

Normal-Form Games and Mixed Nash Equilibria

Thomas Kesselheim

Last Update: April 13, 2025

In this lecture, we introduce the general framework of games. Congestion games, as introduced in the last lecture, are a special case. The notion of pure Nash equilibria readily generalizes but pure Nash equilibria might not exist. Therefore, we will introduce the concept of mixed Nash equilibria, which always exist in games with finitely many players and finitely many strategies.

1 Normal-Form Game

All games that we will consider throughout this course can be represented as *normal-form games*. We are following the standard notation, which is slightly different from the standard notation in congestion games. Normally, this does not cause major confusion.

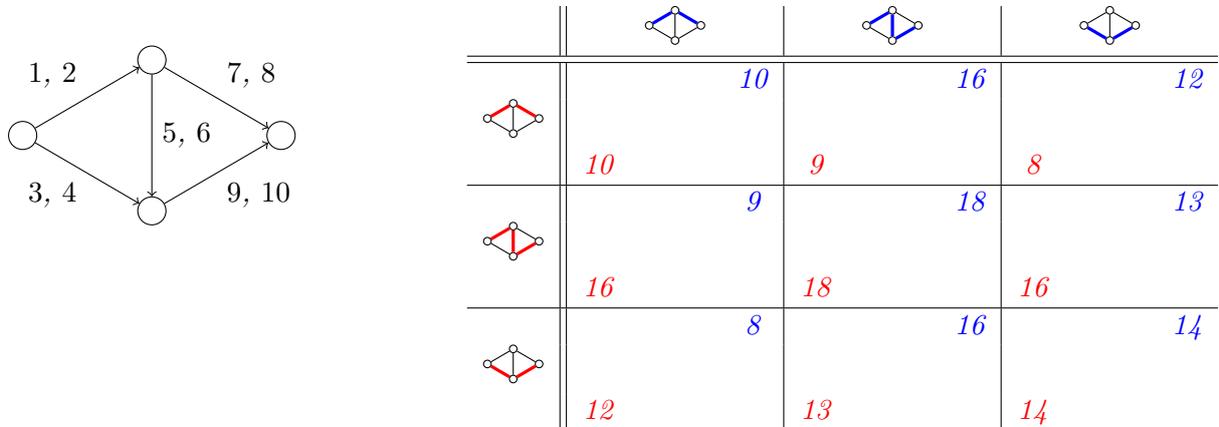
Definition 3.1. A (normal-form, cost-minimization) game is a triple $(N, (S_i)_{i \in N}, (c_i)_{i \in N})$. Here, N is the set of players, $|N| = n$, often $N = \{1, \dots, n\}$. For each player $i \in N$, S_i is the set of (pure) strategies of player i . The set $S = \prod_{i \in N} S_i$ is called the set of states or strategy profiles. For each $i \in N$, $c_i: S \rightarrow \mathbb{R}$ is the cost function of player i . In state $s \in S$, player i has a cost of $c_i(s)$.

Again, we let $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$ denote a state s without the strategy s_i . This notation allows us to concisely define a unilateral deviation of a player. For $i \in N$, let $s \in S$ and $s'_i \in S_i$, then $(s'_i, s_{-i}) = (s_1, \dots, s_{i-1}, s'_i, s_{i+1}, \dots, s_n)$.

One class of such games are congestion games. Here, the cost functions $c_i(s)$ have a particular structure: They are defined via resources and delay functions.

Any game with two players with finitely many strategies can be described by two matrices $A = (a_{s_1, s_2})_{s_1 \in S_1, s_2 \in S_2}$ and $B = (b_{s_1, s_2})_{s_1 \in S_1, s_2 \in S_2}$ (*bimatrix game*). Player 1 (referred to as row player) chooses a row; player 2 (column player) chooses a column. Their costs are given as $c_1(s) = a_{s_1, s_2}$, $c_2(s) = b_{s_1, s_2}$.

Example 3.2 (Congestion Game). A congestion game with two players can also be written as a bimatrix game, see an example of a symmetric network congestion game below with three s - t paths. As we see, the matrices will usually be huge.



Example 3.3 (Battle of the Sexes). Suppose Angelina and Brad go to the movies. Angelina prefers watching movie A, Brad prefers watching movie B. However, both prefer watching a movie together to watching movies separately.

	A	B
A	1 2	6 6
B	5 5	2 1

2 Pure Nash Equilibrium

The notion of best response and pure Nash equilibrium are defined exactly as in the case of congestion games.

Definition 3.4. A strategy s_i is called a best response for player $i \in N$ against a collection of strategies s_{-i} if $c_i(s_i, s_{-i}) \leq c_i(s'_i, s_{-i})$ for all $s'_i \in S_i$. A state $s \in S$ is called a pure Nash equilibrium if s_i is a best response against the other strategies s_{-i} for every player $i \in N$.

So, a pure Nash equilibrium is stable against unilateral deviation. No player can reduce his cost by only changing his only strategy.

Pure Nash equilibria need not be unique.

Example 3.5 (Battle of the Sexes). Recall the game from Example 3.3. We can find its pure Nash equilibria (A, A) and (B, B) by marking best responses with boxes.

