n) \geq u_i(s_1, \ldots, s'_i, \ldots, s_n)$$

for all

$$s_1 \in S_1, \ldots, s_{i-1} \in S_{i-1}, s_{i+1} \in S_{i+1}, \ldots, s_n \in S_n.$$

Example: Sealed-bid second-price auctions

An auction is a game played by potential buyers of an item whose outcome consists of a player and the price the player pays for the item.

For example sellers in *Ebay* sell their items in a process which is essentially equivalent to a *sealed-bid second-price auction*: in such an auction each bidder bids an amount of money *without knowing the other bids.* The item is then sold to the highest bidder at the *price bid by the second highest bidder.* If the value of the item to the winner of the auction is v and the second highest bid was x, the winner's payoff is $v - x$, and everyone else's payoff is zero.

Proposition 1. *Bidding one's valuation of the item is a weakly dominant strategy.*

Proof. Consider a bidder who values the item at v, and consider the payoff P_B of the strategy consisting of bidding B, as a function of the second highest bid x.

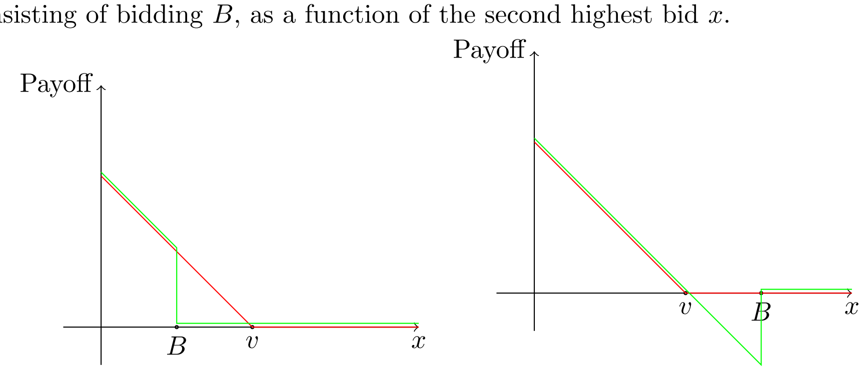

Here the green line is the graph of the payoff P_B of bidding B, and the red line is the graph of the payoff P_v, of bidding v, both as functions of the second highest bid x.

Bidding B produces the payoff

$$P_B(x) = \begin{cases} 0, & \text{if } x > B, \\ v - x & \text{if } x < B, \end{cases}$$

(note the tie breaker when $B = x$).

Now note that $P_v(x) \geq P_B(x)$ for all x:

- if $B \leq v$ then $P_B(x) = P_v(x) = v - x$ for $x \leq B$, and $P_B(x) = 0 \leq P_v(x)$ for $x > B$,
- if $B \geq v$ then $P_B(x) = P_v(x) = v - x$ for $x \leq v$, $P_B(x) = v - x \leq 0 = P_v(x)$ for $v \leq x \leq B$, and $P_B(x) = 0 = P_v(x)$ for $x > B$.

□

Pareto optimality

What is the correct notion of "solution" of a game which is not dominance-solvable? Maybe outcomes which make everyone as happy as possible?

DEFINITION. Consider a game in strategic form $(S_1, \ldots S_n, u_1, \ldots, u_n)$.

A strategy profile is *Pareto optimal* if there is no other strategy profile that all players prefer, i.e., $(s_1, \ldots s_n) \in S_1 \times \cdots \times S_n$ is *Pareto optimal* if for any other $(s_1', \ldots s_n') \in S_1 \times \cdots \times S_n$ there is a $1 \leq j \leq n$ such that $u_j(s_1, \ldots, s_n) > u_j(s_1', \ldots, s_n')$.

Pareto optimal strategies, even if they exist, are not stable! For example, the solution of the Prisoners' Dilemma is not Pareto optimal. So Pareto optimality is not very useful for our purposes.

Stabiliy. Example: the investment game

We are all given the option of investing £10. If at least 90% of us choose to invest, we get back £20, otherwise we lose our £10 investment.

Do you invest?

Now lets do it for real– do you invest?

This poses a coordination problem: similar to adoption of an industry standard (VHS versus Betamax), joining a social network software (facebook versus myspace), bank runs!

Stability: best responses and Nash equilibria

Think of players who can change their strategies once their opponents' strategy is revealed. A stable strategy profile would be one such that no player would change their choice after learning the other players' moves. We call such a pair of strategies a *no regrets* equilibrium.

EXAMPLE.

$$\begin{tikzcd} & l & r \\ U & 3,3 \arrow[r] \arrow[d] & -1,5 \arrow[d] \\ D & 5,-1 \arrow[r] & 0,0 \end{tikzcd}$$

The pair of strategies Dr is the only pair which neither player would want to move away from. Notice that both players would be better off if Ul were played.

Example: no stable strategy profile

$$
\begin{tikzcd}[row sep=small]
 & l & r \\
U & 2,4 \arrow[d] & 1,0 \arrow[l] \\
D & 3,1 \arrow[r] & 0,4 \arrow[u]
\end{tikzcd}
$$

Best response

DEFINITION. Consider the game n-person game $(S_1,\dots,S_n,u_1,\dots,u_n)$.

An action $s_i \in S_i$ is a *best response* to a given set of actions $\{s_j \in S_j\}_{j\neq i}$ if $u_i(s_1,\dots,s_{i-1},s_i,s_{i+1},\dots,s_n) \geq u_i(s_1,\dots,s_{i-1},s,s_{i+1},\dots,s_n)$ for all $s \in S_i$.

EXAMPLE.

	l	m	r
U	1, $\underline{3}$	$\underline{4}$, 2	2, 2
M	$\underline{4}$, 0	0, $\underline{3}$	4, 1
D	2, 5	3, 4	$\underline{5}$, $\underline{6}$

The best responses to strategies l, m and r are M, U and D, respectively. The best responses to strategies U, M and D are l, m and r, respectively.

Nash equilibrium

DEFINITION. Consider the game n-person game $(S_1,\dots,S_n,u_1,\dots,u_n)$.
A strategy profile $(s_1,\dots,s_n)$ is a *Nash equilibrium* if for all $1 \leq i \leq n$, s_i is a best response to $s_1,\dots,s_{i-1},s_{i+1},\dots,s_n$.

EXAMPLE.

	l	m	r
U	1, $\underline{3}$	$\underline{4}$, 2	2, 2
M	$\underline{4}$, 0	0, $\underline{3}$	4, 1
D	2, 5	3, 4	$\underline{5}$, $\underline{6}$

Here Dr is a Nash equilibrium.

PROPOSITION 2. *A Nash equilibrium is a no-regrets equilibrium.*

Example: A coordination game with multiple Nash equilibria

	a	b
I	$\underline{0}$, $\underline{0}$	-1,-3
II	-3, -1	$\underline{1}$, $\underline{1}$

Both Ia and IIb are Nash equilibria!

Example: Best response functions

Alice and Bob collaborate on a project, each devoting up to 4 hours to it. If Alice works $0 \leq x \leq 4$ hours and Bob works $0 \leq y \leq 4$, their combined profit will be $4(x+y)+xy$ which they share equally.

For both the cost of devoting h hours to the project is h^2, hence if they devote x and y hours, respectively, their total utilities are $A(x,y) = 2(x+y)+xy/2-x^2$ and $B(x,y) = 2(x+y)+xy/2-y^2$, respectively.

Note that this is a game with infinite sets of actions!

$A(x,y) = 2(x+y) + xy/2 - x^2$, $B(x,y) = 2(x+y) + xy/2 - y^2$.
Find Alice's best response to Bob's y_0 hours' work, i.e., find the value of x which maximizes $A(x, y_0)$: $\partial A/\partial x = 2 + y/2 - 2x$ and $\partial^2 A/\partial x^2 = -2 < 0$, maximum is at $x = 1 + y_0/4$.
Similarly, for a fixed $0 \le x_0 \le 4$, $B(x_0, y)$ is maximized at $y = 1 + x_0/4$.
If we want x_0 and y_0 to be best responses to each other we solve

$$\begin{cases} x_0 = 1 + y_0/4 \\ y_0 = 1 + x_0/4 \end{cases}$$

and obtain $x_0 = y_0 = 4/3$.
Notice that this solution is not Pareto optimal: If both Alice and Bob devote h hours to the project, they both receive utility worth $h(4 - h/2)$ which is maximized when $h = 4$. They would be better off is someone could impose this better arrangement.

Elimination of dominated strategies does not loose Nash equilibria

PROPOSITION 3. *Consider a game* $(S_1, \ldots, S_n, u_1, \ldots, u_n)$. *If* $s_i, s_i' \in S_i$ *and* $s_i \gg s_i'$, s_i' *cannot occur in a Nash equilibrium.*

PROOF. Let $s_j \in S_j$ for all $j \neq i$. Since $u_i(s_1, \ldots, s_{i-1}, s_i, s_{i+1}, \ldots, s_n) > u_i(s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n)$ for $s_j \in S_j$ $(j \neq i)$, s_i' cannot be a best response to this $\{s_j \in S_j\}_{j \neq i}$ and so it cannot be a best response to anything. □

Example: weak domination and Nash equilibria

Consider the game

	l	r
U	$\underline{1}$, $\underline{1}$	100, 0
D	0, 100	$\underline{100}$, $\underline{100}$

We see that U weakly dominates D and l weakly dominates r.
We can also verify that both Ul and Dr are Nash equilibria.
If we discard the weakly dominated strategies we loose the Pareto optimal Nash equilibrium!

Weird Example: Braess's Paradox

A large number of drivers travel from A to B either through C or D. Travel times are $AC = f, CB = 1.2, AD = 1.2, DB = f'$ where f, f' is the proportion of drivers going through C and D.

Can this be modelled as a game? What are its Nash equilibria?

There are many equilibria corresponding to is $f = 1/2$ giving a travel time of 1.7.

Now if a road from C to D taking 0.1 hours is opened, all will travel through it and resulting time will be 2.1!

CHAPTER 3

Nash equilibria in Economics: monopolies, duopolies and oligopolies

How do producers control their profits?

- Control of production levels (*Cournot model*).
- Control of prices (*Bertrand model*).

In practice both approaches are used by firms, but we study these in isolation.

Throughout this chapter we will consider a product whose cost of production per unit is c.

The Cournot model

Assumption: the price p of each unit of product is a (decreasing) function of the total supply q; specifically we will set $p(q) = a - bq$ in our discussion below.

We must have $p > c$ otherwise there is no profit to be made and hence no production, so $a > c$.

Monopolies. A monopolist faces a simple problem: maximize $f(q) = (p(q)-c)q = (a-bq-c)q$. The maximum occurs at $q = (a-c)/2b$, resulting in a price of $p = a - (a-c)/2 = (a+c)/2$ profit of $(a-c)^2/4b$.

Duopolies. We model two firms producing an identical product.
These two players's strategies amount to deciding their production levels q_1 and q_2.
The price per unit of product is given by $p = a - b(q_1 + q_2)$.
The profits then are

$$(p-c)q_1 = (a - b(q_1+q_2) - c)q_1 = aq_1 - bq_1^2 - bq_1q_2 - cq_1$$

and $(p-c)q_2 = aq_2 - bq_2^2 - bq_1q_2 - cq_2$.

We find Nash equilibria by finding the firms' best responses. If company 2 produces q_2 units, company 1 needs to maximize $aq_1 - bq_1^2 - bq_1q_2 - cq_1$:
differentiate with respect to q_1, set that to be zero and solve for q_1:
from $a - 2bq_1 - bq_2 - c = 0$ and we obtain $\mathrm{BR}_1(q_2) = (a - bq_2 - c)/2b = (a-c)/2b - q_2/2$.

Similarly, the second company's best response to company 1 producing q_1 units is $BR_2(q_1) = (a-c)/2b - q_1/2$.

