Algorithmic Game Theory and Applications

Bayesian Games and First-Price Auctions

First-price auctions (FPA)

Houses in Scotland are sold via *sealed-bid first-price auctions.*

Each bidder submits their bid independently, without seeing the bids of the other bidders.

The winner is the bidder with the highest bid.

> If there are multiple such bidders, one is chosen at random.

The winner needs to pay their bid, all other bidders do not pay anything.

First-Price Auction

There are n bidders from a set $N = \{1, \ldots, n\}$.

There is one item for sale.

Every bidder has a value v_i for the item - this is the bidder's willingness to buy it.

Each bidder chooses a bid $b_i = \beta(v_i)$ according to some function β.

Let $W = \{i : b_i \geq b_j, \forall j\}$ be the set of possible winners of the auction (those with the highest bid).

The utility of bidder i is

- $(v_i - b_i) \cdot \frac{1}{|W|}$ if $i \in W$.
- 0, otherwise.

How should you bid in the FPA?

"I should bid lower than the amount I am willing to spend, to win the item on sale for a smaller price."

"I shouldn't bid too low though, because that increases the chances of not winning the item at all."

"How low should I bid?"

Before we attempt to answer this question, let's ask another one first:

> Could we design a different auction that does not require us to engage in such considerations?
>
> i.e., can we define a *truthful* auction?

Auctions

Auction: A mechanism for buying or selling goods or services by means of eliciting bids from interested parties.

Classic example: Auction of a painting, or art in general.

Most prominent example nowadays: Ad auctions

> Selling advertising space (ad impressions) on online market places (ad exchanges).
>
> In 2022, this accounted for 58% of Google's revenue ($162.45 billion).

Ad Exchanges

Google Ad Exchange (AdX).

Auctions

Auction: A mechanism for buying or selling goods or services by means of eliciting bids from interested parties.

Classic example: Auction of a painting, or art in general.

Most prominent example nowadays: Ad auctions

Selling advertising space (ad impressions) on online market places (ad exchanges).

In 2022, this accounted for 58% of Google's revenue ($162.45 billion).

Ad Exchanges

doubleclick by Google

Facebook Exchange

RIGHTMEDIA from Yahoo!

appnexus

OpenX

Microsoft Exchange

Google

Google Ad Exchange (AdX).

Actually, virtually all of these Ad exchanges use the first-price auction!

How should you bid in the FPA?

"I should bid lower than the amount I am willing to spend, to win the item on sale for a smaller price."

"I shouldn't bid too low though, because that increases the chances of not winning the item at all."

"How low should I bid?"

We will now attempt to answer this question.

Bayesian Games

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players,

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players,

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i,

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players,

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i,

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i,

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players,

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i,

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i,

$p^i : \Theta \rightarrow [0,1]$ is a *joint prior* over types (from the perspective of i), and

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players,

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i,

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i,

$p^i : \Theta \rightarrow [0,1]$ is a *joint prior* over types (from the perspective of i), and

$u = (u_1, \ldots, u_n)$ is the utility function for player i, with $u : A \times \Theta \rightarrow \mathbb{R}_{\geq 0}$.

Bayesian Games

Bayesian Games

A strategy s_i for player i is a function $s_i : \Theta_i \to A_i$ which prescribes an action for every type.

Bayesian Games

A strategy s_i for player i is a function $s_i : \Theta_i \rightarrow A_i$ which prescribes an action for every type.

We will use U_i to denote the *expected utility* of agent i when her type is θ_i, its strategy is s_i and the strategies of the other players are s_{-i}. The expectation is over the common prior p.

Bayesian Games

A strategy s_i for player i is a function $s_i : \Theta_i \rightarrow A_i$ which prescribes an action for every type.

We will use U_i to denote the *expected utility* of agent i when her type is θ_i, its strategy is s_i and the strategies of the other players are s_{-i}. The expectation is over the common prior p.

We write $U_i(s_i, s_{-i}; \theta_i)$.

