Algorithmic Game Theory and Applications

Introduction to the Course

The Team

Aris Filos-Ratsikas
course coordinator, lecturer

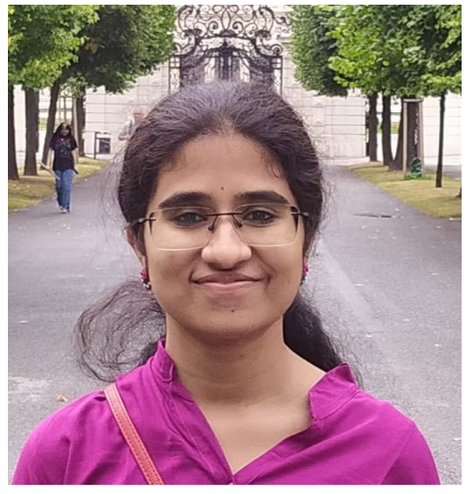

Sreedurga Gogulapati
lecturer

Charalampos Kokkalis
tutor

Lectures and Tutorials

- **Lectures:**

  Mondays 11.10 - 12.00, Weeks 1-10
  Basement Theatre - Adam House

  Thursdays 11.10 - 12.00, Weeks 1-10
  Forrest Hill Drill Hall

- **Tutorials:**

  Group 1
  Fridays 14.10 - 15.00, Weeks 3-10
  4.3. - Lister Learning and Teaching Centre

  Group 2
  Fridays 15.10 - 16.00, Weeks 3-10
  4.3. - Lister Learning and Teaching Centre

More about tutorials

More about tutorials

“Are the tutorials important? Should I attend them really?”

More about tutorials

“Are the tutorials important? Should I attend them really?”

Yes! They are the best preparation for the assignments and the exam.

Past students have reported that actively engaging with the tutorials was a huge plus for their final performance/mark.

More about tutorials

More about tutorials

“Should I prepare for the tutorials before attending, or is simply showing up enough?”

More about tutorials

“Should I prepare for the tutorials before attending, or is simply showing up enough?”

There will most likely be a **large** benefit from attempting the tutorial questions before attending.

More about tutorials

“Should I prepare for the tutorials before attending, or is simply showing up enough?”

There will most likely be a **large** benefit from attempting the tutorial questions before attending.

It’s always better to attend anyway!

Any coding?

Any coding?

“Is there any coding involved in this course?”

Any coding?

“Is there any coding involved in this course?”

Coding is not part of the course.

Any coding?

“Is there any coding involved in this course?”

Coding is not part of the course.

It might still be useful to do some coding to test the algorithms on real instances.

Any coding?

“Is there any coding involved in this course?”

Coding is not part of the course.

It might still be useful to do some coding to test the algorithms on real instances.

Still, the course is highly theoretical.

Pre-requisites

Pre-requisites

- Officially: No prerequisites.

Pre-requisites

- Officially: No prerequisites.
- But:

Pre-requisites

- Officially: No prerequisites.
- But:
  - It is recommended that you have taken and passed Algorithms and Data Structures (ADS).

Pre-requisites

- Officially: No prerequisites.
- But:
  - It is recommended that you have taken and passed Algorithms and Data Structures (ADS).
  - You should have some mathematical maturity, concretely some background in linear algebra, discrete mathematics, and probability theory.

Pre-requisites

- Officially: No prerequisites.
- But:
  - It is recommended that you have taken and passed Algorithms and Data Structures (ADS).
  - You should have some mathematical maturity, concretely some background in linear algebra, discrete mathematics, and probability theory.
  - A background in algorithms and some exposure to computational complexity (NP-completeness, etc, approximation algorithms) is desirable but not required.

Assessment

- Written Exam (80%)
- Coursework (20%)
  - Coursework 1 (10%):
    Released: 31/01/2025
    Due: 27/02/2025
  - Coursework 2 (10%)
    Released: 04/03/2025
    Due: 28/03/2025
  - Submission via Gradescope (via Learn).

Course textbooks

Course textbooks

- No formal textbook for the course.

Course textbooks

- No formal textbook for the course.
- References will be provided for each lecture.

Course textbooks

- No formal textbook for the course.

- References will be provided for each lecture.

- This is because different books cover different topics differently, some better than others.

Course textbooks

- No formal textbook for the course.
- References will be provided for each lecture.
- This is because different books cover different topics differently, some better than others.
- Still, these are some good references for algorithmic game theory.

Keeping in Contact

Keeping in Contact

- How to access the material:

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.
  - Learn page for assessment instructions and details, lecture recordings, and access to teaching tools.

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.
  - Learn page for assessment instructions and details, lecture recordings, and access to teaching tools.
  - Edstem forum for the class available as a direct link from Learn.

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.
  - Learn page for assessment instructions and details, lecture recordings, and access to teaching tools.
- Edstem forum for the class available as a direct link from Learn.
  - Edstem allows you to discuss and ask questions either in your own name, or under and anonymised “handle”.

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.
  - Learn page for assessment instructions and details, lecture recordings, and access to teaching tools.
- Edstem forum for the class available as a direct link from Learn.
  - Edstem allows you to discuss and ask questions either in your own name, or under and anonymised “handle”.
  - It has functionality to ask “private questions”, only visible to the lecturing team.

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.
  - Learn page for assessment instructions and details, lecture recordings, and access to teaching tools.
- Edstem forum for the class available as a direct link from Learn.
  - Edstem allows you to discuss and ask questions either in your own name, or under and anonymised “handle”.
  - It has functionality to ask “private questions”, only visible to the lecturing team.
  - We will keep a close eye on the forum, so please use it!