To find the Nash equilibrium we solve the system of equations $q_1^* = BR_1(q_2^*), q_2^* = BR_2(q_1^*)$, and by symmetry we see that $q_1^* = q_2^*$.
We now obtain $q_1^* = q_2^* = (a-c)/3b$.

The price is now $p = a - 2b(a-c)/3b = a - 2(a-c)/3 = (a+2c)/3$.
Since $a > c$, the duopoly price is lower than the monopoly price of $(a+c)/2$; further, total production is greater than in a monopoly.

This is partly why societies regulate against monopolies: monopolies produce less and sell their products at higher prices to the detriment of consumers.

The total profits are now

$$(p-c)(q_1+q_2) = ((a+2c)/3-c)\,\frac{2(a-c)}{3b} = \frac{2}{9}\frac{(a-c)^2}{b}$$

and so the profit for each company is $(a-c)^2/9b$, less than half the monopoly profit.
Notice that if both companies collude to form a *cartel* and agree to produce $(a-c)/4b$ each, their total profits increase to the monopoly profit

$$(a-(a-c)/2-c)(a-c)/2b = (a-c)^2/4b.$$

Oligopolies and perfect competition. We now model n firms producing an identical product. As before, the firms decide their production levels $q_1, \dots, q_n$.
The price per unit of product is now given by $p = a - b(q_1 + \cdots + q_n)$ and the profits for the ith firm are

$$f_i(q_1,\dots,q_n) = (p-c)q_i = (a - b(q_1+\cdots+q_n) - c)q_i.$$

To maximize this we compute

$$\frac{\partial f_i}{\partial q_i} = a - b(q_1+\cdots+q_n) - c - bq_i$$

and set it to zero.

We now obtain a system of equations

$$\begin{cases} 2q_1 & + & q_2 & + & \dots & q_{n-1} & + & q_n & = & (a-c)/b \\ q_1 & + & 2q_2 & + & \dots & q_{n-1} & + & q_n & = & (a-c)/b \\ & & \vdots & & \dots & & & & & \\ q_1 & + & q_2 & + & \dots & q_{n-1} & + & 2q_n & = & (a-c)/b \end{cases}$$

which has a unique solution

$$q_1^* = \cdots = q_n^* = \frac{a-c}{(n+1)b}.$$

The total production is

$$\frac{n}{n+1}\frac{a-c}{b}$$

giving a price $p = a - \dfrac{n}{n+1}(a-c) = (a+nc)/(n+1)$,
and total profits

$$\left(\frac{a+nc}{n+1} - c\right)\frac{n}{n+1}\frac{a-c}{b} = \left(\frac{a-c}{n+1}\right)^2 \frac{n}{b}.$$

To model *perfect competition* we take $n \to \infty$ to obtain total production $(a-c)/b$, price c and total profit 0!

The Bertrand model

Control prices directly, affecting production levels indirectly through varying demand.

Specifically, we will assume that firms can produce any amount of their products, and that consumers will buy only from the cheapest producer.
We continue to assume that $p = a - bq$, but now we take the price p to be the independent variable and obtain an expression for the demand $q = (a-p)/b$.

We model a *Bertrand duopoly*: the two players first set their prices p_1 and p_2. The profit functions for both company are

$$f_1(p_1, p_2) = \begin{cases} (p_1 - c)\,(a - p_1)/b & \text{if } p_1 < p_2 \\ (p_1 - c)\,(a - p_1)/2b & \text{if } p_1 = p_2 \\ 0 & \text{if } p_1 > p_2 \end{cases}$$

$$f_2(p_1, p_2) = \begin{cases} (p_2 - c)\,(a - p_2)/b & \text{if } p_2 < p_1 \\ (p_2 - c)\,(a - p_2)/2b & \text{if } p_1 = p_2 \\ 0 & \text{if } p_2 > p_1 \end{cases}$$

The only Nash equilibrium here is $p_1^* = p_2^* = c$ for a profit of zero! To see this, we note that if $c < p_1^* \leq p_2^*$, firm 2 profits by changing its price to a bit less than p_1^*, and similarly if $c < p_2^* \leq p_1^*$, firm 1 profits by changing its price to a bit less than p_2^*.

CHAPTER 4

Cooperative games– mixed strategies

Example: matching pennies

Alice and Bob both turn a penny with either heads or tails. If they match, Alice takes them, otherwise Bob takes them.

	H	T
H	1, -1	-1, 1
T	-1, 1	1, -1

Alice hires you as her Game Theory adviser; what do you advise?

Game Theory cannot possibly advise Alice to play H or T; Bob's adviser is just as knowledgable, and if, say, the correct move for Alice is H, Bob will be instructed to play T and win.

There is no correct action here, unless . . . we expand the notion of a strategy: Alice is advised to choose her move *randomly*, say, play H with probability α and T with probability $1-\alpha$. Similarly Bob will choose H with probability β and T with probability $1-\beta$.

Expected utility

Assumption: Given the choice between two uncertain outcomes, players will choose the one with highest expected utility. (The Von-Neumann, Morgenstein expected utility hypothesis.)

Note: With certain outcomes, utility encoded rankings. Now the actual values of the utility function matter!

Mixed strategies and mixed-strategy Nash equilibria

Consider a game $(S_1, \ldots, S_n, u_1, \ldots, u_n)$.

Definition. A *mixed strategy* for player i is a function $p_i : S_i \to [0,1]$ such that $\sum_{s_i \in S_i} p(s_i) = 1$. (We interpret $p(s_i)$ as the probability that player i plays s_i.) Given mixed strategies $p_1, \ldots, p_n$ for players $1, \ldots, n$ we define

$$u_i(p_1, \ldots, p_n) = \sum_{(s_1,\ldots,s_n) \in S_1 \times \cdots \times S_n} p_1(s_1) \ldots p_n(s_n) u_i(s_1, \ldots, s_n)$$

(which is the expected payoff for player i when these mixed strategies are played.)

Definition. Mixed strategies $p_1, \ldots, p_n$ for players $1, \ldots, n$ are a *mixed strategy Nash equilibrium* of the game if for all $1 \le i \le n$ and any mixed strategy p_i' for player i we have

$$u_i(p_1, \ldots, p_{i-1}, p_i, p_{i+1} \ldots p_n) \ge u_i(p_1, \ldots, p_{i-1}, p_i', p_{i+1} \ldots p_n).$$

(Each p_i gives a response with highest expected value to the other mixed strategies.)

Example: matching pennies again

Alice plays p, $p(H) = \alpha$, $p(T) = 1 - \alpha$;
Bob plays q, $q(H) = \beta$, $q(T) = 1 - \beta$.

We have

$$u_1(p,q) = \alpha\beta + (1-\alpha)(1-\beta) - \alpha(1-\beta) - (1-\alpha)\beta = (4\beta - 2)\alpha + 1 - 2\beta$$

$u_2(p,q) = -u_1(p,q)$.

To maximize u_1 for a given β, we take $\alpha = 1$ if $\beta > 1/2$, $\alpha = 0$ if $\beta < 1/2$ and any α when $\beta = 1/2$. Similarly, u_2 is maximized by $\beta = 0$ if $\alpha > 1/2$, $\beta = 1$ if $\alpha < 1/2$ and any β when $\alpha = 1/2$.

The only mixed strategy Nash equilibrium occurs with $\alpha = \beta = 1/2$.

Example: The battle of the sexes

Alice and Bob decide to go on a date together: Alice suggests going to the pub and Bob suggests watching a film. Before they can agree, the MAS348 lecture is over and they rush to different lectures. Realizing they don't have a way to contact each other before their date, they have to decide independently where to go. They are playing the following game:

	Pub	Film
Pub	3,1	0,0
Film	0,0	1,3

This game has two Nash equilibria in pure strategies, namely (Pub, Pub) and $(Film, Film)$. Are there any other Nash equilibria?

	Pub	Film
Pub	3,1	0,0
Film	0,0	1,3

Suppose Alice goes to the pub with probability p and Bob goes to the film with probability q. Alice's expected utility is $3p(1-q) + (1-p)q = p(3-4q) + q$ and Bob's expected utility is $p(1-q) + 3(1-p)q = q(3-4p) + p$. To find Alice's best responses we distinguish between three cases:

$q < 3/4$ and Alice's expected utility is maximized when $p = 1$,
$q > 3/4$ and Alice's expected utility is maximized when $p = 0$,
$q = 3/4$ now Alice's expected utility is $3/4$, regardless of p.

Similarly, to find Bob's best responses we distinguish between three cases:

$p < 3/4$ and Bob's expected utility is maximized when $q = 1$,
$p > 3/4$ and Bob's expected utility is maximized when $q = 0$,
$p = 3/4$ now Bob's expected utility is $3/4$, regardless of q.

The strategy profile corresponding to $p = q = 3/4$ is a Nash equilibrium!

Note: in a Nash equilibrium each player had an expected utility independent of the strategy adopted by the other player– *this is not a coincidence.*

Notation. For any action s of a player in a game, $\widehat{s}$ denotes the mixed strategy which plays s with probability 1.

THEOREM 4 (The Indifference Principle). *Consider a game (S_1, S_2, u_1, u_2). Let (p, q) be a mixed strategy profile which is a Nash equilibrium. For any $s \in S_1$ with $p(s) > 0$ we have $u_1(\widehat{s}, q) = u_1(p, q)$.*

PROOF. Let $\mathrm{Supp}(p) = \{s \in S_1 \mid p(s) > 0\}$ (the “support” of p).

We have $u_1(p,q) = \sum_{t \in \mathrm{Supp}(p)} p(t) u_1(\widehat{t}, q)$.
If the theorem does not hold, there must be an $s \in \mathrm{Supp}(p)$ for which $u_1(\widehat{s}, q) > u_1(p,q)$, contradicting the fact that p is a best response to q. □

Example: Rock-Scissors-Paper

Consider the game

	Rock	Scissors	Paper
Rock	0, 0	1, -1	-1, 1
Scissors	-1, 1	0, 0	1, -1
Paper	1, -1	-1, 1	0, 0

There are no dominant strategies, nor are there Nash equilibria in pure strategies. To find mixed strategy profiles which are Nash equilibria, find a mixed strategy (p_1, p_2, p_3) for Alice which gives Bob equal utility for all his pure strategies: this happens in $p_2 - p_3 = -p_1 + p_3 = p_1 - p_2$, which together with $p_1 + p_2 + p_3 = 1$ yields the mixed strategy $(1/3, 1/3, 1/3)$. Similarly, Bob's strategy is $(1/3, 1/3, 1/3)$.

Example: Alice and Bob play tennis

Bob is at the net and Alice needs to decide to hit the ball to Bob's right or left; and Bob needs to decide to jump to left or right. The probabilities of Bob responding are as follows:

	left	right
left	50%	20%
right	10%	80%

and we assume Alice and Bob are playing the following game

	left	right
left	0.5, 0.5	0.8, 0.2
right	0.9, 0.1	0.2, 0.8

	left	right
left	0.5, 0.5	0.8, 0.2
right	0.9, 0.1	0.2, 0.8

There are no dominant strategies, nor are there Nash equilibria in pure strategies.
When is $((p, 1-p), (q, 1-q))$ a Nash equilibrium?

this happens when $0.5p + 0.1(1-p) = 0.2p + 0.8(1-p)$, which gives $p = 0.7$ and $0.5q + 0.8(1-q) = 0.9q + 0.2(1-q)$, which gives $q = 0.6$.

Example: The attrition game

Consider the game

	Stay	Leave
Stay	-2, -2	1, 0
Leave	0, 1	0, 0

There are two pure-strategy Nash equilibria: (Stay, Leave) and (Leave, Stay). There is also a mixed-strategy Nash equilibrium: ((1/3,2/3), (1/3, 2/3)).