By picture

θ_{i1}	θ_{i2}	θ_{i3}	θ_{i4}	θ_{i5}	θ_{i6}

By picture

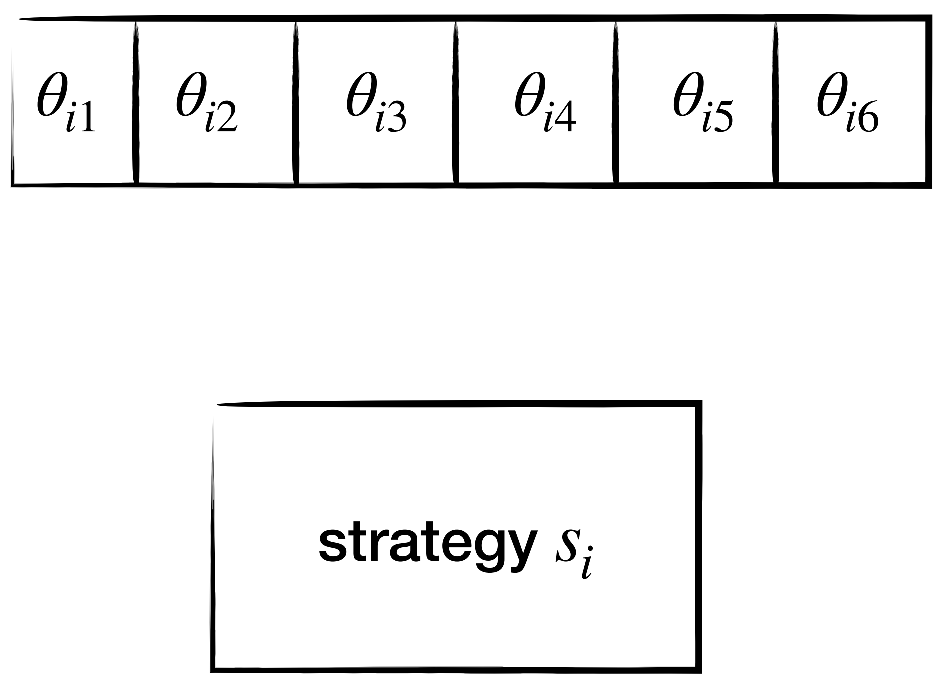

θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i

By picture

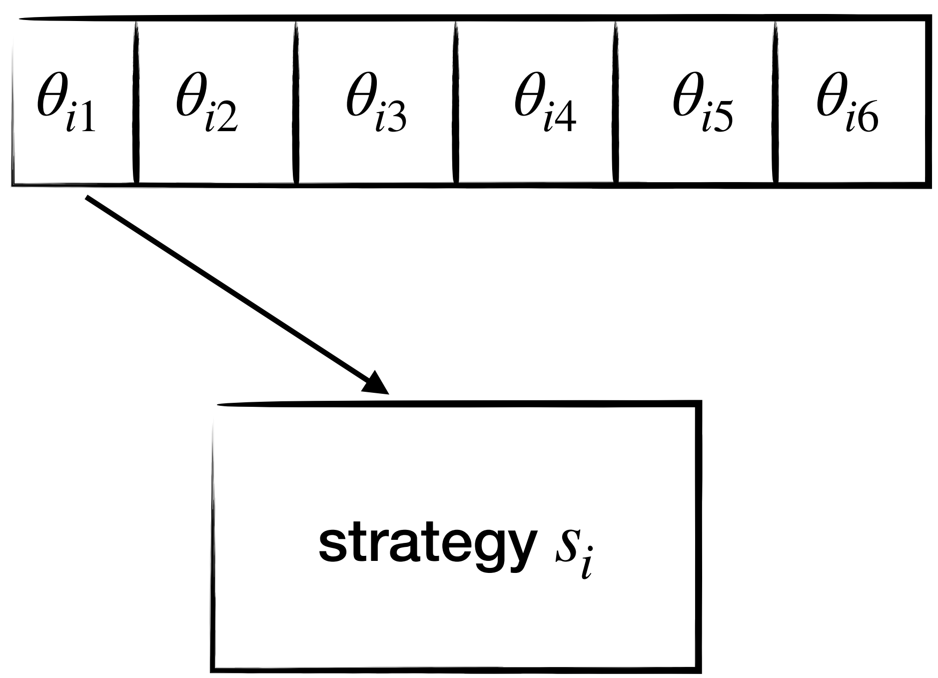

θ_{i1}
θ_{i2}
θ_{i3}
θ_{i4}