Keeping in Contact

- How to access the material:
  - Course webpage for material, slides, and general instructions.
  - Learn page for assessment instructions and details, lecture recordings, and access to teaching tools.
- Edstem forum for the class available as a direct link from Learn.
  - Edstem allows you to discuss and ask questions either in your own name, or under and anonymised “handle”.
  - It has functionality to ask “private questions”, only visible to the lecturing team.
  - We will keep a close eye on the forum, so please use it!
- Questions after the lectures are very much welcome!

What is Game Theory?

What is Game Theory?

Here's a vague definition:

What is Game Theory?

Here's a vague definition:

"Game Theory is the formal study of interaction between goal-oriented (or rational, or strategic) 'agents' (or 'players') and the strategic scenarios that arise in such settings".

What is Game Theory?

Here's a vague definition:

"Game Theory is the formal study of interaction between goal-oriented (or rational, or strategic) 'agents' (or 'players') and the strategic scenarios that arise in such settings".

Kousha Etessami

Example 1:
Auction of a video game

Example 1:
Auction of a video game

Scenario:

Example 1:
Auction of a video game

Scenario:

Your lecturer in AGTA bought NBA2K 25 and now he wants to sell NBA2K 24, which is of no longer use to him.

Example 1:
Auction of a video game

Scenario:

Your lecturer in AGTA bought NBA2K 25 and now he wants to sell NBA2K 24, which is of no longer use to him.

You are willing to spend up to a certain amount of money for this game, between £0 and £100. To figure out this amount, use the following link and generate a random number between 0 and 100.

https://www.calculator.net/random-number-generator.html

Remember this number and try to keep it hidden from the people around you!

Example 1:
Auction of a video game

Scenario:

Your lecturer in AGTA bought NBA2K 25 and now he wants to sell NBA2K 24, which is of no longer use to him.

You are willing to spend up to a certain amount of money for this game, between £0 and £100. To figure out this amount, use the following link and generate a random number between 0 and 100.

https://www.calculator.net/random-number-generator.html

Remember this number and try to keep it hidden from the people around you!

You are trying to maximise happiness, where

happiness = amount_willing_to_spend - amount_spent if you win and

happiness = 0 if you don't win.

Example 1:
Auction of a video game

Scenario:

Your lecturer in AGTA bought NBA2K 25 and now he wants to sell NBA2K 24, which is of no longer use to him.

You are willing to spend up to a certain amount of money for this game, between £0 and £100. To figure out this amount, use the following link and generate a random number between 0 and 100.

https://www.calculator.net/random-number-generator.html

Remember this number and try to keep it hidden from the people around you!

You are trying to maximise happiness, where

happiness = amount_willing_to_spend - amount_spent if you win and

happiness = 0 if you don't win.

Rules of the auction: The highest bid wins, the payment is that bid.

Example 2:
Setting the temperature

Example 2:
Setting the temperature

Scenario:

Example 2:
Setting the temperature

Scenario:

Your lecturer in AGTA decides that the temperature in the lecture room is too low (or too high) and would like to adjust it.

Example 2:
Setting the temperature

Scenario:

Your lecturer in AGTA decides that the temperature in the lecture room is too low (or too high) and would like to adjust it.

He is democratic, so he would like your input on which temperature to set the thermostat to.

Example 2:
Setting the temperature

Scenario:

Your lecturer in AGTA decides that the temperature in the lecture room is too low (or too high) and would like to adjust it.

He is democratic, so he would like your input on which temperature to set the thermostat to.

He will set the temperature to the average of the temperatures reported by the students.

Example 2:
Setting the temperature

Scenario:

Your lecturer in AGTA decides that the temperature in the lecture room is too low (or too high) and would like to adjust it.

He is democratic, so he would like your input on which temperature to set the thermostat to.

He will set the temperature to the average of the temperatures reported by the students.

You prefer it if it’s a little bit cold, let’s say at 17 degrees.

Example 2:
Setting the temperature

Scenario:

Your lecturer in AGTA decides that the temperature in the lecture room is too low (or too high) and would like to adjust it.

He is democratic, so he would like your input on which temperature to set the thermostat to.

He will set the temperature to the average of the temperatures reported by the students.

You prefer it if it's a little bit cold, let's say at 17 degrees.

What are you going to report to your lecturer as your proposed temperature? Why?