Change this to

	Stay	Leave
Stay	-2, -2	x, 0
Leave	0, 1	0, 0

for $x > 1$. We preserve the pure-strategy Nash equilibria, but the mixed-strategy Nash equilibrium is now: ((1/3,2/3), (x/(2+x), 2/(2+x))).

Change this to

	Stay	Leave
Stay	-2, -2	1, 0
Leave	y, 1	y, 0

for $0 < y < 1$. How does the mixed-strategy Nash equilibrium change as y increases?

Dominating mixed strategies

Even when no action strictly dominates another, a mixed strategy might.

Consider the game

	L	R
U	3, 1	0, 2
M	1, 1	1, 0
D	0, 0	3, 1

None of the actions is dominated, but playing U and D with probability 1/2 each dominates M. After eliminating M, Bob's R dominates L, and we obtain a solution (D, R) for this game.

Existence of Nash equilibria

Example (A game with no Nash equilibrium). Two players choose an integer. The player with the biggest integer gets 1, the other 0, and in case of a tie both get 0.

Example (Another game with no Nash equilibrium). $n > 1$ players choose a number in the open interval $(0, 1)$ and the mean μ of these numbers is computed. The players closest to $\mu/2$ win, the others lose.

Theorem 5 (Nash's Theorem). *Any game $(S_1, \ldots, S_n, u_1, \ldots, u_n)$ with finite strategy sets has at least one Nash equilibrium.*

The proof of this theorem is quite intricate!

CHAPTER 5

Two-person zero-sum games

"Maxminization"

You play a game and want to ensure you don't do too badly.
You find the worst possible outcome of each of your actions and choose the action with highest possible *worst* outcome.
E.g., in a 2-person game (S, T, u_1, u_2), if we restrict to pure strategies Alice plays the $s \in S$ which maximizes $\min_{t \in T} u_1(s,t)$ and Bob plays the $t \in T$ which maximizes $\min_{s \in S} u_2(s,t)$.

This guarantees Alice at least $\max_{s \in S} \min_{t \in T} u_1(s,t)$ and
Bob $\max_{t \in T} \min_{s \in S} u_2(s,t)$.

"Maxminimizing" with mixed strategies

If P and Q denote the mixed-strategies for Alice and Bob, Alice could play the $p \in P$ which maximizes $\min_{q \in Q} u_1(p,q)$ and
Bob plays the $q \in Q$ which maximizes $\min_{p \in P} u_2(p,q)$.

Zero-sum games

Definition. A two-person zero-sum game a game in strategic form (S_1, S_2, u_1, u_2) where $u_1(s_1, s_2) + u_2(s_1, s_2) = 0$ for all $s_1 \in S_1$ and $s_2 \in S_2$. ("My gain is your loss, your gain is my loss.")

In a zero-sum game a player's best response is always the one that does most damage to the other!

Games in which one side wins and the other loses, or both sides draw, can be modeled as zero-sum games, e.g., chess and Rock-Scissors-Paper.

Wars are seldom zero-sum games.

Unlike cooperative games, zero-sum games can be solved, in some sense.

Pure strategies: looking for saddle points

Example. Consider the zero-sum game

	a	b	c
I	1,-1	1,-1	8, -8
II	5, -5	2, -2	4, -4
III	7, -7	0, 0	0, 0

We don't need to keep track of both sides' payoffs, and by convention we keep track of the row-player payoffs only. The abbreviated form of the game is

	a	b	c
I	1	1	8
II	5	2	4
III	7	0	0

and we set $u = u_1$ to be the row-players's utility function.

How should both sides maxminimize this? Consider the case where only pure strategies are played.

Bob's responses to each of Alice's strategies are the ones resulting in minimal score for Alice:

	a	b	c
I	$\boxed{1}$	$\boxed{1}$	8
II	5	$\boxed{2}$	4
III	7	$\boxed{0}$	$\boxed{0}$

So Alice would choose the strategy $s = s^*$ which maximizes

$$\max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t),$$

i.e., $s^* = II$, and choosing this guarantees her a payoff of 2.

Alice's responses to each of Bob's strategies are the ones resulting in maximal score for Alice:

	a	b	c
I	1	1	$\boxed{8}$
II	5	$\boxed{2}$	4
III	$\boxed{7}$	0	0

So Bob would choose the strategy $t = t^*$ which minimizes

$$\min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t),$$

i.e., $t^* = b$, and this strategy guarantees a payoff of -2 for Bob.

Note that (II, b) is a Nash equilibrium; this is because (II, b) is a *saddle-point* for u, i.e., $u(II, b) \geq u(s, b)$ for all $s \in \{I, II, III\}$ and $u(II, b) \leq u(II, t)$ for all $t \in \{a, b, c\}$.

In this example

$$\max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t) = \min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t) = 2$$

This implies that for s^* and t^* such that

$$\min_{t\in\{a,b,c\}} u(s^*,t) = \max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t)$$

$$\max_{s\in\{I,II,III\}} u(s,t^*) = \min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t),$$

(s^*, t^*) is a Nash equilibrium.

As long as Bob plays t^* Alice cannot get more than 2,
as long as Alice plays s^* Bob cannot get more than -2.

Another example

Consider the zero-sum game

	a	b	c
I	2, -2	5, -5	0, 0
II	3, -3	1, -1	2, -2
III	4, -4	3, -3	6, -6

Knowing that this game is zero-sum, we don't need to keep track of both sides' payoffs, and by convention we keep track of the row-player payoffs only. The abbreviated form of the game is

	a	b	c
I	2	5	0
II	3	1	2
III	4	3	6

and we set $u = u_1$ to be the row-players's utility function.

Bob's responses to each of Alice's strategies are the ones resulting in minimal score for Alice:

	a	b	c
I	2	5	$\boxed{0}$
II	3	$\boxed{1}$	2
III	4	$\boxed{3}$	6

So Alice would choose the strategy $s = s^*$ which maximizes $\min_{t\in\{a,b,c\}} u(s,t)$,. i.e. the value $s = s^*$ for which

$$\max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t) = \min_{t\in\{a,b,c\}} u(s^*,t),$$

i.e., $s^* = III$, and choosing this guarantees her a payoff of 3.

Alice's responses to each of Bob's actions are the ones resulting in maximal score for Alice:

	a	b	c
I	2	$\boxed{5}$	0
II	3	1	2
III	$\boxed{4}$	3	$\boxed{6}$

So Bob would choose the strategy $t = t^*$ which minimizes $\max_{s\in\{I,II,III\}} u(s,t)$, i.e., the value $t = t^*$ for which

$$\min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t) = \max_{s\in\{I,II,III\}} u(s,t^*),$$

i.e., $t^* = a$, and this strategy guarantees a payoff of -4 for Bob.

Note that

$$\min_{t\in\{a,b,c\}} \max_{s\in\{I,II,III\}} u(s,t) = 4 > 3 = \max_{s\in\{I,II,III\}} \min_{t\in\{a,b,c\}} u(s,t).$$

Note also that (III, a) is *not* a Nash equilibrium. Is this the best that can be achieved?

minmax $\geq$ maxmin

LEMMA 6. *Consider a finite zero-sum game (S,T,u). We have*

$$\min_{t\in T}\max_{s\in S} u(s,t) \geq \max_{s\in S}\min_{t\in T} u(s,t).$$

PROOF. Write $m = \max_{s\in S}\min_{t\in T} u(s,t)$. For all $s \in S$ and $t' \in T$ we have $u(s,t') \geq \min_{t\in T} u(s,t)$ so $\max_{s\in S} u(s,t') \geq \max_{s\in S}\min_{t\in T} u(s,t) = m$ so in particular $\min_{t\in T}\max_{s\in S} u(s,t) \geq m$. □

THEOREM 7. *Consider a finite zero-sum game (S,T,u). Let $s^* \in S$ and $t^* \in T$ be such that*

$$\min_{t\in T}\max_{s\in S} u(s,t) = \max_{s\in S} u(s,t^*)$$

and

$$\max_{s\in S}\min_{t\in T} u(s,t) = \min_{t\in T} u(s^*,t).$$

A strategy profile (s^*,t^*) is a saddle-point for u (and hence a Nash equilibrium) if and only if

$$\min_{t\in T}\max_{s\in S} u(s,t) = \max_{s\in S}\min_{t\in T} u(s,t).$$

Proof. Write $m = \max_{s\in S}\min_{t\in T} u(s,t)$ and $M = \min_{t\in T}\max_{s\in S} u(s,t)$.

Assume first that (s^*, t^*) is a saddle-point. Then for all $s \in S$ and $t \in T$ we have

$$u(s^*, t) \geq u(s^*, t^*) \geq u(s, t^*)$$

hence

$$\min_{t\in T} u(s^*, t) \geq u(s^*, t^*) \geq \max_{s\in S} u(s, t^*)$$

and

$$m = \max_{s\in S}\min_{t\in T} u(s,t) \geq u(s^*, t^*) \geq \min_{t\in T}\max_{s\in S} u(s,t) = M$$

and since $m \leq M$, we get equalities throughout.

Assume now that $m = M$. For all $s \in S$ and $t \in T$ we have

$$u(s^*, t) \geq \min_{\tau\in T} u(s^*, \tau) = m = M = \max_{\sigma\in S} u(\sigma, t^*) \geq u(s, t^*)$$

and in particular $u(s^*, t) \geq u(s^*, t^*)$ and $u(s, t^*) \leq u(s^*, t^*)$. □

Mixed strategies

Consider a finite zero-sum game (S, T, u) with $S = \{s_1, \ldots, s_m\}$ and $T = \{t_1, \ldots, t_n\}$.

The row player's set of mixed strategies is

$$\Delta^R = \{(p_1, \ldots, p_m) \,|\, p_1, \ldots, p_m \geq 0, p_1 + \cdots + p_m = 1\}$$

(s_i is played with probability p_i.) and the column player's set of mixed strategies is

$$\Delta^C = \{(q_1, \ldots, q_n) \,|\, q_1, \ldots, q_n \geq 0, q_1 + \cdots + q_n = 1\}$$

(t_i is played with probability q_i.)

The mixed strategy pair $(p, q) \in \Delta^R \times \Delta^C$ yields a payoff of

$$u(p,q) = \sum_{1\leq i\leq m}\sum_{1\leq j\leq n} p_i q_j u(s_i, t_j)$$

for the row player and $-u(p,q)$ for the column player.

Definition. Consider a zero sum game (S, T, u) with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. Let

$$\overline{V} = \min_{y\in\Delta^C}\max_{x\in\Delta^R} u(x,y)$$

and

$$\underline{V} = \max_{x\in\Delta^R}\min_{y\in\Delta^C} u(x,y).$$

An optimal strategy for the column player is a $y^* \in \Delta^C$ for which $\overline{V} = \max_{x\in\Delta^R} u(x, y^*)$.

An optimal strategy for the row player is a $x^* \in \Delta^R$ for which $\underline{V} = \min_{y\in\Delta^C} u(x^*, y)$.

Previous results hold for mixed strategies as well:

Lemma 8. *Consider a finite zero sum game (S, T, u) with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. We have*

$$\min_{y\in\Delta^C}\max_{x\in\Delta^R} u(x,y) \geq \max_{x\in\Delta^R}\min_{y\in\Delta^C} u(x,y).$$

Proof. (Essentially the same proof as for pure strategies.) Write $m = \max_{x\in\Delta^R}\min_{y\in\Delta^C} u(x,y)$. For all $x \in \Delta^R$ and $y' \in \Delta^C$ we have $u(x, y') \geq \min_{y\in\Delta^C} u(x,y)$ so $\max_{x\in\Delta^R} u(x, y') \geq \max_{x\in\Delta^R}\min_{y\in\Delta^C} u(x,y) = m$

so in particular $\min_{y\in\Delta^C}\max_{x\in\Delta^R} u(x,y) \geq m$. □ □

THEOREM 9. *Consider a finite zero-sum game (S, T, u) with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. Let $x^* \in \Delta^R$ and $y^* \in \Delta^C$ be optimal strategies, i.e.,*

$$\min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) = \max_{x \in \Delta^R} u(x, y^*)$$

and

$$\max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y) = \min_{y \in \Delta^C} u(x^*, y).$$

The strategy profile (x^*, y^*) is a Nash equilibrium if and only if

$$\min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y).$$

PROOF. Write $m = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y)$ and $M = \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y)$.