θ_{i5}
θ_{i6}
strategy s_i

By picture

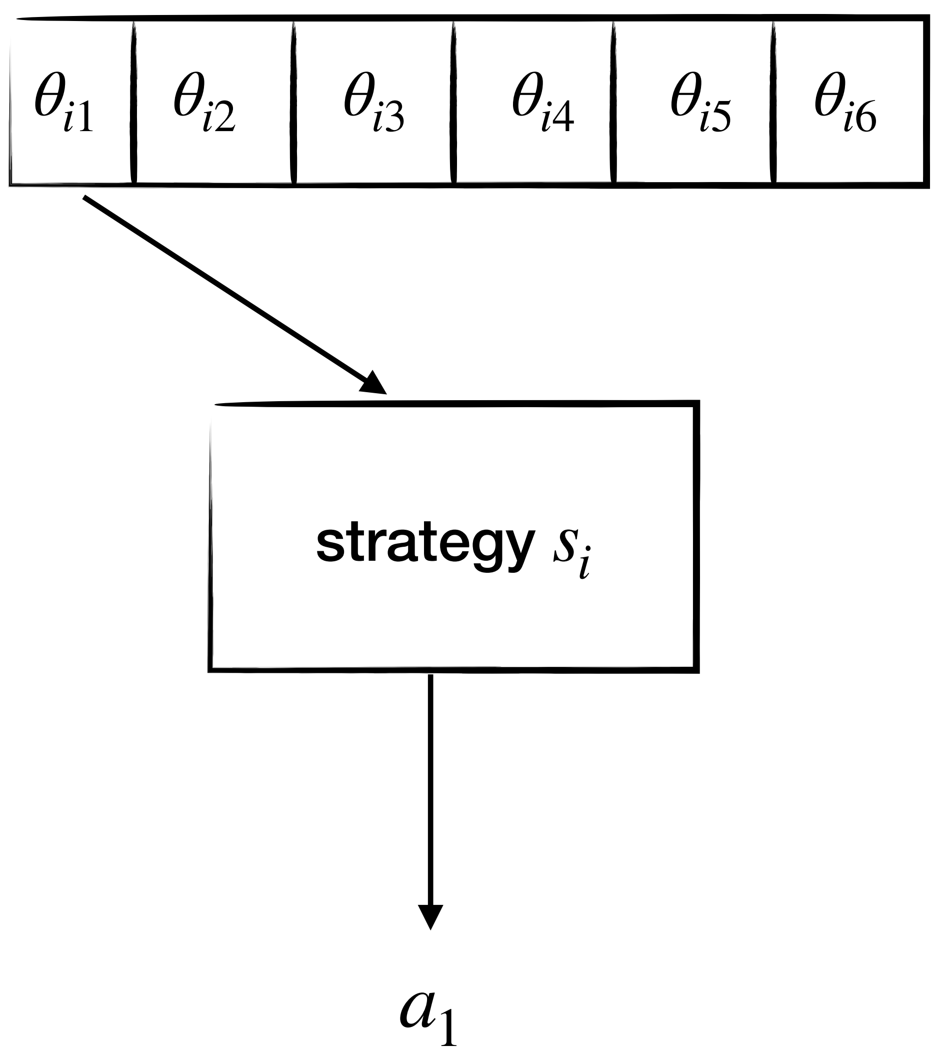

θ_{i1}
θ_{i2}
θ_{i3}
θ_{i4}
θ_{i5}
θ_{i6}
strategy s_i
a_1

By picture

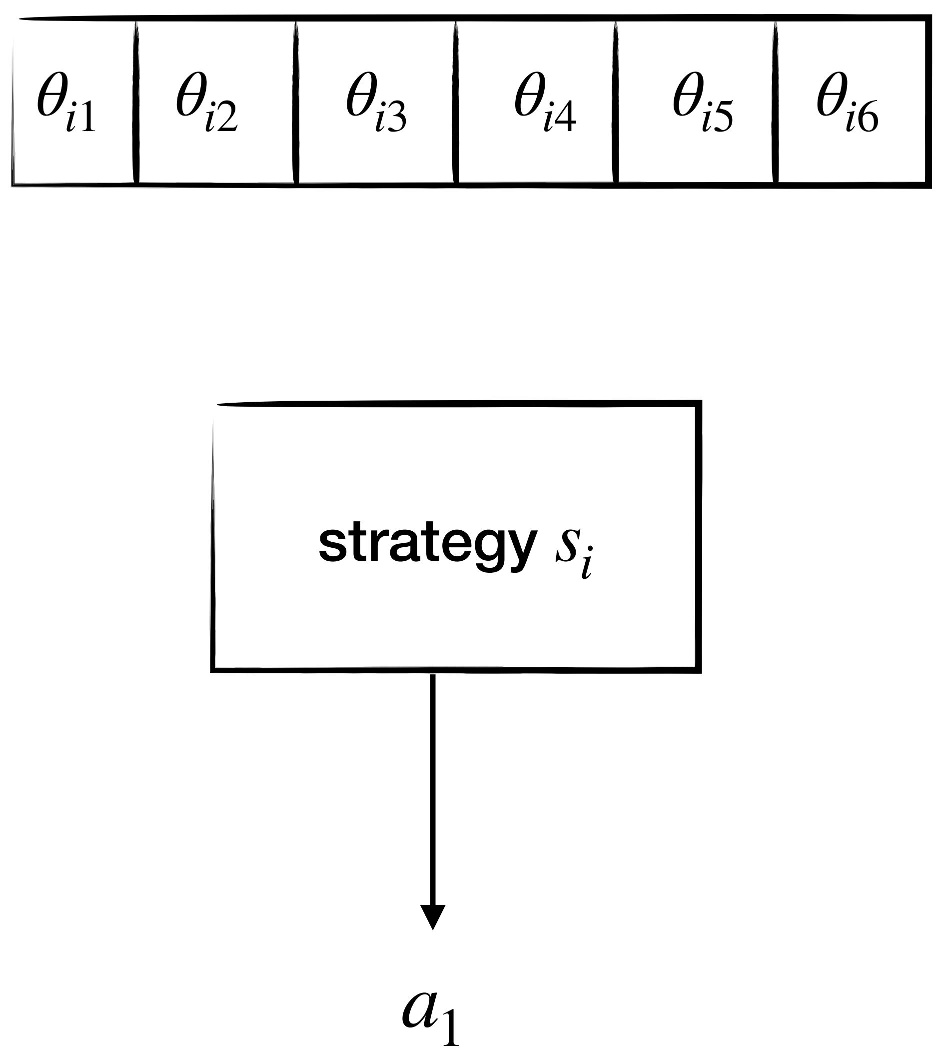
θ_{i1}
θ_{i2}
θ_{i3}
θ_{i4}
θ_{i5}
θ_{i6}
strategy s_i
a_1