	A	B
A	1 2	6 6
B	5 5	2 1

Not every game has a pure Nash equilibrium.

Example 3.6 (Inspection Game). Consider a game between a train traveler and a ticket inspector. The traveler can either (C)omply and buy a ticket before boarding the train or (D)efect and not buy one. The inspector can decide to be (L)azy and not inspect or she can decide to (I)nspect. The inspector prefers to inspect only if the traveller is not holding a ticket. The traveller, in contrast, prefers to buy a ticket only if she will be inspected.

	C	D
L	15 0	0 10
I	15 1	60 6

Having marked best responses, we see that there is no cell (corresponding to a pure state) in which both players are playing a best response simultaneously. Therefore, there is no pure Nash equilibrium.

3 Mixed Nash Equilibrium

We will now define a different equilibrium concept that generalizes pure Nash equilibria. Rather than making a fixed (deterministic) choice, the players can randomize over their strategies.

Definition 3.7. A mixed strategy σ_i for player i is a probability distribution over the set of pure strategies S_i .

We will only consider the case of finitely many pure strategies and finitely many players. In this case, we can write a mixed strategy σ_i as $(\sigma_{i,s_i})_{s_i \in S_i}$ with $\sum_{s_i \in S_i} \sigma_{i,s_i} = 1$. The cost of a mixed state σ for player i is defined as the *expected cost* if we draw for each player i a strategy according to distribution σ_i . Formally,

$$c_i(\sigma) = \mathbf{E}_{s_1 \sim \sigma_1, \dots, s_n \sim \sigma_n} [c_i(s)] = \sum_{s_1 \in S_1} \cdots \sum_{s_n \in S_n} \sigma_{1,s_1} \cdots \sigma_{n,s_n} \cdot c_i(s) .$$

Any pure strategy $s_i \in S_i$ can also be represented as a mixed strategy by setting $\sigma_{i,s_i} = 1$ and $\sigma_{i,s'_i} = 0$ for $s'_i \neq s_i$. Therefore, effectively, we have now defined a larger strategy space. On this space, best response and Nash equilibrium are just defined the same way.

Definition 3.8. A mixed strategy σ_i is a (mixed) best-response strategy against a collection of mixed strategies σ_{-i} if $c_i(\sigma_i, \sigma_{-i}) \leq c_i(\sigma'_i, \sigma_{-i})$ for all other mixed strategies σ'_i . A mixed state σ is called a mixed Nash equilibrium if σ_i is a best-response strategy against σ_{-i} for every player $i \in N$.

A key observation is that it is enough to only consider deviations to pure strategies.

Lemma 3.9. A mixed strategy σ_i is a best-response strategy against σ_{-i} if and only if $c_i(\sigma_i, \sigma_{-i}) \leq c_i(s'_i, \sigma_{-i})$ for all pure strategies $s'_i \in S_i$.

Proof. The “only if” part is trivial: Every pure strategy is also a mixed strategy.

For the “if” part, let σ_{-i} be an arbitrary mixed strategy profile for all players except for i . Furthermore, let σ_i be a mixed strategy for player i such that $c_i(\sigma_i, \sigma_{-i}) \leq c_i(s'_i, \sigma_{-i})$ for all pure strategies $s'_i \in S_i$.

Observe that for any mixed strategy σ'_i , we have $c_i(\sigma'_i, \sigma_{-i}) = \sum_{s'_i \in S_i} \sigma'_{i,s'_i} c_i(s'_i, \sigma_{-i}) \geq \min_{s'_i \in S_i} c_i(s'_i, \sigma_{-i})$. Using $\min_{s'_i \in S_i} c_i(s'_i, \sigma_{-i}) \geq c_i(\sigma_i, \sigma_{-i})$, we are done. \square

This lemma has a very important consequence: If a pure strategy s_i is a best response against pure strategies s_{-i} in the sense of Definition 3.4 then it is also a best response against s_{-i} in the sense of Definition 3.8. So, it is well defined if we simply talk about “best responses”. For this reason, **every pure Nash equilibrium is also a mixed Nash equilibrium.**

4 Mixed Best Responses as Probability Distributions over Pure Best Responses

Let us now see an example of a mixed Nash equilibrium and how we can find it. We will make use of the following lemma.

Lemma 3.10. A mixed strategy σ_i is a best-response strategy against σ_{-i} if and only if every strategy in the support of σ_i , i.e., every $s_j \in S_i$ with $\sigma_{i,s_j} > 0$, is a best response against σ_{-i} .

Proof. First suppose σ_i is a distribution over pure best responses. Then for every pure strategy $s'_i \in S_i$ we have

$$c_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_{i,s_i} c_i(s_i, \sigma_{-i}) \leq \sum_{s_i \in S_i} \sigma_{i,s_i} c_i(s'_i, \sigma_{-i}) = c_i(s'_i, \sigma_{-i}).$$

Now we can invoke Lemma 3.9 to conclude that σ_i is a best response to σ_{-i} .