Example 2:
Setting the temperature

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

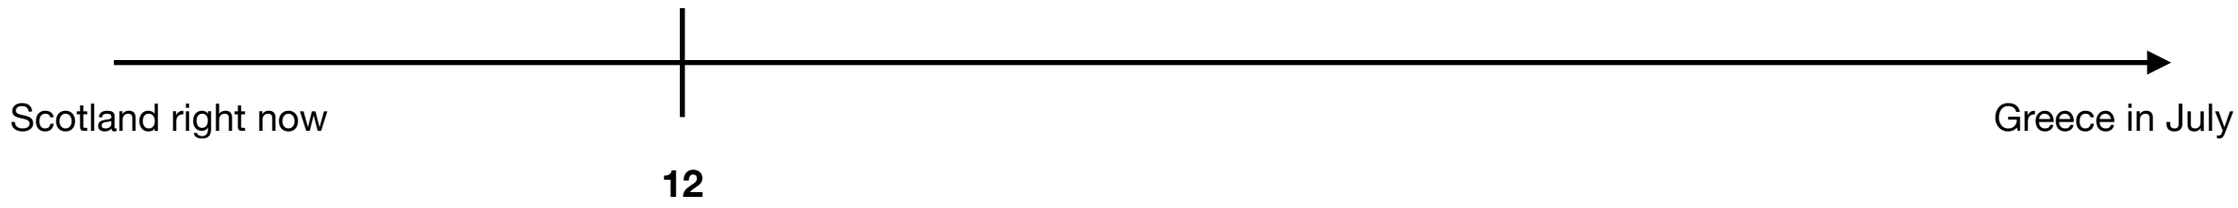

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

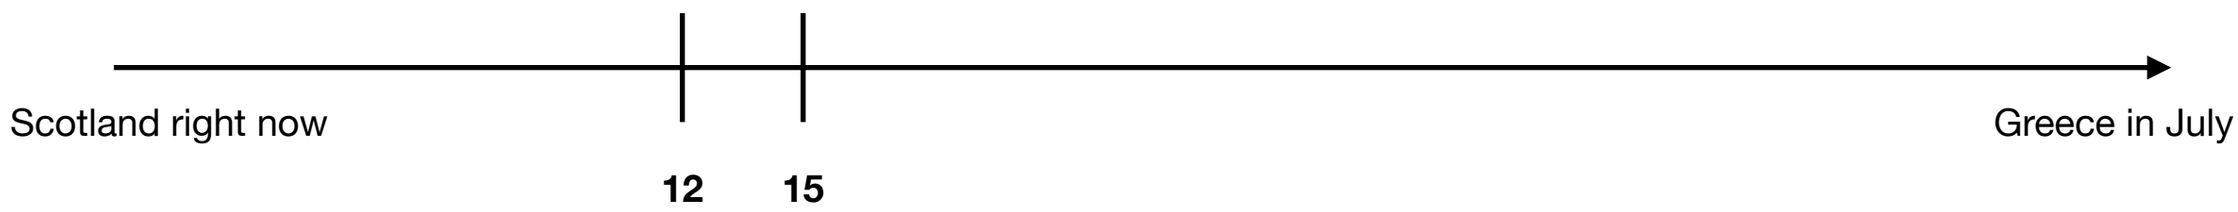

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

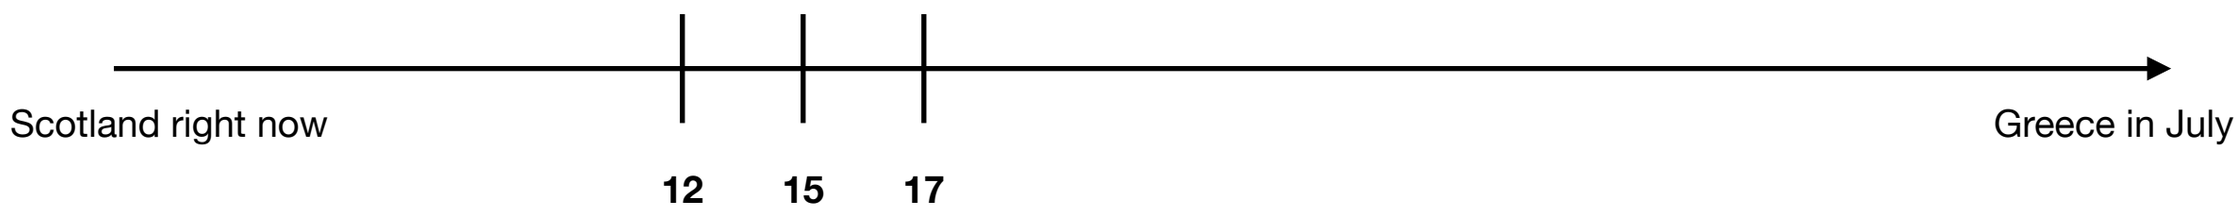

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

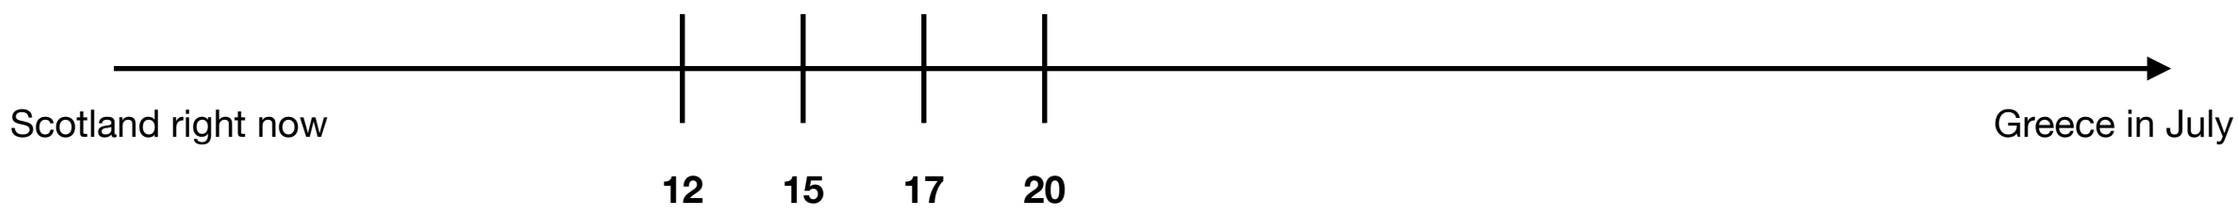

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

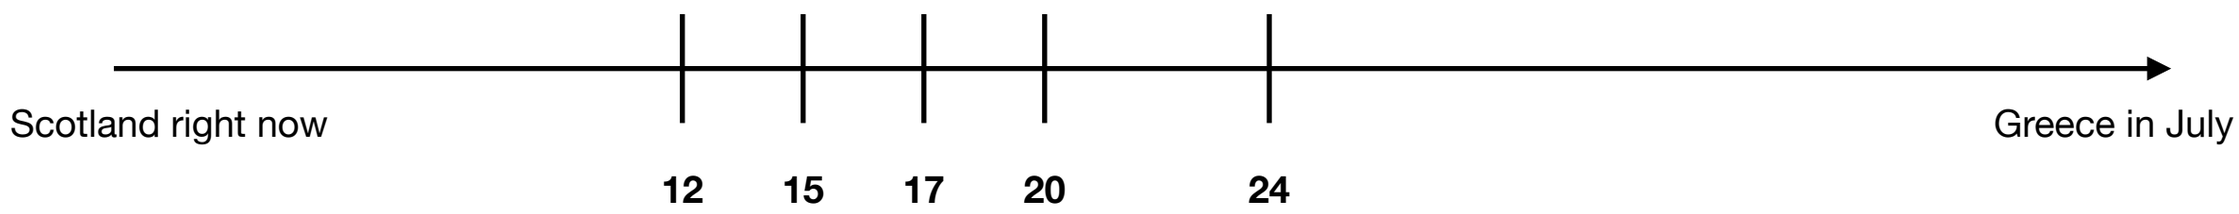

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

