

# Networks and Markets — Lecture 2: Games

Nikhil Garg, Cornell Tech

**Reading: Games (Kleinberg & Easley, Ch. 6)**

“*All models are wrong. Some models are useful!*” — George Box

“*The map is not the territory.*” — Borges

## 1 Informal analysis of the mechanism from last lecture

Last lecture discussed ways to allocate *paper-discussion day slots*. Abstractly, there is a set of days  $D$ , each day  $d \in D$  has a capacity  $c_d$  (number of available slots), and each person/team wants one slot. Each person  $i$  has a preference ranking  $\succ_i$  over the days in  $D$ .

A *mechanism* asks people to report information (a *strategy*), and then uses those reports to produce an allocation. We say a mechanism is *strategy-proof* if truth-telling is a (weakly) dominant strategy: no matter what others report, you cannot benefit from misreporting your preferences.

### 1.1 Attempt 1: Random selection

- Mechanism: ignore preferences and assign slots uniformly at random.
- Pros: “fair” ex ante (symmetry); trivially strategy-proof (there is nothing you can say to affect the outcome).
- Con: preferences are ignored, so the outcome can be very inefficient.

### 1.2 Attempt 2: Report a single favorite day (serial dictatorship with one choice)

Each person reports only a single day  $s_i \in D$  (intended to be their favorite day).

- Randomize people into an order (lottery numbers).
- When it is your turn, you get  $s_i$  if it still has capacity; otherwise you are *skipped*.
- At the end, remaining slots are assigned uniformly at random among the skipped people.

For the purpose of analysis, suppose that everyone knows their lottery numbers before reporting their favorite day.

**Pros.** This takes some preferences into account (at least the top choice), and the lottery order gives a notion of “fair-ish” access to scarce days.

This mechanism is *not* strategy-proof.

**A 3-person counterexample.** Consider three people and three days  $\{M, T, W\}$ , each with one slot. Suppose the lottery order is 1, 2, 3, and true preferences are:

person	1st	2nd	3rd
1	$M$	$T$	$W$
2	$M$	$T$	$W$
3	$M$	$W$	$T$

If everyone truthfully reports their favorite day ( $s_1 = s_2 = s_3 = M$ ), then person 1 takes  $M$ , and people 2 and 3 are both skipped. The remaining days  $\{T, W\}$  are then assigned uniformly at random, so person 2 gets  $T$  with probability  $1/2$  and  $W$  with probability  $1/2$ .

But person 2 can do better by *lying*: report  $s_2 = T$ . Then person 2 receives  $T$  for sure (since it is still available when person 2 is processed). Thus, truthful reporting is not a dominant strategy in Attempt 2.

### 1.3 Attempt 3: Report a full preference list (serial dictatorship with full rankings)

Each person reports a full ranking over days. The mechanism:

- Randomize people into an order (lottery numbers).
- When it is your turn, assign you the highest-ranked day on your list that still has capacity.

**Pros.** This uses preference information more fully than Attempt 2, while keeping the same “fair-ish” lottery symmetry.

**Why Attempt 3 is strategy-proof (informal).** Fix the lottery order and fix everyone else’s reported rankings. When person  $i$  is reached, the set of remaining days with available capacity is already determined by earlier people; person  $i$ ’s report cannot change which days remain available.

The mechanism then gives person  $i$  the *top-ranked* remaining day according to their submitted list. Therefore, to get the best possible outcome given the remaining set, person  $i$  should rank those remaining days in the same order as their true preferences, i.e., report truthfully. Any misreport can only cause the mechanism to select a day that is weakly worse (in  $i$ ’s true ranking) than the best remaining day.

## 2 Games

A *game* specifies:

1. participants (players),
2. actions available to each participant,
3. rewards (payoffs) as a function of the chosen actions.

### 2.1 Warm-up (single decision maker)

Suppose there are 30 lecture days in a semester. Each student can choose to attend  $a$  days.