Next suppose σ_i is *not* a distribution over pure best responses. Then there exists strategies $s_i, s'_i \in S_i$ such that $\sigma_{i,s_i} > 0$ but $c_i(s_i, \sigma_{-i}) > c_i(s'_i, \sigma_{-i})$. We construct distribution σ'_i from σ_i by increasing σ_{i,s'_i} by σ_{i,s_i} and decreasing σ_{i,s_i} to zero. Then,

$$c_i(\sigma'_i, \sigma_{-i}) < c_i(\sigma_i, \sigma_{-i}),$$

and so σ_i cannot be a best response. □

As a consequence of this lemma, we can compute mixed Nash equilibria by choosing probabilities for one player that will make the other player indifferent between his pure strategies.

Example 3.11 (Inspection Game). *Recall the game from Example 3.6.*

		C	D	
L		15	0	1 - p
	I	0	10	p
		15	60	
		1	6	
		1 - q	q	

We consider the case that both p and q are in $(0, 1)$, meaning that both players mix both their strategies. It is easy to see that otherwise the equilibrium property cannot hold. The lemma tells us to compute probabilities $(1 - p, p)$ for the row player and $(1 - q, q)$ for the column player that make the respective other player indifferent between the two strategies.

To determine the probabilities of the column player, we compute the expected costs for the pure strategies of the row player, equate them, and solve for q :

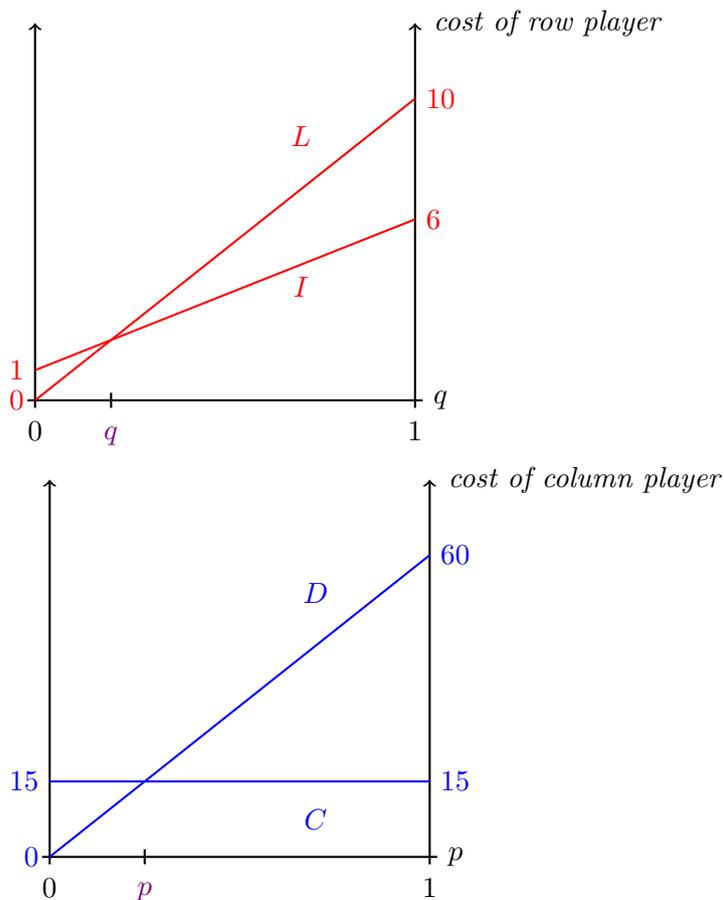
$$\begin{aligned} c_{\text{row}}(L, (1 - q, q)) &= c_{\text{row}}(I, (1 - q, q)) \\ \Leftrightarrow 0 \cdot (1 - q) + 10 \cdot q &= 1 \cdot (1 - q) + 6 \cdot q \\ \Leftrightarrow q &= \frac{1}{5}. \end{aligned}$$

Similarly, to determine the probabilities for the row player:

$$\begin{aligned} c_{\text{col}}(C, (1 - p, p)) &= c_{\text{col}}(D, (1 - p, p)) \\ \Leftrightarrow 15 \cdot (1 - p) + 15 \cdot p &= 0 \cdot (1 - p) + 60 \cdot p \\ \Leftrightarrow p &= \frac{1}{4}. \end{aligned}$$

We obtain the mixed Nash equilibrium in which the row player mixes between L and I with probabilities $(3/4, 1/4)$ and the column player mixes between C and D with probabilities $(4/5, 1/5)$.

There is also a nice pictorial interpretation of this argument. For each player, the expected cost of a strategy depends on the choice of probabilities of the respective other player. We find the mixed Nash equilibrium at the respective intersection of both lines.



One may wonder at this point whether it is actually reasonable to assume that player roll dice when making their decisions. An equivalent interpretation is: There is a population of inspectors, 25% of them inspect, 75% do not. Among the travelers, 80% buy a ticket, 20% do not. Now, an inspector and a traveler are matched at random. On average, they are happy with their choice.

Acknowledgments

Parts of these lecture notes are based on earlier versions by Berthold Vöcking, Martin Hoefer and Paul Dütting.

References and Further Reading

- Philip D. Straffin. Game Theory and Strategy, The Mathematical Association of America, fifth printing, 2004. (For basic concepts)