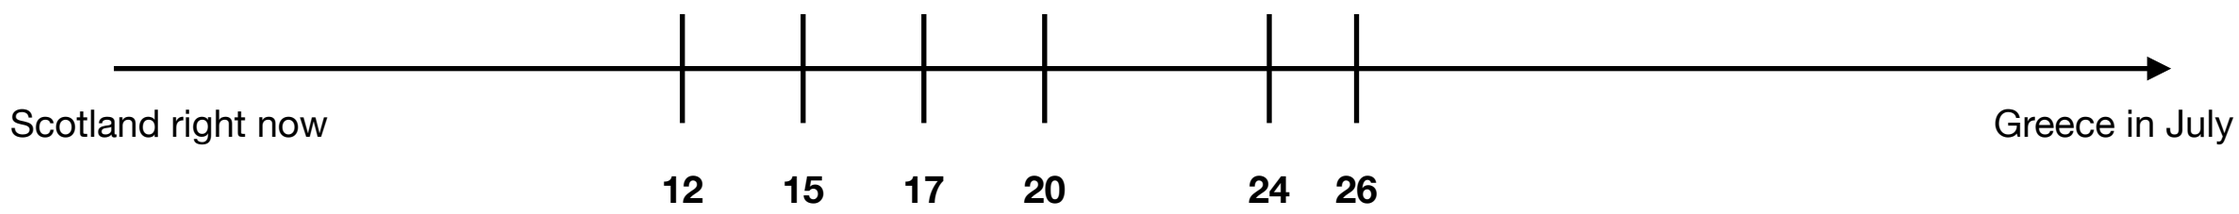

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

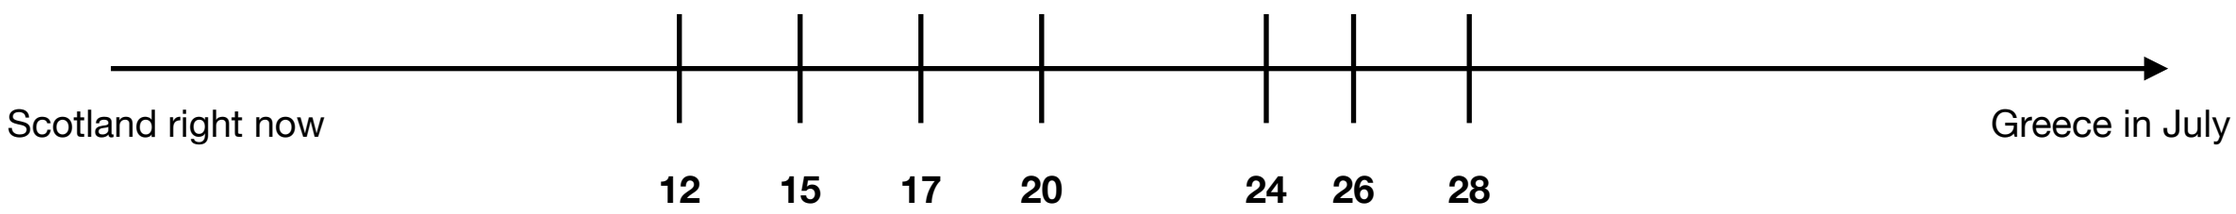

Example 2:
Setting the temperature

Scenario:

What if your lecturer instead chose the median temperature out of all the reports from the students?

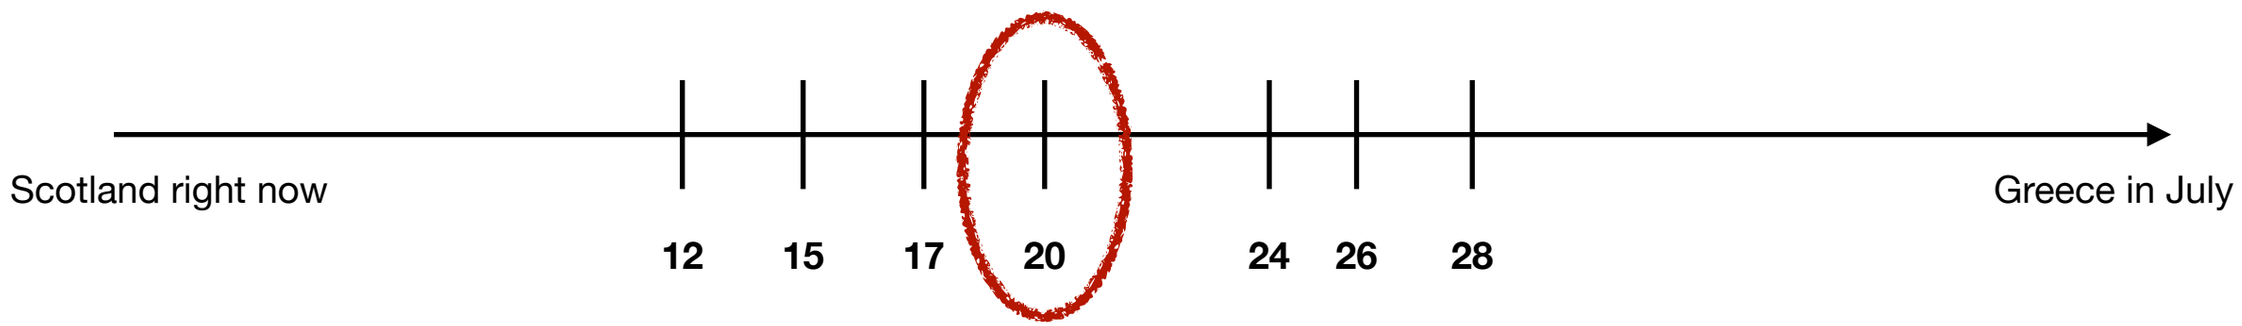

Example 3:
Cheating Partners

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

You and your classmate found the AGTA assignment impossible to do, and decided to bribe the course tutor to do it for you.