Assume first that (x^*, y^*) is a Nash equilibrium. Then for all $x \in \Delta^R$ and $y \in \Delta^C$ we have

$$u(x^*, y) \geq u(x^*, y^*) \geq u(x, y^*)$$

hence

$$\min_{y \in \Delta^C} u(x^*, y) \geq u(x^*, y^*) \geq \max_{x \in \Delta^R} u(x, y^*)$$

and

$$m = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y) \geq u(x^*, y^*) \geq \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y) = M$$

and since $m \leq M$, we get equalities throughout.

Assume now that $m = M$. For all $x \in \Delta^R$ and $y \in \Delta^C$ we have

$$u(x^*, y) \geq \min_{\overline{y} \in \Delta^C} u(x^*, \overline{y}) = m = M = \max_{\overline{x} \in \Delta^R} u(\overline{x}, y^*) \geq u(x, y^*)$$

and in particular $u(x^*, y) \geq u(x^*, y^*)$ and $u(x, y^*) \leq u(x^*, y^*)$. □

Two-person zero-sum games have mixed-strategy Nash-equilibria!

We now state one of the most important results in game theory: John von Neumann's *Minimax Theorem* (1928)

THEOREM 10 (The Minimax Theorem). $\overline{V} = \underline{V}$.

COROLLARY 11. *Two-person finite zero-sum games have at least one mixed-strategy Nash-equilibrium: any pair of optimal strategies is a Nash equilibrium.*

DEFINITION. Let $G = (S, T, u)$ be a zero sum game with sets Δ^R and Δ^C of mixed strategies for the row and column players, respectively. The value of G is defined as the common value of

$$\overline{V} = \min_{y \in \Delta^C} \max_{x \in \Delta^R} u(x, y)$$

and

$$\underline{V} = \max_{x \in \Delta^R} \min_{y \in \Delta^C} u(x, y).$$

Zero sum games in matrix form

A finite zero sum game (S, T, u) with $S = \{s_1, \ldots, s_m\}$ and $T = \{t_1, \ldots, t_n\}$ can be represented by a $m \times n$ matrix $M = ((u(s_i, t_j)))$ and mixed strategy sets

$$\Delta_m = \left\{ \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix} : x_1, \ldots, x_m \geq 0 \text{ and } x_1 + \cdots + x_m = 1 \right\} \subseteq \mathbb{R}^m$$

and

$$\Delta_n = \left\{ \begin{bmatrix} y_1 \\ \vdots \\ y_n \end{bmatrix} : y_1, \ldots, y_n \geq 0 \text{ and } y_1 + \cdots + y_n = 1 \right\} \subseteq \mathbb{R}^n$$

for the row and column players.

Now the payoff of a strategy pair $(x, y) \in \Delta_m \times \Delta_n$ is $x^t My$ for the row player and $-x^t My$ for the column player.

Example: symmetric games

A zero-sum game $G = (S, T, u)$ is *symmetric* if $S = T$ and for all $s_1, s_2 \in S$, $u(s_2, s_1) = -u(s_1, s_2)$. Let A be the matrix associated with a symmetric game $G = (\{1, \ldots, n\}, \{1, \ldots, n\}, u)$.

(a) The value V of a symmetric game is zero.

(b) An optimal strategy for one player is also an optimal strategy for the other player.

Here $C = R = \{1, \ldots, n\}$.

(a) For any $p \in \Delta^R$ we have $p^T Ap = (p^T Ap)^T = p^T A^T p = -p^T Ap$, hence $p^T Ap = 0$ so

$V = \min_{y \in \Delta^C} \max_{x \in \Delta^R} x^T Ay \geq \min_{y \in \Delta^C} y^T Ay = 0$ and $V = \max_{x \in \Delta^R} \min_{y \in \Delta^C} x^T Ay \leq \max_{x \in \Delta^R} x^T Ax = 0$.

(b) If $p \in \Delta^R$ is optimal for the row player, $p^T A$ has non-negative entries, and so $(p^T A)^T = A^T p = -Ap$ has non-negative entries and p is optimal for the column player.

Example: a symmetric game

Verify that the zero-sum game

	A	B	C
I	0	2	-1
II	-2	0	3
III	1	-3	0

has Nash-equilibrium $((1/2, 1/6, 1/3), (1/2, 1/6, 1/3))$.

PROPOSITION 12. *Let $(\{s_1, \ldots, s_m\}, \{t_1, \ldots, t_n\}, u)$ be a zero-sum game. Let M be the $m \times n$ matrix with $M_{ij} = u(s_i, t_j)$. Suppose that $p^* \in \Delta_m$ and $q^* \in \Delta_n$ are such that the minimal coordinate in ${p^*}^t M$ and the maximal coordinate in Mq^* both equal v. Then the value of the game is v and (p^*, q^*) is an optimal strategy.*

PROOF. Let V be value of the game.

$$V = \max_{p \in \Delta_m} \min_{q \in \Delta_n} p^t Mq \leq \max_{p \in \Delta_m} p^t Mq^* \leq \max_{p \in \Delta_m} p^t \begin{bmatrix} v \\ \vdots \\ v \end{bmatrix} = v$$

$$V = \min_{q\in\Delta_n}\max_{p\in\Delta_m} p^t M q \geq \min_{q\in\Delta_n} {p^*}^t M q \geq \min_{q\in\Delta_n} (v,\dots,v)q = v$$

hence $v = V$.

Also

$${p^*}^t M q^* \leq {p^*}^t \begin{bmatrix} v \\ \vdots \\ v \end{bmatrix} = v$$

$${p^*}^t M q^* \geq (v,\dots,v)q^* = v$$

hence ${p^*}^t M q^* = v$. □

Example: finding the value of a game

Consider the following zero-sum game

	A	B	C
I	2	-1	-1
II	-2	0	3
III	1	2	1

We verify that (p^*, q^*) is an optimal strategy profile where $p^* = (0,0,1)$ and $q^* = (2/5, 0, 3/5)$. Write

$$M = \max\{u(\widehat{I}, q^*), u(\widehat{II}, q^*), u(\widehat{III}, q^*)\} = \max\{1/5, 1, 1\} = 1$$

and

$$m = \min\{u(p^*, \widehat{A}), u(p^*, \widehat{B}), u(p^*, \widehat{C})\} = \min\{1, 2, 1\} = 1.$$

$$\underline{V} = \max_{p\in\Delta^R}\min_{q\in\Delta^C} u(p,q) = \max_{p\in\Delta^R}\min_{t\in T} u(p,\widehat{t}) \geq \min_{t\in T} u(p^*,\widehat{t}) = 1$$

$$\overline{V} = \min_{q\in\Delta^C}\max_{p\in\Delta^R} u(p,q) = \min_{q\in\Delta^C}\max_{s\in S} u(\widehat{s},q) \leq \max_{s\in S} u(\widehat{s},q^*) = 1$$

So $1 \leq \underline{V} \leq \overline{V} \leq 1$ hence $\underline{V} = \overline{V}$ and value of this game is 1.

Example: finding the value of a game

Consider the following zero-sum game

	A	B	C	D	E
I	-1	2	-2	0	1
II	-2	-1	3	2	0
III	2	1	0	-1	-2
IV	0	0	2	1	1
V	1	-1	0	-2	1

We verify that (p^*, q^*) is an optimal strategy profile where $p^* = (5/52, 0, 11/52, 34/52, 2/52)$ and $q^* = (21/52, 12/52, 0, 3/52, 16/52)$.

Compute

$$\begin{bmatrix} -1 & 2 & -2 & 0 & 1 \\ -2 & -1 & 3 & 2 & 0 \\ 2 & 1 & 0 & -1 & -2 \\ 0 & 0 & 2 & 1 & 1 \\ 1 & -1 & 0 & -2 & 1 \end{bmatrix} \begin{bmatrix} 21/52 \\ 12/52 \\ 0 \\ 3/52 \\ 16/52 \end{bmatrix} = \begin{bmatrix} 19/52 \\ -\mathbf{48/52} \\ 19/52 \\ 19/52 \\ 19/52 \end{bmatrix}$$

Bob's q^* guarantees not to lose more that 19/52.

Compute

$$\begin{bmatrix} 5/52 & 0 & 11/52 & 34/52 & 2/52 \end{bmatrix} \begin{bmatrix} -1 & 2 & -2 & 0 & 1 \\ -2 & -1 & 3 & 2 & 0 \\ 2 & 1 & 0 & -1 & -2 \\ 0 & 0 & 2 & 1 & 1 \\ 1 & -1 & 0 & -2 & 1 \end{bmatrix} =$$

$$\begin{bmatrix} 19/52 & 19/52 & \mathbf{58/52} & 19/52 & 19/52 \end{bmatrix}$$

So Alices p^* guarantees her at least 19/52.

The value of this game is 19/52.

A very different example: A duel

Alice and Bob fight a duel at dawn. They start at a distance of 1 unit apart, each armed with a pistol that can fire once, and they walk toward each other. Each can shoot at will: if the person who shoots first hits the target, he survives and the other dies; otherwise if the first person to shoot misses, the second person shoots the first point blank. If both decide to shoot at the same time, a fair coin is tossed to select who shoots first.

Alice's probability of hitting Bob at a distance of d is given by $A(d)$ and Bob's probability of hitting Alice at a distance of d is given by $B(d)$, and we assume that $A(d)$ and $B(d)$ are decreasing continuous functions of d and that $A(0) = B(0) = 1$.

How should this duel proceed?

Consider a strategy profile (x, y) (Alice shoots at a distance x and Bob shoots at a distance y.)

The probability of Alice surviving is

$$p(x, y) = \begin{cases} A(x) & \text{if } x > y \\ 1 - B(y) & \text{if } x < y \end{cases}$$

and that of Bob is $q(x, y) = 1 - p(x, y)$.

Alice wants to shoot at a distance x^* for which

$$\sup_{0 \le x \le 1} \inf_{0 \le y \le 1} p(x, y) = \inf_{0 \le y \le 1} p(x^*, y).$$

Define $\beta(x) = \inf_{0 \le y \le 1} p(x, y)$.

If $A(x) > 1-B(x)$, $\beta(x) = \inf_{x<y\le 1} p(x, y) = 1-B(x)$ and if $A(x) \le 1-B(x)$, then $\beta(x) = A(x)$.

Let $z \in [0, 1]$ be such that $1 - B(z) = A(z)$. We have

$$\beta(x) = \begin{cases} 1 - B(x) & \text{if } x < z \\ A(x) & \text{if } x \ge z \end{cases}$$

Alice will look for an x which maximizes $\beta(x)$ and this occurs at $x = z$. A similar analysis shows that Bob will also shoot at distance z.

CHAPTER 6

Sequential games

An example

Alice and Bob play the following game: Alice goes first and chooses A, B or C. If she chose A, the game ends and both get 0. If she chose B, Bob can either choose D resulting in utilities of 2 and 3, or he can choose E, resulting in utilities -2 and 2. If Alice chose C, Bob can either choose F resulting of utilities of 3 and 5, or G, resulting in utilities of -6 and 6. We can summarize this game with a rooted, directed tree as follows:

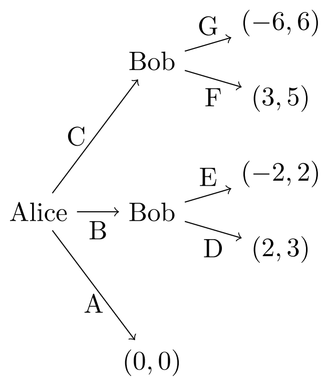

Example: Solving games using backward induction

If Bob has to choose between D and E, he will choose the outcome with highest utility for him, namely D, and if he needs to choose between F and G, he will choose G. So when Alice has to choose between A, B and C, she is choosing between utilities 0, 2 and -6, so she will choose B. The optimal strategies are for Alice to choose B and then for Bob to choose D, resulting in a payoff of 2 for Alice and 3 for Bob.