### 2.1.1 Model 1: Fixed benefit per lecture attended

The overall benefit is proportional to the fraction of lectures attended, i.e.,  $\frac{10a}{30}$ . (Think of the benefit as learning benefit, or attendance grade). Suppose also there is a cost  $c$  per lecture attended (e.g., opportunity cost of time). And the net reward is benefit minus cost.

In formalism, we have one player (the student). Let the action be

$$a \in \{0, 1, \dots, 30\}$$

with reward

$$r(a) = \left(\frac{10a}{30}\right) - ca = \frac{10}{30}a - ca = a\left(\frac{10}{30} - c\right),$$

where  $c$  is a constant. Then, the optimal action that maximizes the reward is

$$a^* = \begin{cases} 30 & \text{if } c < \frac{10}{30}, \\ 0 & \text{if } c > \frac{10}{30}, \\ \text{any value in } \{0, \dots, 30\} & \text{if } c = \frac{10}{30}. \end{cases}$$

### 2.1.2 Model 2: Decreasing marginal returns

Now, suppose that at some point, attending more lectures yields diminishing returns in learning. For example, the benefit could be proportional to the square root of the fraction of lectures attended, i.e.,  $10\left(\frac{a}{30}\right)^{\frac{1}{2}}$ .

Then, the reward is

$$r(a) = 10\left(\frac{a}{30}\right)^{1/2} - ca$$

To find the optimal action, we can take the derivative of  $r(a)$  with respect to  $a$  and set it to zero:

$$r'(a) = \frac{10}{2 \cdot 30^{1/2}} a^{-1/2} - c = 0$$

Solving for  $a$ , we get

$$a^* = \left(\frac{10}{2c \cdot 30^{1/2}}\right)^2 = \frac{100}{4c^2 \cdot 30} = \frac{25}{30c^2} = \frac{5}{6c^2}$$

Thus, the optimal number of lectures to attend depends on the cost  $c$ . For example, if  $c^2 = \frac{5}{60}$ , then  $a^* = 10$  lectures.

## 2.2 Two-player game: “Exam” vs “Presentation”

There is a group of two students (Person 1 and Person 2) who must each choose between two activities to prepare for: studying for an exam (which is individual) or preparing a presentation (which is joint). The payoffs depend on the combination of choices made by both students.

Each player chooses one of two actions:

- $E$ : study for the exam,
- $P$ : prepare the presentation.

If they both prepare for the presentation, they each get 100 points on the presentation but only 80 on the exam, for an overall grade of 90.

If they both study for the exam, they each get 92 on the exam but only 84 on the presentation, for an overall grade of 88.

If one studies for the exam while the other prepares the presentation, they both get 92 on the presentation. The one who studied for the exam gets 92 on the exam, while the one who prepared the presentation gets 80 on the exam, for overall grades of 92 and 86, respectively.

Payoffs (*Person 1*, *Person 2*):

	2 : <i>E</i>	2 : <i>P</i>
1 : <i>E</i>	(88, 88)	(92, 86)
1 : <i>P</i>	(86, 92)	(90, 90)

**Best responses / dominant strategies.** If Person 1 plays *E*, Person 2 prefers *E* over *P* ( $88 > 86$ ). If Person 1 plays *P*, Person 2 prefers *E* over *P* ( $92 > 90$ ). Thus *E* is a *dominant strategy* for Person 2. By symmetry, *E* is also dominant for Person 1.

**Nash equilibrium (pure).** Thus, the unique pure Nash equilibrium is  $(E, E)$  with payoff (88, 88). Neither player can improve their payoff by unilaterally changing their action.

**Pareto Dominance.** An outcome  $x$  *Pareto dominates* outcome  $y$  if both players are at least as well off in  $x$  as in  $y$ , and at least one player is strictly better off in  $x$  than in  $y$ . An outcome is *Pareto optimal* if there is no other outcome that Pareto dominates it.