By picture

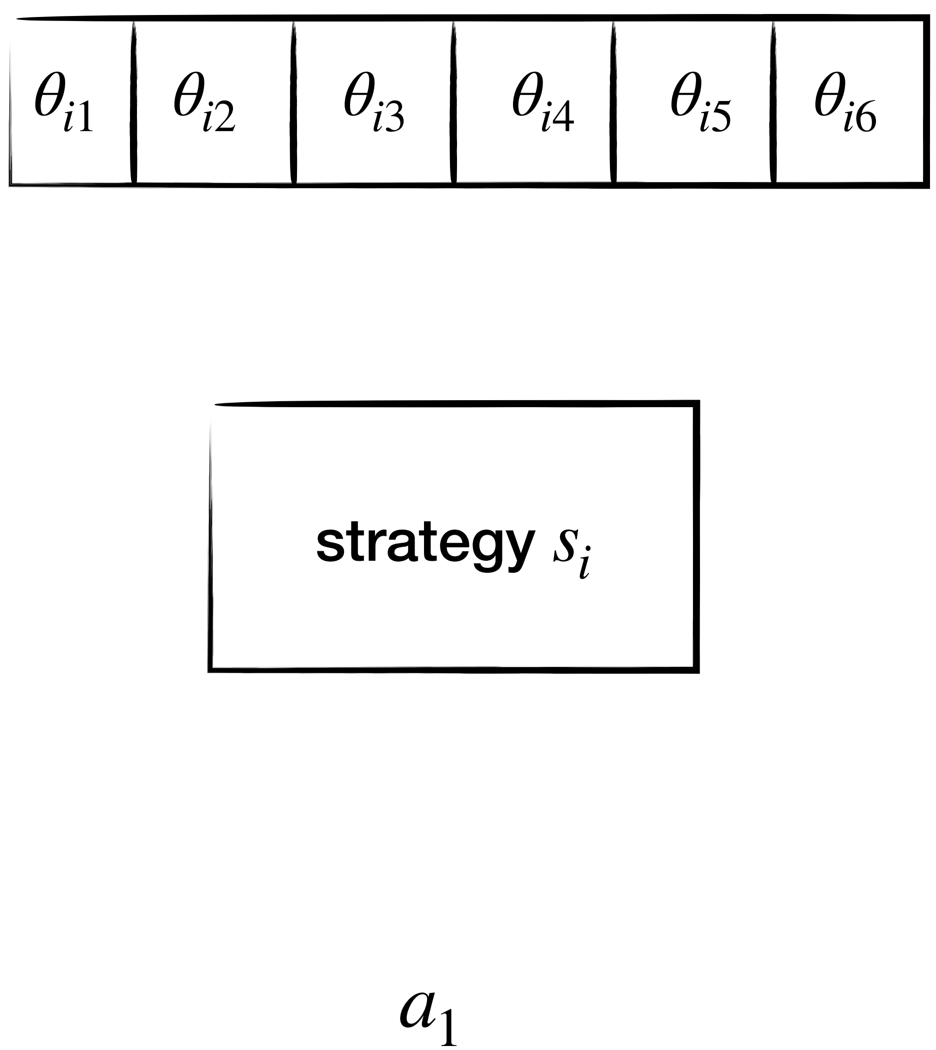
θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i
a_1