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

You and your classmate found the AGTA assignment impossible to do, and decided to bribe the course tutor to do it for you.

The tutor accepted the bribe (!) but he was also lazy, so he only provided only one set of solutions for both of you.

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

You and your classmate found the AGTA assignment impossible to do, and decided to bribe the course tutor to do it for you.

The tutor accepted the bribe (!) but he was also lazy, so he only provided only one set of solutions for both of you.

Unluckily your assignments were both marked by the same marker, who figured out that they are identical.

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

You and your classmate found the AGTA assignment impossible to do, and decided to bribe the course tutor to do it for you.

The tutor accepted the bribe (!) but he was also lazy, so he only provided only one set of solutions for both of you.

Unluckily your assignments were both marked by the same marker, who figured out that they are identical.

You were called to appear before the academic integrity officer, separately.

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

The officer explains the different outcomes:

Example 3:
Cheating Partners

<u>(Completely Hypothetical) Scenario:</u>

The officer explains the different outcomes:

- Paying to have your assignment done is a serious academic offence, and if proven, you would be expelled by the university.

Example 3:
Cheating Partners

<u>(Completely Hypothetical) Scenario:</u>

The officer explains the different outcomes:

- Paying to have your assignment done is a serious academic offence, and if proven, you would be expelled by the university.

- Otherwise, if it is deemed that you and your classmate collaborated or used ChatGPT, that would mean failing the course, but being allowed to repeat it next year.

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

The officer explains the different outcomes:

- Paying to have your assignment done is a serious academic offence, and if proven, you would be expelled by the university.

- Otherwise, if it is deemed that you and your classmate collaborated or used ChatGPT, that would mean failing the course, but being allowed to repeat it next year.

The officer strongly suspects that the first type of academic offence has taken place, but they cannot prove it. So they offer you the following deal:

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

The officer explains the different outcomes:

- Paying to have your assignment done is a serious academic offence, and if proven, you would be expelled by the university.
- Otherwise, if it is deemed that you and your classmate collaborated or used ChatGPT, that would mean failing the course, but being allowed to repeat it next year.

The officer strongly suspects that the first type of academic offence has taken place, but they cannot prove it. So they offer you the following deal:

- If you provide evidence that your classmate paid for the assignment, you will get a 0 for the assignment, but you can still continue with the course this year.

Example 3:
Cheating Partners

(Completely Hypothetical) Scenario:

The officer explains the different outcomes:

- Paying to have your assignment done is a serious academic offence, and if proven, you would be expelled by the university.

- Otherwise, if it is deemed that you and your classmate collaborated or used ChatGPT, that would mean failing the course, but being allowed to repeat it next year.

The officer strongly suspects that the first type of academic offence has taken place, but they cannot prove it. So they offer you the following deal:

- If you provide evidence that your classmate paid for the assignment, you will get a 0 for the assignment, but you can still continue with the course this year.

- However, your classmate has been offered the same deal. If they provide evidence against you, then you will not be expelled, but you will have to repeat the whole year (all courses).

Example 3: Cheating Partners

You \ Classmate	Take the deal	Don't take the deal
Take the deal	Repeat the whole year / Repeat the whole year	Get expelled / Fail the assignment
Don't take the deal	Fail the assignment / Get expelled	Fail the course / Fail the course

Example 3: Cheating Partners

Classmate / You	Take the deal	Don't take the deal
Take the deal	Repeat the whole year / Repeat the whole year	Get expelled / Fail the assignment
Don't take the deal	Fail the assignment / Get expelled	Fail the course / Fail the course

What would you do?

What is Game Theory?

Here's a vague definition:

"Game Theory is the formal study of interaction between goal-oriented (or rational, or strategic) 'agents' (or 'players') and the strategic scenarios that arise in such settings".

Kousha Etessami

What is Algorithmic Game Theory?

What is Algorithmic Game Theory?

Here's another vague definition:

What is Algorithmic Game Theory?

Here's another vague definition:

"Algorithmic Game Theory is concerned with computational and algorithmic questions related to game theory. In particular, can we design efficient algorithms for finding good ways to play (for 'solving games') or prove that such algorithms are unlikely to exist? Can we use algorithmic measures (like approximation) to measure how good the outcome of a game is?".

Some brief history

Some brief history

People have engaged in strategic games for centuries.

Some brief history

People have engaged in strategic games for centuries.

Examples of games date back to at least the 1700s.

Some brief history

People have engaged in strategic games for centuries.

Examples of games date back to at least the 1700s.

A comprehensive theory of strategic games was initiated by the work of John von Neumann (1928, 1944).

Some brief history

People have engaged in strategic games for centuries.

Examples of games date back to at least the 1700s.

A comprehensive theory of strategic games was initiated by the work of John von Neumann (1928, 1944).

A crucial contribution was made by John Nash in his PhD thesis in 1950.