In this game, the following outcomes are Pareto optimal:

- $(P, P)$  with payoff (90, 90).
- $(E, P)$  with payoff (92, 86).
- $(P, E)$  with payoff (86, 92).

The only outcome that is not Pareto optimal is the Nash equilibrium  $(E, E)$  with payoff (88, 88), since it is Pareto dominated by  $(P, P)$ .

**Social optimality.** The socially optimal outcome is the one that maximizes total payoff. Here, the total payoffs are:

- $(E, E)$ :  $88 + 88 = 176$ ,
- $(E, P)$ :  $92 + 86 = 178$ ,
- $(P, E)$ :  $86 + 92 = 178$ ,
- $(P, P)$ :  $90 + 90 = 180$ .

Thus, the socially optimal outcome is  $(P, P)$  with total payoff 180. However, note that this outcome is not a Nash equilibrium, since both players can improve their payoffs by unilaterally switching to *E*. This is true even though  $(P, P)$  is better for both players than  $(E, E)$ ! (The outcome Pareto dominates  $(E, E)$ .)

**Price of anarchy.** We have a socially optimal outcome  $(P, P)$  with total payoff 180, but the only Nash equilibrium is  $(E, E)$  with total payoff 176.

The price of anarchy measures how the efficiency of the worst Nash equilibrium compares to the social optimum.

Thus, price of anarchy (PoA) =  $\frac{\text{social optimum}}{\text{worst Nash equilibrium}} = \frac{180}{176} \approx 1.023$ .

### 2.3 Rock–Paper–Scissors and mixed Nash

So far, we have only seen pure strategy Nash equilibria. Next, we consider a game that has no pure strategy Nash equilibria, but has a mixed strategy Nash equilibrium, in which players randomize over their actions.

Payoffs (row player, column player), with win = 1 and loss = 0:

	$R$	$P$	$S$
$R$	(0, 0)	(0, 1)	(1, 0)
$P$	(1, 0)	(0, 0)	(0, 1)
$S$	(0, 1)	(1, 0)	(0, 0)

**No pure strategy Nash equilibria.** Suppose the row player plays  $R$ . Then, the column player prefers  $P$  over  $R$  and  $S$ . Suppose the row player plays  $P$ . Then, the column player prefers  $S$  over  $R$  and  $P$ . Suppose the row player plays  $S$ . Then, the column player prefers  $R$  over  $P$  and  $S$ . (Symmetrically, the same reasoning applies for the column player.) Thus, there is no pure strategy Nash equilibrium. For any pure strategy profile, at least one player can improve their payoff by switching actions.

**Mixed strategy Nash equilibrium.** Let the row player play  $R, P, S$  with probabilities  $p_R, p_P, p_S$ , respectively, and let the column player play  $R, P, S$  with probabilities  $q_R, q_P, q_S$ , respectively.

What are the conditions for a mixed strategy Nash equilibrium? Each player must be indifferent between their actions, given the other player's mixed strategy, so that no player can improve their expected payoff by changing their own mixed strategy.

For the row player, the expected payoffs for playing  $R, P, S$  are:

$$\text{Payoff}(R) = 0 \cdot q_R + 0 \cdot q_P + 1 \cdot q_S = q_S,$$

$$\text{Payoff}(P) = 1 \cdot q_R + 0 \cdot q_P + 0 \cdot q_S = q_R,$$

$$\text{Payoff}(S) = 0 \cdot q_R + 1 \cdot q_P + 0 \cdot q_S = q_P.$$

(Recall  $q_x$  is the probability that the column player plays action  $x$ .)

For the row player to be indifferent between  $R, P, S$ , we need

$$\text{Payoff}(R) = \text{Payoff}(P) = \text{Payoff}(S)$$

$$q_S = q_R = q_P.$$

Since  $q_R + q_P + q_S = 1$ , we have  $q_R = q_P = q_S = \frac{1}{3}$ .

By symmetry, the same reasoning applies for the column player, yielding  $p_R = p_P = p_S = \frac{1}{3}$ .