There is an outcome which is better for both Alice and Bob: If Alice chooses C and Bob F, Alice's payoff is $3 > 2$ and Bob's payoff is $5 > 3$. The problem is that if Alice chooses C, Bob will chose G for a payoff of $6 > 5$. But suppose we change the payoff $(-6, 6)$ to $(-6, 4)$: are we hurting Bob? No! Now the optimal strategy for both players leads to the optimal payoff $(3, 5)$. Surprisingly, lowering one's payoffs in a sequential game can lead to a better outcome.

Directed rooted trees

Definition. A *rooted tree* is a directed graph with a distinguished vertex, the *root*, which is connected to every other vertex with a unique directed path.
Vertices with no arrows coming out from them are *leaves*.
Leaves will encode payoffs, and the other vertices will be labeled by the players of the game.
Edges coming out of a vertex will be labeled by the possible decisions at that stage of the game.
The *rank* of a rooted tree is the maximum of all lengths of directed paths in the tree.

Example: "The hungry lions" game

A pride of lions consist of a highly hierarchical group of lions. When prey is caught, the top lion eats first, followed by lion number two, and so on. Once an especially large and tasty prey was caught, so large and tasty that it would be impossible for lions to stop eating it, and they would fall asleep. If a lion falls asleep, he will be eaten by the lion immediately below him in the hierarchy, but not by any other lion. This lion will then also fall asleep. Should the top lion eat the prey?

Suppose there are n lions. If presented with the choice of eating the prey, should lion number n eat it? Yes, no one will bother to eat it. Should lion $n-1$ eat? No! It would fall asleep and be eaten. Should the $n-2$ lion eat? We know that lion $n-1$ won't eat it, so it is safe for it to eat and fall asleep, etc. So the top lion should eat if and only if n is odd!

Cournot duopoly revisited: the Stackelberg model

Two firms produce an identical product. Unlike in the Cournot model, the second firm will choose to produce q_2 *after* the first firm produces q_1. As before, the cost of production per unit is denoted c, and the price per unit of product is given by $p = a - b(q_1 + q_2)$.

The second company's best response to company 1 producing q_1 units is $\mathrm{BR}_2(q_1) = (a-c)/2b - q_1/2$ resulting in total production of $q = q_1 + q_2 = q_1 + (a-c)/2b - q_1/2 = (a-c)/2b + q_1/2$. The first firm would then make a profit of

$$P(q_1) = (p-c)q_1 = \left(a - b\left(\frac{a-c}{2b} + q_1/2\right) - c\right) q_1 =$$
$$\left(\frac{a-c}{2} - \frac{q_1 b}{2}\right) q_1 = \frac{b}{2}\left(\frac{a-c}{b} - q_1\right) q_1$$

and $P(q_1)$ is maximized when $q_1 = (a-c)/2b$.

The second firm will produce $q_2 = \mathrm{BR}_2(q_1) = (a-c)/2b - q_1/2 = (a-c)/2b - (a-c)/4b = (a-c)/4b$, for a total production of $3(a-c)/4b$ which is higher than the Cournot Duopoly production of $2(a-c)/3b$.

The profit of the first firm is $\left(a - \frac{3}{4}(a-c) - c\right)\frac{a-c}{2b} = \frac{(a-c)^2}{4b}$ which is higher than Cournot profit of $(a-c)^2/9b$.

Moving first gives an advantage.

Solving strictly competitive games with no chance moves

We now study more familiar games, similar to and chess and checkers.

In these games two players take turns making moves, until the game comes to an end and the outcome of the game is announced: either player I wins or player II wins or there is a draw.

(We shall call the player who moves first "player I" and the other player "player II".)

Definition. A *strictly competitive* game is a 2-player game in which for any two outcomes (u_1, u_2) and (u_1', u_2'), $u_1 > u_1'$ implies $u_2 < u_2'$.

Since the preferences of player II in a strictly competitive game can be inferred from the preferences of player I, we can, and shall, keep track only of the outcomes of player I, e.g., the outcomes in a game of chess can be thought as being "player I wins", "the game is drawn", and "player I loses".

Subgames and subtrees

DEFINITION. Let T be the tree associated with a sequential game G. For any vertex v in T we can construct a rooted subtree T' of T whose root is v, whose set of vertices V' consist of all vertices in T which can be reached by a directed path starting at v, and whose edges are all edges in T connecting vertices in V'. Such a subtree defines in turn a new game, the *subgame of G starting at v*.

Example: 2-pile Nim

There are two piles of stones, Alice and Bob take turns to make moves, each move consists of choosing a pile and taking way any number of stones from it. The person who takes the last stone wins.

Suppose that the piles have initially 5 and 8 stones and Alice goes first. Show that she has a strategy which ensures her victory.

Suppose that the piles have initially 7 stones each and Alice goes first. Show that Bob has a strategy which guarantees him victory.

(Hint: analyze the game which starts with one stone in each pile.)

Zermelo's theorem

DEFINITION. A sequential game is *finite* if its associated tree is finite.

THEOREM 13 (Zermelo's theorem). *Consider a finite, sequential game (with no chance moves) and let $\mathcal{S}$ the set of outcomes of the game. For any $S \subseteq \mathcal{S}$ either*

(a) *player I can force an outcome in S, or*

(b) *player II can force an outcome not in S.*

PROOF. Let T be the tree of this game; we proceed by induction on the rank r of T. If $r = 1$, T consist of a root and leaves only and either player I can choose a leaf with outcome in S or she is forced to choose a leaf with outcome not in S.

Assume now that $r > 1$ and that the result holds for all games with trees of rank less than r. Let $G_1, \ldots, G_n$ be the subgames resulting after player I makes her move, and let $T_1, \ldots, T_n$ be the corresponding trees. Note that the ranks of $T_1, \ldots, T_n$ are less than r, and we may apply the induction hypothesis to each of $G_1, \ldots, G_n$. We have two cases: either

(i): Player II can force an outcome not in S in each one of the games $G_1, \ldots, G_n$, or

(ii): for some G_i, player II cannot force an outcome not in S.

If (i) holds, conclusion (b) follows, while if (ii) holds player I has a strategy which starts with a move to game G_i which forces an outcome in S. □

Strictly competitive, finite, sequential games have a value

Let the outcomes of such game be $u_1 < u_2 < \cdots < u_k$ where $<$ denotes player I's preference. For any $1 \le i \le k$, let $W_i = \{u_i, \ldots, u_k\}$ and $L_i = \{u_1, \ldots, u_i\}$.

There exists a largest $1 \le i \le k$ for which player I can force an outcome in W_i;
player I cannot force an outcome in W_{i+1}, so player II can force an outcome in L_i.

We conclude that player I has a strategy that guarantees her at least u_i and player II has a strategy that guarantees him player I will not do better than u_i.

Chess-like games

COROLLARY 14. *In chess, either white can force a win, or black can force a win, or both sides can force at least a draw.*

PROOF. Here $\mathcal{S} = \{\text{white wins}, \text{draw}, \text{black wins}\}$. Apply Zermelo's Theorem to $S_1 = \{\text{white wins}\}$ and $S_2 = \{\text{white wins, draw}\}$ and deduce that

(a): either white can force a win or black can force a win or a draw, and

(b): either black can force a win or white can force a win or a draw.

So we either have a strategy which guarantees victory by one side, or, if this fails, both sides have a strategy that guarantees at least a draw. ☐

Chomp!

A game of Chomp! starts with a rectangular array of pieces and two players alternate taking. A turn consists of a player choosing a piece and removing that piece and all other pieces above or to the right of the chosen piece. The player who takes the last piece loses the game.

THEOREM 15. *The player who goes first has a strategy which guarantees her victory.*

PROOF. We use the following *strategy-stealing argument*.

Assume that the Theorem fails; Zermelo's theorem implies that the second player has a strategy which guarantees him victory. If the first player chooses the upper-right piece, there is a choice of piece P for the second player which is the first step in a winning strategy. Let the first player then start with a choice of P, and let him follow the winning strategy with the roles reversed! ☐

Although it is known that there exists a winning strategy for Player I in Chomp!, it is not known what that strategy is.

(You can play this game in `http://www.math.ucla.edu/ tom/Games/chomp.html` `http://www.math.ucla.edu/ tom/Ga`

Sequential games in strategic form

So far we described sequential games as rooted trees T where

(a) vertices which were not leaves were labeled with players,
(b) arrows originating from a vertex were labeled with actions actions available to the player corresponding to the vertex,
(c) arrows either pointed to leaves labeled with payoffs, or to other vertices as in (a)

DEFINITION. A *strategy* for player i in the sequential game described by the rooted tree T is a function which takes any vertex v labeled i to an action labeled on an arrow originating from v.

Example

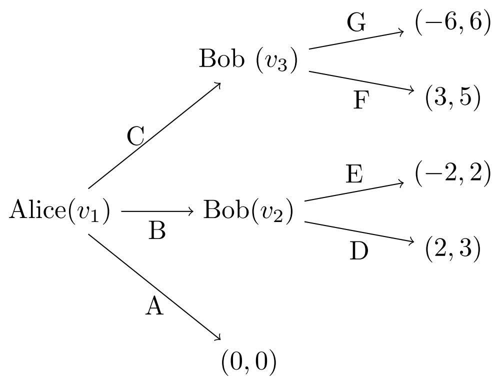

Alice has three strategies: $[v_1 \to A]$, $[v_1 \to B]$, and $[v_1 \to C]$. Bob has four strategies strategies $[v_2 \to D, v_3 \to F]$, $[v_2 \to D, v_3 \to G]$, $[v_2 \to E, v_3 \to F]$, $[v_2 \to E, v_3 \to G]$.

We can now give the strategic form of this game as follows

	$[D,F]$	$[D,G]$	$[E,F]$	$[E,G]$
A	0, 0	0, 0	0, 0	0, 0
B	2, 3	2, 3	-2, 2	-2, 2
C	3, 5	-6, 6	3, 5	-6, 6

Note that we have a pure-strategy Nash equilibrium $(B, [D,G])$, corresponding to our previous solution. There is also a new Nash equilibrium $(A, [E,G])$.

Example: A bargaining process

Alice and Bob will share £1, and Alice starts by issuing a take-it-or-leave-it offer of $0 \le s \le 1$. If Bob accepts, Alice gets $1-s$, Bob gets s, otherwise both get 0. We apply backward induction and see that Bob should accept any offer of $s > 0$. (This is not observed in the real world!)

We change the game, so that if Bob rejects the offer, the game is played again: now £δ is shared and Bob makes a take-it-or-leave it offer of $0 \le t \le \delta$. Backward induction now shows Alice should offer δ, and the payoffs are $(1-\delta, \delta)$.

Assume now that there are three iterations, in which £$1, \delta, \delta^2$ are shared in the 1st, 2nd and third iterations. Alice will offer (almost) 0 in the last iteration, and receive δ^2. Now Bob has to offer Alice δ^2 in the second iteration and he receives $\delta(1-\delta)$. Now Alice has to offer Bob $\delta(1-\delta)$ and she receives $1 - \delta(1-\delta) = 1 - \delta + \delta^2$.

What if this process can go on forever? Now payoffs are $(1/(1+\delta), \delta/(1+\delta))$, which for δ close to 1, is close to $(1/2, 1/2)$.
What if Alice and Bob have different discount factors δ_1 and δ_2?