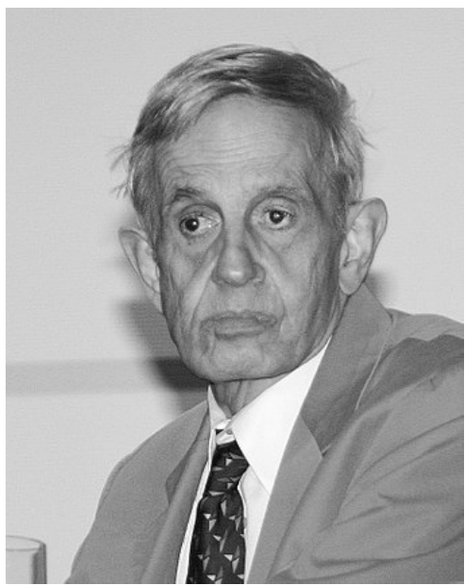

Some brief history

People have engaged in strategic games for centuries.

Examples of games date back to at least the 1700s.

A comprehensive theory of strategic games was initiated by the work of John von Neumann (1928, 1944).

A crucial contribution was made by John Nash in his PhD thesis in 1950.

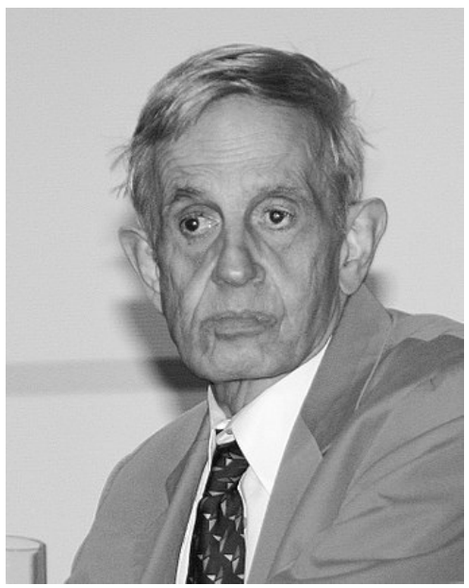

Many important results followed, which we will examine more closely as we go along.

Some brief history

People have engaged in strategic games for centuries.

Examples of games date back to at least the 1700s.

A comprehensive theory of strategic games was initiated by the work of John von Neumann (1928, 1944).

A crucial contribution was made by John Nash in his PhD thesis in 1950.

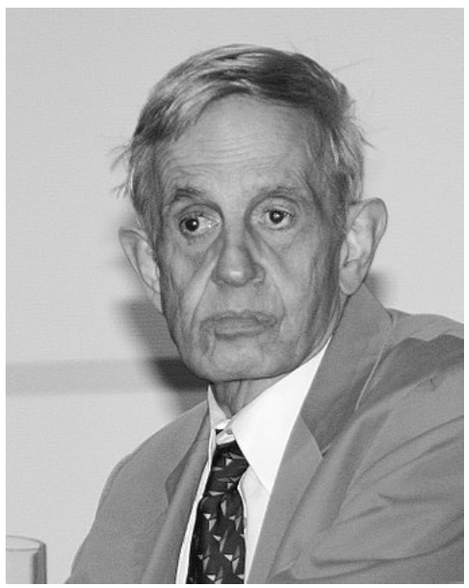

Many important results followed, which we will examine more closely as we go along.

Algorithmic Game Theory flourished in the mid to late 1990s.

Some brief history

People have engaged in strategic games for centuries.

Examples of games date back to at least the 1700s.

A comprehensive theory of strategic games was initiated by the work of John von Neumann (1928, 1944).

A crucial contribution was made by John Nash in his PhD thesis in 1950.

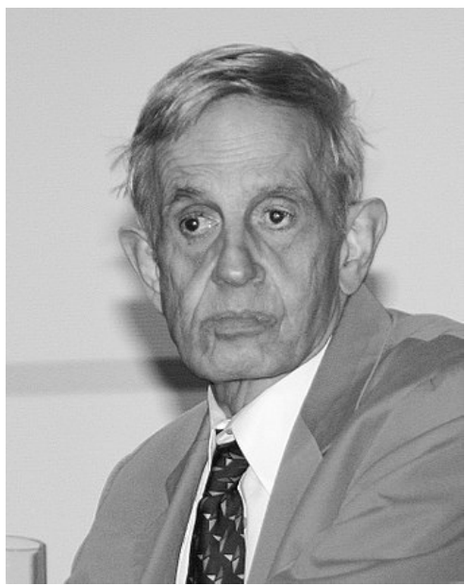

Many important results followed, which we will examine more closely as we go along.

Algorithmic Game Theory flourished in the mid to late 1990s.

It is still very much an active area of research, and as relevant for applications as ever!

The formal definition of a (strategic) game

Definition: A game in normal or strategic form is a tuple $(N, S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$ where

The formal definition of a (strategic) game

Definition: A game in normal or strategic form is a tuple $(N, S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$ where

1. $N = \{1,\ldots,n\}$ is a set of players (sometimes called "agents").

The formal definition of a (strategic) game

Definition: A game in normal or strategic form is a tuple $(N, S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$ where

1. $N = \{1,\ldots,n\}$ is a set of players (sometimes called "agents").

2. For each player $i \in N$, there is a set S_i of (pure) strategies.

The formal definition of a (strategic) game

Definition: A game in normal or strategic form is a tuple $(N, S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$ where

1. $N = \{1,\ldots, n\}$ is a set of players (sometimes called "agents").

2. For each player $i \in N$, there is a set S_i of (pure) strategies.

A vector $(s_1, s_2, \ldots, s_n) \in S_1 \times S_2 \times \ldots \times S_n = S$ is called a strategy profile.

The formal definition of a (strategic) game

Definition: A game in normal or strategic form is a tuple $(N, S_1, S_2, \ldots, S_n, u_1, u_2, \ldots, u_n)$ where

1. $N = \{1,\ldots, n\}$ is a set of players (sometimes called "agents").

2. For each player $i \in N$, there is a set S_i of (pure) strategies.

A vector $(s_1, s_2, \ldots, s_n) \in S_1 \times S_2 \times \ldots \times S_n = S$ is called a strategy profile.