Monopolist vs. new entrant

Bob has a monopoly in the market of wireless widgets, and Alice is considering competing with Bob. If Alice does not enter the market, her payoff is zero, and Bob makes a three billion pound profit. If Alice enters the market, Bob can either fight her off by reducing prices so that he makes zero profit and Alice loses a billion pounds, or Bob can choose not to fight her and they both end up making a profit of a billion pounds. Here is the tree describing this sequential game

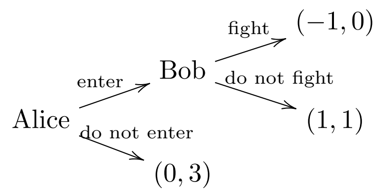

We use backward induction to solve this game an we see that the optimal strategies are for Alice to enter the market, and for Bob not to fight her off.

Consider the strategic form of this game

	fight	do not fight
enter	-1, 0	1, 1
do not enter	0, 3	0, 3

We have two Nash equilibria at (enter, do not fight) and at (do not enter, fight). The first corresponds to the solution above, what is the significance of the second? For Bob choosing to fight amounts to announcing that, no matter what Alice does, he will fight her off, but this is not credible as we know that it would be foolish for Bob to fight Alice if she enters the market.

His strategic choice is *not credible.* The situation could be different if the game were to repeat, e.g., if Bob would have to face multiple potential competitors: now he may choose to fight to set an example to the others.

So in sequential games, strategic Nash equilibria do not necessarily correspond to feasible choice of strategy.

Imperfect information and information sets

We now deal with games in which some of the actions of players are not known to the other players: these are *games of imperfect information.*

For example, any game in strategic form is a game of imperfect information: players do not know the other players actions until the game is finished.

Definition. Let G be a sequential game represented by a tree T. An *information set* for player i is a set of vertices V such that

(a) each vertex in V is labeled i,
(b) the set of arrows starting at each vertex in V is identical.

We partition all vertices which are not leaves into information sets, i.e., every such vertex is in precisely one information set.

A sequential game with partial information has as its strategies for player i the set of all functions which take an *information set* to an arrow starting at a vertex in that information set.

Note that in a game of perfect information all information sets contain one vertex.

An information set of a sequential game is a set of vertices belonging to one player which are indistinguishable to that player.

Example: matching pennies as a sequential game.

Example: Alice and Bob compete, Part I

Alice Ltd. is considering entering a new market in Freedonia dominated by Bob plc. If Alice enters the market, Bob can decide to resist or accommodate Alice's entrance. When Alice enters the market, she can choose to do so aggressively (e.g., lots of advertising) or not. If at the time of

Alice's second decision, she knows Bob's choice of strategy, we have a sequential game described by the following tree.

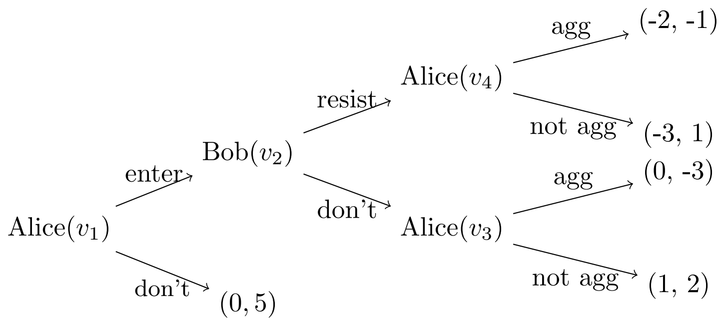

The backward induction solution of this game is straightforward: Alice will enter, Bob will not resist and Alice will not advertise aggressively.

The normal from of this game is as follows:

	resist	don't resist
enter, aggressive at v_3, aggressive at v_4	(-2, -1)	(0, -3)
enter, not aggressive at v_3, aggressive at v_4	(-2, -1)	(1,2)
enter, aggressive at v_3, not aggressive at v_4	(-3, -1)	(0, -3)
enter, not aggressive at v_3, not aggressive at v_4	(-3, -1)	(1, 2)
don't enter, aggressive at v_3, aggressive at v_4	(0,5)	(0,5)
don't enter, not aggressive at v_3, aggressive at v_4	(0,5)	(0,5)
don't enter, aggressive at v_3, not aggressive at v_4	(0,5)	(0,5)
don't enter, not aggressive at v_3, not aggressive at v_4	(0,5)	(0,5)

Example: Alice and Bob compete, Part II

Assume now that when Alice enters the market she does not know whether Bob decided to resist or not; Alice makes her second decision without knowing whether she is in vertex v_3 or v_4. This implies that any of its strategies has the same values on v_3 or v_4, and we can express this formally by saying that $\{v_3, v_4\}$ is an information set of Alice.

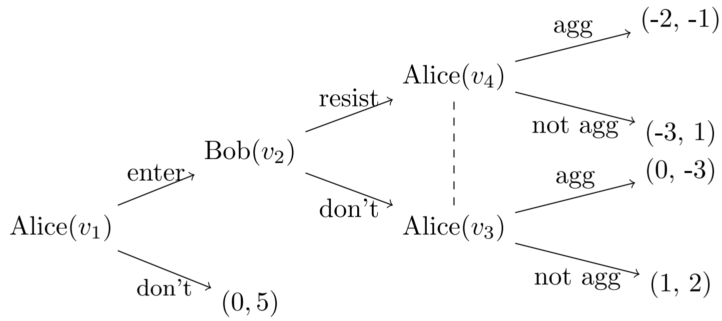

The normal form of this modified game is as follows:

	resist	don't resist
enter, aggressive at v_3, aggressive at v_4	(-2, -1)	(0, -3)
enter, not aggressive at v_3, not aggressive at v_4	(-3, -1)	(1, 2)
don't enter, aggressive at v_3, aggressive at v_4	(0,5)	(0,5)
don't enter, not aggressive at v_3, not aggressive at v_4	(0,5)	(0,5)

We cant solve this game with backward induction anymore, but we can find its Nash equilibria: any strategy profile where Alice stays out is a Nash equilibrium, and there is an additional one: ((enter, not aggressive at v_3, not aggressive at v_4) , dont resist).

An example

Consider the following game with perfect information.

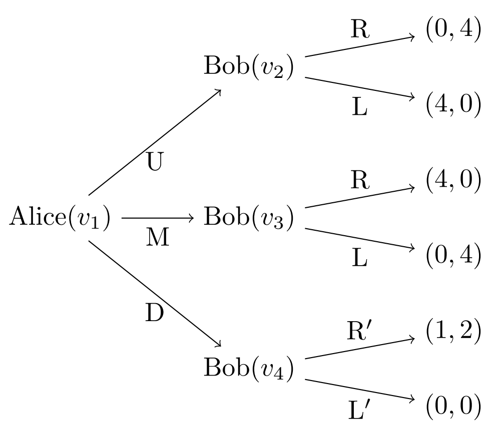

We easily solve this game using backward induction and obtain a solution $([D], [R'])$.

We now change the game and put vertices v_2 and v_3 in one information set: now Bob doesn't distinguish between these two vertices. Now Alice can mix strategies U and M with equal probabilities to produce an expected payoff of $2 > 1$. Formally, the game in strategic form is

	L L′	L R′	R L′	R R′
U	4, 0	4, 0	0, 4	0, 4
M	0, 4	0, 4	4, 0	4, 0
D	0, 0	1, 2	0, 0	1, 2

Alice's mixed strategy $(1/2, 1/2, 0)$ dominates strategy D and we reduce to the game

	L L′	L R′	R L′	R R′
U	4, 0	4, 0	0, 4	0, 4
M	0, 4	0, 4	4, 0	4, 0

	L L′	L R′	R L′	R R′
U	4, 0	4, 0	0, 4	0, 4
M	0, 4	0, 4	4, 0	4, 0

Now the first two and last two columns are indistinguishable, and if we ignore duplication we end up with the game

	L	R
U	4, 0	0, 4
M	0, 4	4, 0

which has mixed-strategy Nash equilibrium $((1/2, 1/2), (1/2, 1/2))$.

Subgame perfect Nash equilibria

Definition. Consider a sequential game with perfect information. A strategy profile which is a Nash equilibrium of the game is *subgame perfect* if its restriction to any subgame is also a Nash equilibrium.

Example: Market monopolist and new entrant, revisited.

Recall the game

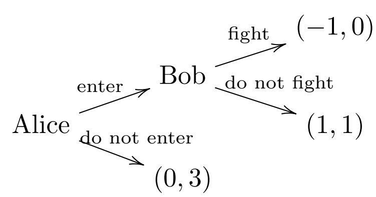

whose strategic form is

	fight	do not fight
enter	-1, 0	1, 1
do not enter	0, 3	0, 3

The game has two Nash equilibria: (*enter*, *do not fight*), (*do not enter*, *fight*). There is only one subgame to consider here, the one involving Bob's decision to fight or not, and clearly, *not fight* is the only Nash equilibrium. So (*enter*, *do not fight*) is a subgame perfect Nash equilibrium, and (*do not enter*, *fight*) is not.

Theorem 16. *The backward induction solutions of a finite game of perfect information are subgame perfect Nash equilibria.*

Proof. Let T be the tree of this game; we proceed by induction on the rank r of T. Call the player who moves at the root of T player I. If $r = 1$, T consist of a root and leaves only and player I will choose the strategy with highest payoff. This is clearly a subgame perfect Nash equilibrium.

Assume now that $r > 1$ and that the result holds for all games with trees of rank less than r. Let $G_1, \ldots, G_n$ be the subgames resulting after player I makes her move, and let $T_1, \ldots, T_n$ be the corresponding trees. Note that the ranks of $T_1, \ldots, T_n$ are less than r, and if we apply the induction hypothesis to each of $G_1, \ldots, G_n$ we obtain backward induction solutions of these which are subgame perfect Nash equilibria resulting in payoffs $u_1, \ldots, u_n$ for player I. Now player I needs to choose to move to one of the subgames $G_1, \ldots, G_n$, and she will choose the one with highest payoff for her. This results in a Nash equilibrium strategy. Note that a proper subgame of G is a subgame of one of $G_1, \ldots, G_n$ so the strategy is subgame perfect, because its restriction to $G_1, \ldots, G_n$ is subgame perfect and so its restriction to any subgame is a Nash equilibrium. $\square$

To define subgame perfect solutions of games without perfect information we need to modify the concept of a subgame.

Definition. Let G be a sequential game represented by a tree T. A *subgame* of G is a game corresponding to a subtree of T' of T with the property that every information set is entirely in T' or entirely outside T'.

Example

Consider the following game.

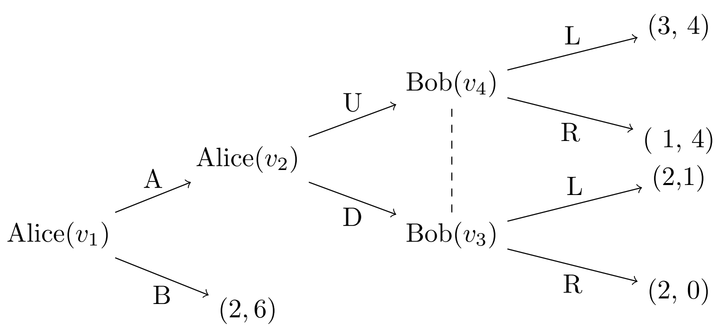

$(\{v_3, v_4\}$ is an information set)

This game has two subgames: the game itself, and the game corresponding to the subtree with root v_2. The proper subgame has strategic form:

	L	R
U	3, 4	1, 4
D	2, 1	2, 0

and has a Nash-equilibrium (U, L) resulting in payoff $(3, 4)$. Now at the root of the game Alice would choose to move A because $3 > 2$. The resulting Nash Equilibrium $(v_1 \to A, v_2 \to U,\ v_3 \to, v_4 \to L)$ is subgame perfect.