3. For each player $i \in N$, there is a payoff (or utility) function $u_i : S \to \mathbb{R}$ which assigns a numerical value $u_i(s_1, s_2, \ldots, s_n)$ to player i for a given strategy profile $(s_1, s_2, \ldots, s_n)$.

Example 3: Cheating Partners

Classmate / You	Take the deal	Don't take the deal
Take the deal	Repeat the whole year / Repeat the whole year	Get expelled / Fail the assignment
Don't take the deal	Fail the assignment / Get expelled	Fail the course / Fail the course

What would you do?

Example 3: Cheating Partners

Classmate / You	Take the deal	Don't take the deal
Take the deal	5 / 5	-100 / 1
Don't take the deal	1 / -100	10 / 10

What would you do?

Mixed Strategies

Mixed Strategies

The goal of each agent $i \in N$ is to choose a strategy $s_i \in S_i$ which maximises the agent's payoff.

Mixed Strategies

The goal of each agent $i \in N$ is to choose a strategy $s_i \in S_i$ which maximises the agent's payoff.

This is if we are only considering pure strategies (no randomisation).

Mixed Strategies

The goal of each agent $i \in N$ is to choose a strategy $s_i \in S_i$ which maximises the agent's payoff.

This is if we are only considering pure strategies (no randomisation).

Sometimes it will make sense to consider randomised strategies, e.g.,

Mixed Strategies

The goal of each agent $i \in N$ is to choose a strategy $s_i \in S_i$ which maximises the agent's payoff.

This is if we are only considering pure strategies (no randomisation).

Sometimes it will make sense to consider randomised strategies, e.g.,

"Bid £33 with probability 1/2, and £4 with probability 1/2".

Mixed Strategies

The goal of each agent $i \in N$ is to choose a strategy $s_i \in S_i$ which maximises the agent's payoff.

This is if we are only considering pure strategies (no randomisation).

Sometimes it will make sense to consider randomised strategies, e.g.,

> *"Bid £33 with probability 1/2, and £4 with probability 1/2".*

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

Mixed Strategies

The goal of each agent $i \in N$ is to choose a strategy $s_i \in S_i$ which maximises the agent's payoff.

This is if we are only considering pure strategies (no randomisation).

Sometimes it will make sense to consider randomised strategies, e.g.,

> *"Bid £33 with probability 1/2, and £4 with probability 1/2".*

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

For example, if S_i is finite and contains m pure strategies, then $x_i = (x_{i1}, x_{i2}, \ldots, x_{im})$, where x_{ij} is the probability of playing pure strategy $s_j \in S_i$.

Support of Mixed Strategies

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

For example, if S_i is finite and contains m pure strategies, then $x_i = (x_{i1}, x_{i2}, \ldots, x_{im})$, where x_{ij} is the probability of playing pure strategy $s_j \in S_i$.

Support of Mixed Strategies

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

For example, if S_i is finite and contains m pure strategies, then $x_i = (x_{i1}, x_{i2}, \ldots, x_{im})$, where x_{ij} is the probability of playing pure strategy $s_j \in S_i$.

It could be the case that $x_{ij} = 0$ for some j. That means that the pure strategy $s_j \in S_i$ is played with 0 probability.

Support of Mixed Strategies

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

For example, if S_i is finite and contains m pure strategies, then $x_i = (x_{i1}, x_{i2}, \ldots, x_{im})$, where x_{ij} is the probability of playing pure strategy $s_j \in S_i$.

It could be the case that $x_{ij} = 0$ for some j. That means that the pure strategy $s_j \in S_i$ is played with 0 probability.

The support of a mixed strategy x_i is the set of all pure strategies that are played with strictly positive probability.

A little bit of geometry

A little bit of geometry

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

A little bit of geometry

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

Example 1: $S_i = \{s_1, s_2\}$

A little bit of geometry

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

Example 1: $S_i = \{s_1, s_2\}$

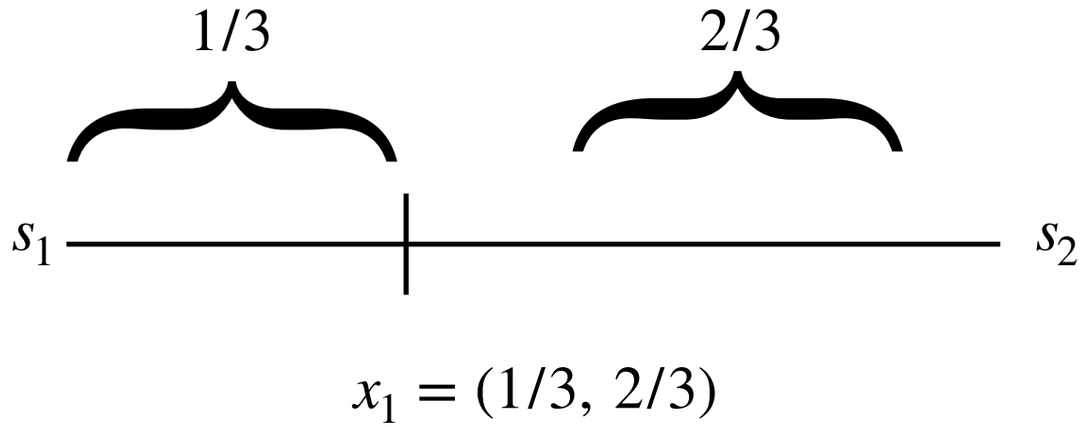

A little bit of geometry

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

Example 1: $S_i = \{s_1, s_2\}$

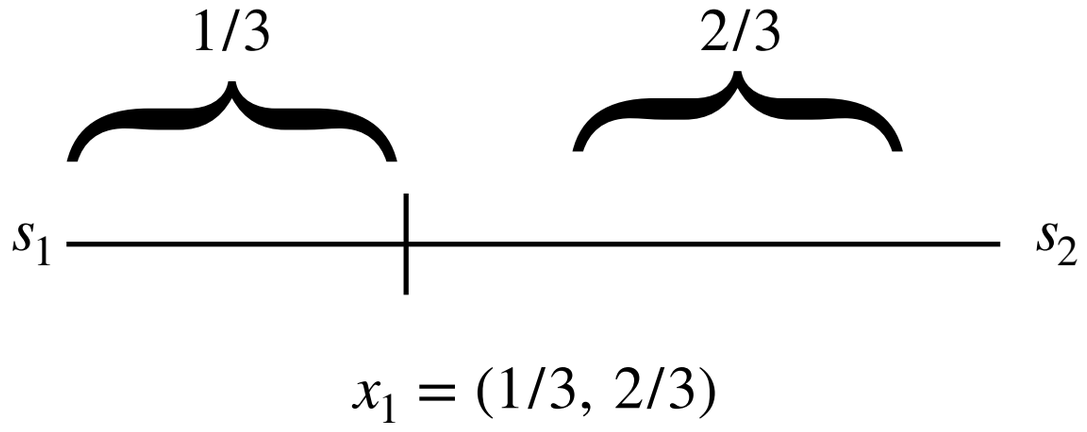

Example 2:

$S_i = \{s_1, s_2, s_3\}$

A little bit of geometry

A mixed strategy $x_i \in \Delta(S_i)$ is a probability distribution over the possible pure strategies in S_i.

Example 1: $S_i = \{s_1, s_2\}$

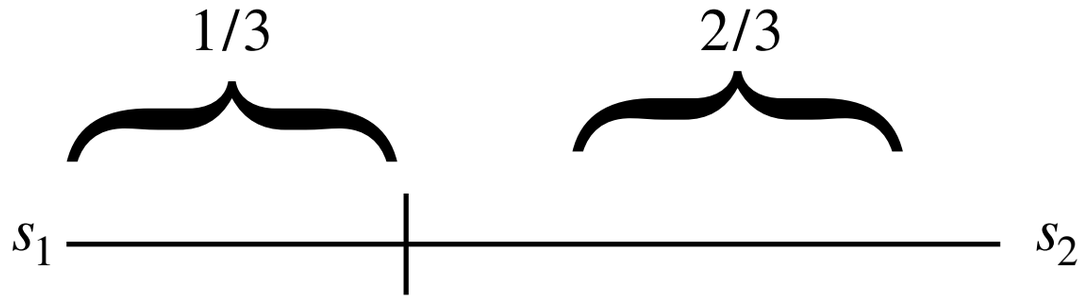

$x_1 = (1/3,\ 2/3)$

Example 2:

$S_i = \{s_1, s_2, s_3\}$

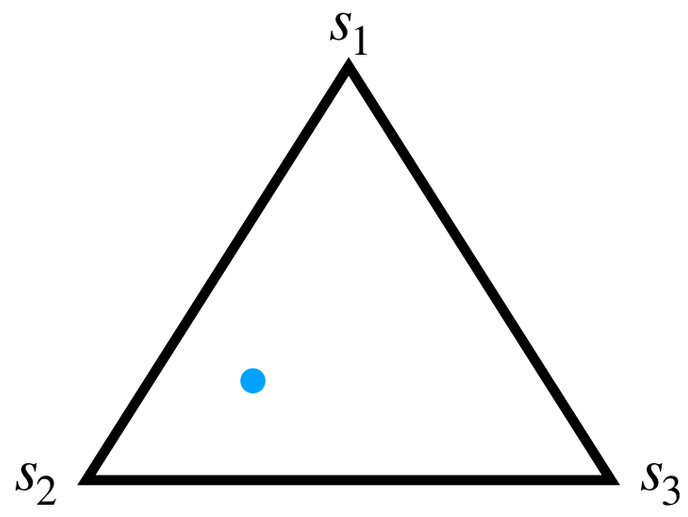

Expected Payoff/Utility

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

Expected Payoff/Utility

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

A mixed strategy profile $(x_1, x_2, \ldots, x_n)$ is a vector of mixed strategies, one for each player.

Expected Payoff/Utility

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

A mixed strategy profile $(x_1, x_2, \ldots, x_n)$ is a vector of mixed strategies, one for each player.

We can view a mixed strategy profile $(x_1, x_2, \ldots, x_n)$ as a distribution over pure strategy profiles $(s_1, s_2, \ldots, s_n)$.

Expected Payoff/Utility

A mixed strategy x_i is a probability distribution over the possible pure strategies in S_i.

A mixed strategy profile $(x_1, x_2, \ldots, x_n)$ is a vector of mixed strategies, one for each player.

We can view a mixed strategy profile $(x_1, x_2, \ldots, x_n)$ as a distribution over pure strategy profiles $(s_1, s_2, \ldots, s_n)$.

The expected utility of a player for a mixed strategy profile is

$$u_i[x_1, \ldots, x_n] = \mathbb{E}_{(s_1,\ldots,s_n)\sim(x_1,\ldots,x_n)}[u_i(s_1, \ldots, s_n)]$$

Example 3: Cheating Partners

Classmate / You	Take the deal	Don't take the deal
Take the deal	5 / 5	-100 / 1
Don't take the deal	1 / -100	10 / 10

Assume that we have the mixed strategies (1/2,1/2), (1/2,1/2)

Example 3: Cheating Partners

Classmate / You	Take the deal	Don't take the deal
Take the deal	5, 5	-100, 1
Don't take the deal	1, -100	10, 10

Assume that we have the mixed strategies (1/2,1/2), (1/2,1/2)

The expected utility of each player is

$$\frac{1}{4}\cdot 5+\frac{1}{4}\cdot 1+\frac{1}{4}\cdot(-100)+\frac{1}{4}\cdot 10$$