CHAPTER 7

Repeated Games

Example: Repeated Prisoners' Dilemma

Consider the following version of Prisoners' Dilemma

	n	c
N	2, 2	-1, 3
C	3, -1	0, 0

How do we model playing Prisoners' Dilemma twice?
Let $S_A = \{N, C\}$ and $S_B = \{n, c\}$.

If the game is repeated n times, Alice and Bob have

$$2^{4^0} \times 2^{4^1} \times \dots 2^{4^{n-1}}$$

strategies available to them, and for $n = 3$ this is already quite big ($2^{21} = 2,097,152$).

In each round of the game there are 4 possible strategy profiles $(N, n), (N, c), (C, n), (C, c)$. When Alice makes her choice in the ith stage of the game, there are 4^{i-1} possible histories, and she has to choose between N and C for each of them, and there are $K_i = 2^{4^{i-1}}$ possible choices. The total number of choices is

$$K_1 \times K_2 \times \dots K_n = 2^{4^0} \times 2^{4^1} \times \dots 2^{4^{n-1}}.$$

Payoffs of infinitely repeated games

If the game is played finitely many times, we can add up the payoffs of each stage.

What if the game is played infinitely many times, or potentially infinitely many times (e.g., after each repetition a coin is tossed and the game is stopped if Heads occurs.)

We'll use two methods for handling potentially infinite streams of payoffs. Given a sequence of intermediate payoffs $r_0, r_1, \dots$ we can define
the average payoff: $\lim_{m\to\infty}(r_0 + \dots + r_{m-1})/m$, and
the future discounted payoff: $r_0 + \beta r_1 + \dots$ where $0 \le \beta < 1$ (notice that this series converges!)

There are two main motivations behind the second definition: (a) the game repeated every fixed period of time, β is the *present value* of a payoff of 1 occurring one period of time into the future, (b) the game is repeated at each stage with probability β.

We refer to games with average reward as being *infinite repetitions,* and those which use the the future discounted reward as *indefinite repetitions.*

A finitely repeated game

Use backward induction for finite repetition.

Proposition 17. *Consider Prisoner's Dilemma repeated a finite number of times k and we add up the payoffs at each stage. There is a unique subgame-perfect Nash equilibrium in which the players play (C, c) at each stage.*

Proof. Induction on k. The case $k = 1$ follows from the elimination of the dominated strategies n and N.
Assume that $k > 1$ and that the statement holds for $k - 1$.

Consider the kth and last stage of the game: regardless of previous history, Alice and Bob find themselves playing the case $k = 1$ and hence they play the dominant strategies (C, c) and both receive a payoff of 0.
Since the kth stage of the game does not affect previous payoffs, there is no advantage in deviating from the Nash equilibria of the first $k - 1$ repetitions of the game.

This shorter game has repeated (C, c) as a unique Nash equilibrium. □

Nash equilibria of repeated games

Theorem 18 (1). *Let $G = (S, T, u_1, u_2)$ be a game with (pure- or mixed-strategy) Nash equilibrium (s, t). In the finite, infinite and indefinite repetitions of the game, playing (s, t) repeatedly is a subgame perfect Nash equilibrium.*

Proof. If the game is played a finite number of times k, and Alice plays $s_1, \dots, s_k$, against Bob's t, her payoff is $u_1(s_1, t) + \dots + u_1(s_k, t) \le u_1(s, t) + \dots + u_1(s, t)$. Similarly, Bob does not gain from deviating from t, and we conclude that playing (s, t) repeatedly is a Nash Equilibrium. This argument works for all subgames, so it is a subgame perfect Nash equilibrium.

Subgames of the infinite and indefinite repeated games are identical to the whole game, so the restriction of the strategy of repeated (s, t) gives the same strategy on all subgames. So if the proposed strategy is a Nash equilibrium, it is automatically subgame perfect.

Suppose that the first player deviates and plays s_k at the kth stage of the game. In the infinitely repeated game her payoff is

$$\lim_{k\to\infty} \frac{u_1(s_0, t) + \dots + u_1(s_{k-1}, t)}{k} \le \lim_{k\to\infty} \frac{u_1(s, t) + \dots + u_1(s, t)}{k}$$

and in the indefinite repeated game the payoff is $\displaystyle\sum_{k=0}^{\infty} \beta^k u_1(s_k, t) \le \sum_{k=0}^{\infty} \beta^k u_1(s, t)$. □

Consider an indefinitely repeated game of Prisoner's Dilemma where after each repetition the game is stopped with probability $1/3$ (and we measure payoffs with the the future discounted reward method with $\beta = 2/3$). The sets of strategies of Alice and Bob are now infinite; nevertheless we find some Nash equilibria.

Consider the following strategies:

HAWK: Always confess.

DOVE: Never confess.

GRIM: Don't confess until other person confesses, after that always confess.

Not surprisingly (HAWK, HAWK) is a (subgame-perfect) Nash equilibrium. (DOVE, DOVE) is not.

GRIM vs. GRIM

Proposition 19. *The strategy profile (GRIM, GRIM) is a (subgame-perfect) Nash equilibrium.*

PROOF. If both players play GRIM, their payoffs are $2 + 2\beta^1 + \cdots = 2/(1-\beta) = 6$. If Alice deviates from GRIM for the first time at the kth stage of the game, her payoff is at most

$$2(1+\beta+\cdots+\beta^{k-2}) + 3\beta^{k-1} = 2\frac{1-\beta^{k-1}}{1-\beta} + 3\beta^{k-1} =$$

$$6\left(1-\left(\frac{2}{3}\right)^{k-1}\right) + 3\left(\frac{2}{3}\right)^{k-1} = -3\left(\frac{2}{3}\right)^{k-1} + 6 < 6.$$

□

The "folk theorem"– first version (indefinite games)

Given a game $G = (S, T, u_1, u_2)$ and $0 < \beta < 1$ we define $G^\infty(\beta)$ to be the indefinitely repeated game G with discount factor β.

THEOREM 20 (2). *Let $G = (S, T, u_1, u_2)$ be a game with (pure- or mixed-strategy) Nash equilibrium (p, q). If (s, t) is a pure strategy profile with $u_1(s,t) > u_1(p,q)$ and $u_2(s,t) > u_2(p,q)$, then there exists a $0 \le \beta_0 < 1$ such that for all $\beta_0 \le \beta < 1$ there is subgame perfect Nash equilibrium of $G^\infty(\beta)$ with same payoff as that of playing (s,t) repeatedly.*

PROOF. Let $\mathcal{G}_1$ be the strategy for player 1 in which, if player 2 ever deviated from t she plays p and she plays s as long as player 2 sticks to t. Let $\mathcal{G}_2$ be the strategy for player 2 in which, if player 1 ever deviated from s he plays q and he plays t as long as player 1 sticks to s.

We first show that if β is close enough to 1, then $(\mathcal{G}_1, \mathcal{G}_2)$ is a Nash equilibrium for $G^\infty(\beta)$.

If $(\mathcal{G}_1, \mathcal{G}_2)$ is played, the players get a payoff of $\sum_{i=0}^\infty \beta^i u_1(s,t) = u_1(s,t)/(1-\beta)$ and $\sum_{i=0}^\infty \beta^i u_2(s,t) = u_2(s,t)/(1-\beta)$, respectively.

If player 1 deviates from playing s by playing s' at the k stage of the game and s_i $(i > k)$ thereafter, she gets the payoff

$$\begin{aligned}
&\sum_{i=0}^{k-1} \beta^i u_1(s,t) + \beta^k u_1(s',t) + \sum_{i=k+1}^{\infty} \beta^i u_1(s_i, q) \\
&\le \sum_{i=0}^{k-1} \beta^i u_1(s,t) + \beta^k u_1(s',t) + \sum_{i=k+1}^{\infty} \beta^i u_1(p, q) \\
&= \sum_{i=0}^{k-1} \beta^i u_1(s,t) + \beta^k u_1(s',t) + u_1(p,q)\frac{\beta^{k+1}}{1-\beta}
\end{aligned}$$

We need to consider values of β for which

$$\sum_{i=0}^{\infty} \beta^i u_1(s,t) > \sum_{i=0}^{k-1} \beta^i u_1(s,t) + \beta^k u_1(s',t) + u_1(p,q)\frac{\beta^{k+1}}{1-\beta}$$

i.e.,

$$u_1(s,t)\frac{\beta^k}{1-\beta} > \beta^k u_1(s',t) + u_1(p,q)\beta^{k+1}\frac{1}{1-\beta},$$

i.e.,

$$u_1(s,t)\frac{1}{1-\beta} > u_1(s',t) + u_1(p,q)\frac{\beta}{1-\beta}$$

which simplifies to $u_1(s,t) > (1-\beta)u_1(s',t) + \beta u_1(p,q)$. As $\beta \to 1$, the right hand side converges to $u_1(p,q) < u_1(s,t)$, hence a continuity argument shows that there exists a $0 \le \beta_1 < 1$ for which the inequality holds for all $\beta_1 < \beta < 1$.

A similar argument shows that there exists a $0 \leq \beta_2 < 1$ such that whenever $\beta_2 < \beta < 1$, player 2 would not deviate from playing $\mathcal{G}_2$. We deduce that with $\beta_0 = \max\{\beta_1, \beta_2\}$ and $\beta \in (\beta_0, 1)$, $(\mathcal{G}_1, \mathcal{G}_2)$ is a Nash equilibrium whose payoff is identical to the payoff of playing repeated (s, t) repeatedly.

To show that $(\mathcal{G}_1, \mathcal{G}_2)$ is a *subgame-perfect* Nash equilibrium we need to show that it is a Nash equilibrium of its subgames: such a subgame occurs either after a defection or not. After a defection $(\mathcal{G}_1, \mathcal{G}_2)$ calls for playing (p, q) in every repetition of the game, which is a Nash equilibrium by Theorem 1, and if no defection occurs the subgame is identical to the original game. □

Notation: G^∞

Given a game $G = (S, T, u_1, u_2)$ we define G^∞ to be the infinitely repeated game G with average reward payoff.

Minimax values

Definition. Given a finite game $G = (S, T, u_1, u_2)$ we define the *minimax values* of players 1 and 2 to be $\min_{t\in T} \max_{s\in S} u_1(s, t)$ and $\min_{s\in S} \max_{t\in T} u_2(s, t)$.

Thus the minimax value of a player is the worst possible payoff the other player can inflict.

Example. The minimax values in

	a	b	c
I	1, 0	6, 4	0, 9
II	2, 1	0, 2	3, 0
III	3, 7	2, 3	4, 0

are $3 = u_1(III, a)$ and $2 = u_2(II, b)$.

Convex sets

Definition. A *convex* set in $\mathbb{R}^d$ is a subset S of $\mathbb{R}^d$ with the property that for any $v, w \in S$, $tv + (1 - t)w \in S$ for all $0 \leq t \leq 1$.

Examples of convex sets: Δ_d, vector-subspaces of $\mathbb{R}^d$, cones, intersections of convex sets, etc.

Definition. The *convex hull* of a set $A \subseteq \mathbb{R}^d$ is, equivalently,

(a) the intersection of all convex sets containing A,
(b) the smallest convex set containing A,
(c) the subset C of $\mathbb{R}^d$ defined as the set of all points of the form $\lambda_1 a_1 + \cdots + \lambda_k a_k$ where $k \geq 1$, $\lambda_1, \ldots, \lambda_k$ are non-negative real numbers, $\lambda_1 + \cdots + \lambda_k = 1$ and $a_1, \ldots, a_k \in A$.

The cooperative payoff region of games

Definition. Let $G = (S, T, u_1, u_2)$ be a game. The *cooperative payoff region* of G is the convex hull of $\{(u_1(s, t), u_2(s, t)) \mid s \in S, t \in T\} \subseteq \mathbb{R}^2$.

Definition. Let $A, B \subseteq \mathbb{R}^n$. We say that A is dense in B if for all $\epsilon > 0$ and all $b \in B$ there is an $a \in A$ whose distance to b is less than ϵ.

Examples of dense sets: $\mathbb{Q}$ is dense in $\mathbb{R}$ and more generally $\mathbb{Q}^n$ is dense in any subset of $\mathbb{R}^n$.

The “folk theorem”– second version (infinite games)

THEOREM 21. *Let* $G = (S, T, u_1, u_2)$ *be a game with minimax values* $\mu_1 = u_1(\sigma_1, \tau_1)$ *and* $\mu_2 = u_2(\sigma_2, \tau_2)$. *Let* R *be the cooperative payoff region of* G *and let* $A = \{(x, y) \in R \,|\, x > \mu_1, y > \mu_2\}$. *Let* B *be the set of payoffs of Nash equilibria of* G^∞. *Then* B *is dense in* A.

PROOF. Consider $(x, y) \in A$ of the form $(x, y) = \lambda_1(u_1(s_1, t_1), u_2(s_1, t_1)) + \cdots + \lambda_k(u_1(s_k, t_k), u_2(s_k, t_k))$ where $k \geq 1$, $s_1, \ldots, s_k \in S$, $t_1, \ldots, t_k \in T$ and $\lambda_1, \ldots \lambda_k$ are non-negative *rational* numbers in $[0, 1]$ adding up to 1. The set of all such points is dense in A. Also, we can find a positive common denominator N such that $\lambda_1 = m_1/N, \ldots, \lambda_k = m_k/N$ for integers $m_1, \ldots, m_k$.

Let s be the strategy for player 1 which consists of playing s_1 for m_1 turns followed by s_2 for m_2 turns, etc., ending with playing s_k for m_k turns and repeating this pattern cyclically.

Let t be the strategy for player 2 which consists of playing t_1 for m_1 turns followed by t_2 for m_2 turns, etc., ending with playing t_k for m_k turns and repeating this pattern cyclically.

Let $\mathcal{G}_1$ be the strategy for player 1 in which, if player 2 ever deviated from t she plays σ_2 at every turn and she plays s as long as player 2 sticks to t. Let $\mathcal{G}_2$ be the strategy for player 2 in which, if player 1 ever deviated from s he plays τ_1 at every turn and he plays t as long as player 1 sticks to s.

We show $(\mathcal{G}_1, \mathcal{G}_2)$ is a Nash equilibrium for G^∞.

If $(\mathcal{G}_1, \mathcal{G}_2)$ is played, the players get a payoff of $\frac{1}{N} \sum_{i=1}^k m_i u_1(s_i, t_i) = x$ and $\frac{1}{N} \sum_{i=1}^k m_i u_2(s_i, t_i) = y$, respectively.

If player 1 deviates from playing s, player 2 punishes her and her subsequent payoffs are at most $\mu_1 < x$ and thus in the limit as the number of plays increases her average score converges to a limit below x.

If player 2 deviates from playing t, player 1 punishes him and his subsequent payoffs are at most $\mu_2 < y$ and thus in the limit as the number of plays increases his average score converges to a limit below y. $\square$

The proof above does not show that the dense set of payoffs occur as payoffs of *subgame perfect* Nash equilibria, but in fact they are.

CHAPTER 8

Bayesian Games

A Bayesian game is, roughly speaking, a game in which players do not have full information about the game they are playing.

Example: Alice and Bob face a showdown

PC Bob faces an armed suspect called Alice and they have to decide simultaneously whether to shoot or not. The suspect is either a criminal (with probability p) or not (with probability $1 - p$). The sheriff would rather shoot if the suspect shoots, but not otherwise. A criminal would rather shoot even if the sheriff does not. An innocent suspect would rather not shoot even if the sheriff shoots.

Bob doesn't know which of the following two games is being played

Game against innocent person

	shoot	dont shoot
shoot	-3, -1	-1, -2
dont shoot	-2, -1	0, 0

Game against criminal

	shoot	dont shoot
shoot	0, 0	2, -2
dont shoot	-2, -1	-1, 1

he knows that he is playing the first with probability $1 - p$ and the second with probability p.

What should Bob do? He should put himself in Alice's shoes: there is no uncertainty for her. In the first game her *dont shoot* strategy dominates her *shoot* strategy and in the second game her *shoot* strategy dominates her *dont shoot* strategy. So Bob is really playing the following games

Game against innocent person

	shoot	dont shoot
dont shoot	-2, -1	0, 0

Game against criminal

	shoot	dont shoot
shoot	0, 0	2, -2

Now if Bob shoots, his expected payoff is $-(1 - p)$ and $-2p$ if he doesnt. So he shoots when $-(1 - p) \geq -2p$, i.e., when $p \geq 1/3$.

Definition. A (two-player) Bayesian game consists of:

(a) a set of actions available to each player (S and T),
(b) a finite set of types for each player (Θ_A and Θ_B),
(c) a set of possible states of the world Ω, and functions $\tau_A : \Omega \to \Theta_A$ and $\tau_B : \Omega \to \Theta_B$.
(d) a payoff function for each player whose domain is $S \times T \times \Omega$,
(e) all probabilities $\mathrm{P}(\omega|\theta)$ for all $\theta \in \Theta_A \cup \Theta_B$ and $\omega \in \Omega$.

PC Bob revisited:

(a) Both Alice and Bob have set of actions $S = T = \{\text{shoot}, \text{don't shoot}\}$.
(b) $\Theta_B = \{\text{PC}\}$ and $\Theta_A = \{\text{innocent}, \text{criminal}\}$.
(c) $\Omega = \{\omega_1, \omega_2\}$, $\tau_B(\omega_1) = \tau_B(\omega_2) = \text{PC}$, $\tau_A(\omega_1) = \text{innocent}$, $\tau_A(\omega_2) = \text{criminal}$.
(d) $u_A(\text{shoot}, \text{shoot}, \omega_1) = -3$,
$u_B(\text{shoot}, \text{shoot}, \omega_1) = -1$, etc.

(e) $\mathrm{P}(\omega_1|\mathrm{PC}) = 1 - p$, $\mathrm{P}(\omega_2|\mathrm{PC}) = p$, $\mathrm{P}(\omega_1|\text{innocent}) = 1$, $\mathrm{P}(\omega_2|\text{innocent}) = 0$, $\mathrm{P}(\omega_1|\text{criminal}) = 0$, $\mathrm{P}(\omega_2|\text{criminal}) = 1$.

Strategies in Bayesian Games

Definition. A *strategy* for a player in a Bayesian game is a function from the set of its types to the set of its actions. Using our notation for Alice and Bob, a strategy for Alice is an element of S^{Θ_A} and a strategy for Bob is an element in T^{Θ_B}. A strategy profile is a choice of a strategy for each player, i.e., an element in $S^{\Theta_A} \times T^{\Theta_B}$.

Given a strategy profile $(s(-), t(-)) \in S^{\Theta_A} \times T^{\Theta_B}$, Alice's expected payoff when she has type $\theta_A \in \Theta_A$ is

$$\sum_{\omega \in \Omega} \mathrm{P}(\omega|\theta_A) u_A(s(\theta_A), t(\tau_B(\omega)), \omega)$$

and Bob's expected payoff when he has type $\theta_B \in \Theta_B$ is

$$\sum_{\omega \in \Omega} \mathrm{P}(\omega|\theta_B) u_B(s(\tau_A(\omega)), t(\theta_B), \omega).$$

Bayes-Nash equilibrium

Definition. The strategy profile $(s(-), t(-)) \in S^{\Theta_A} \times T^{\Theta_B}$ is a *Bayes-Nash equilibrium* if $s = s(\theta_A)$ maximizes

$$\sum_{\omega \in \Omega} \mathrm{P}(\omega|\theta_A) u_A(s, t(\tau_B(\omega)), \omega)$$

for all $\theta_A \in \Theta_A$ and $t = t(\theta_B)$ maximizes

$$\sum_{\omega \in \Omega} \mathrm{P}(\omega|\theta_B) u_B(s(\tau_A(\omega)), t, \omega)$$

for all $\theta_B \in \Theta_B$.

This definition says that a strategy profile is a Bayes-Nash equilibrium if *each type separately* would not change its action as specified in the corresponding strategy.

So for practical uses we can look for Bayes-Nash equilibria by replacing a given Bayesian game with a new game whose set of players are all the types in the original game.

Example: Alice and Bob have a dispute

Example. Alice has a dispute with Bob, who is either strong (S) or weak (W). Alice believes that Bob is strong with probability p. Each person can either fight or yield (henceforth abbreviated F and Y). The outcome of the confrontation is given as follows:

strong Bob

	F	Y
F	-1,1	1,0
Y	0,1	0,0

weak Bob

	F	Y
F	1,-1	1,0
Y	0,1	0,0

Here the sets of actions for both players is $\{F, Y\}$, $\Theta_A = \{A\}$ and $\Theta_B = \{S, W\}$ $\Omega = \{\omega_1, \omega_2\}$, $\tau_A(\omega_1) = \tau_A(\omega_2) = A$, $\tau_B(\omega_1) = S$, $\tau_B(\omega_2) = W$. Alice has two strategies: $(A \to F)$ and $(A \to Y)$. Bob has four strategies $(S \to F, W \to F)$, $(S \to F, W \to Y)$, $(S \to Y, W \to F)$, and $(S \to Y, W \to Y)$. We compute expected payoffs of all strategy profiles:

	$(S \to F, W \to F)$	$(S \to F, W \to Y)$
$(A \to F)$	1-2p, (1,-1)	1-2p,(1,0)
$(A \to Y)$	0, (1,1)	0, (1,0)

	$(S \to Y, W \to F)$	$(S \to Y, W \to Y)$
$(A \to F)$	1, (0,-1)	1, (0,0)
$(A \to Y)$	0, (0,1)	0, (0,0)

If $p < 1/2$, $1 - 2p > 0$ and the first row dominates the second, and we obtain a Bayes-Nash equilibrium at $((A \to F), (S \to F, W \to Y))$. If $p > 1/2$, $1 - 2p < 0$ and we obtain a Bayes-Nash equilibrium at $((A \to Y), (S \to F, W \to F))$.

Example: More information may hurt

Alice and Bob play one of the games below:

Game G_1

	L	M	R
U	1, 2x	1,0	1, 3x
D	2,2	0,0	0,3

Game G_2

	L	M	R
U	1, 2x	1,3x	1, 0
D	2,2	0,3	0,0

where $0 \leq x < 1/2$. Neither know which game is played, both assign probability 1/2 to either game. Here each player has one type, $\Omega = \{G_1, G_2\}$.

Strategies can be identified with actions and the expected payoffs of this Bayesian game are given by

	L	M	R
U	1, 2x	1,3x/2	1, 3x/2
D	2,2	0,3/2	0,3/2

Column L dominates the others and we obtain a Bayes-Nash equilibrium (D,L) which results in a payoff of 2 for Bob.

Game G_1

	L	M	R
U	1, 2x	1,0	1, 3x
D	2,2	0,0	0,3

Game G_2

	L	M	R
U	1, 2x	1,3x	1, 0
D	2,2	0,3	0,0

Suppose now that Bob knows which game he is playing: now Bob has two types, say B_1 and B_2, $\tau_B(G_1) = B_1$ and $\tau_B(G_2) = B_2$. Also, now Bob has 9 strategies, $(L, L), (L, M), \ldots$.

For B_1, strategy R is a dominant strategy, and for B_2, M is a dominant strategy, so we can eliminate strategy L. Now Alice strategy U dominates, and we end up in a Bayes-Nash equilibrium $(U, (R, M))$ which give Bob a payoff of $3x < 2$!