By picture

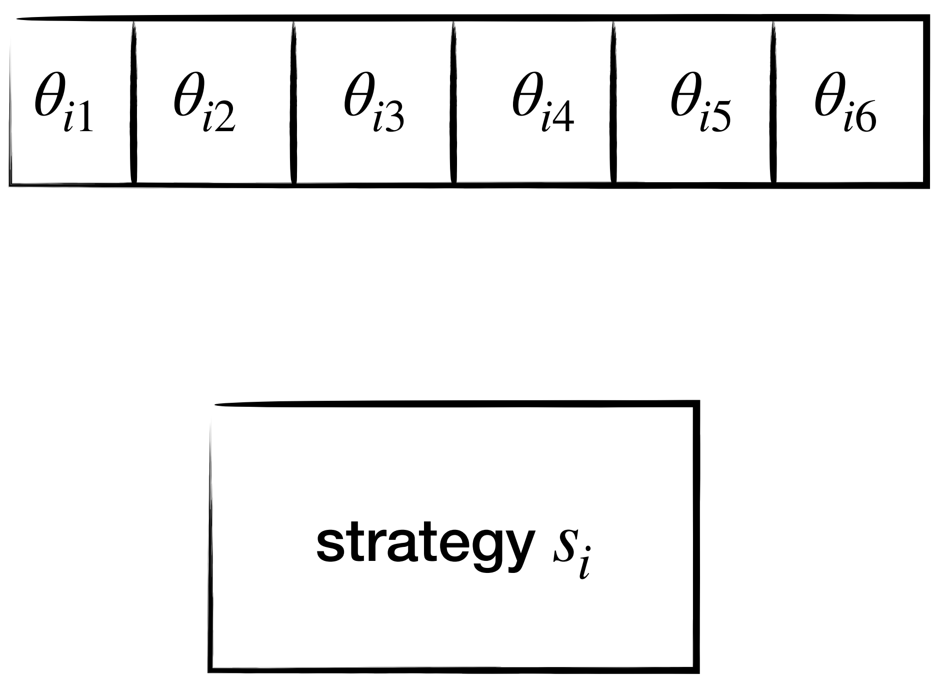
θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i

By picture

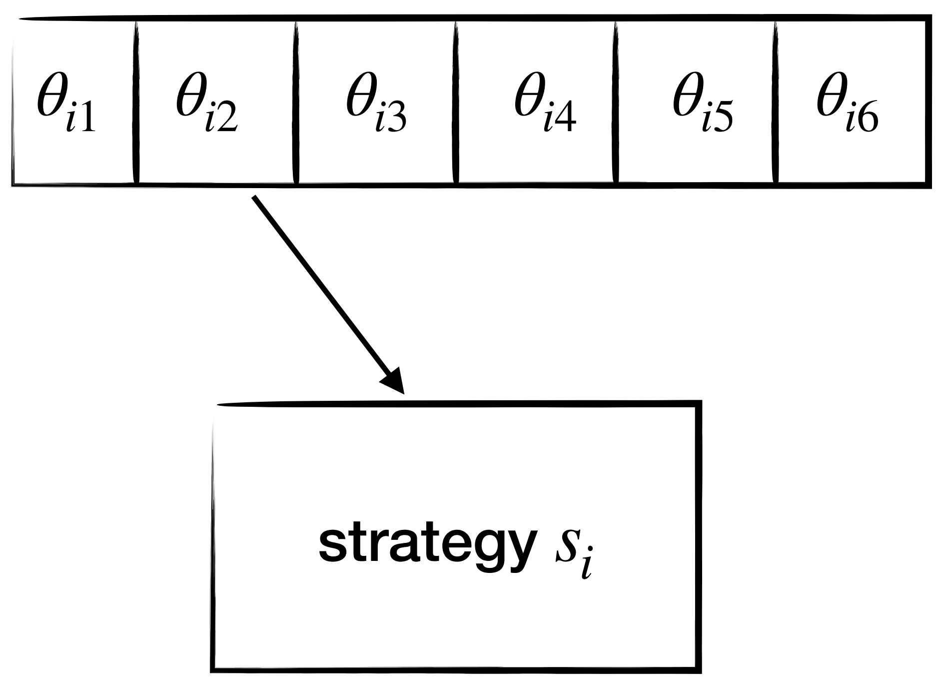

θ_{i1}
θ_{i2}
θ_{i3}
θ_{i4}
θ_{i5}
θ_{i6}
strategy s_i

By picture

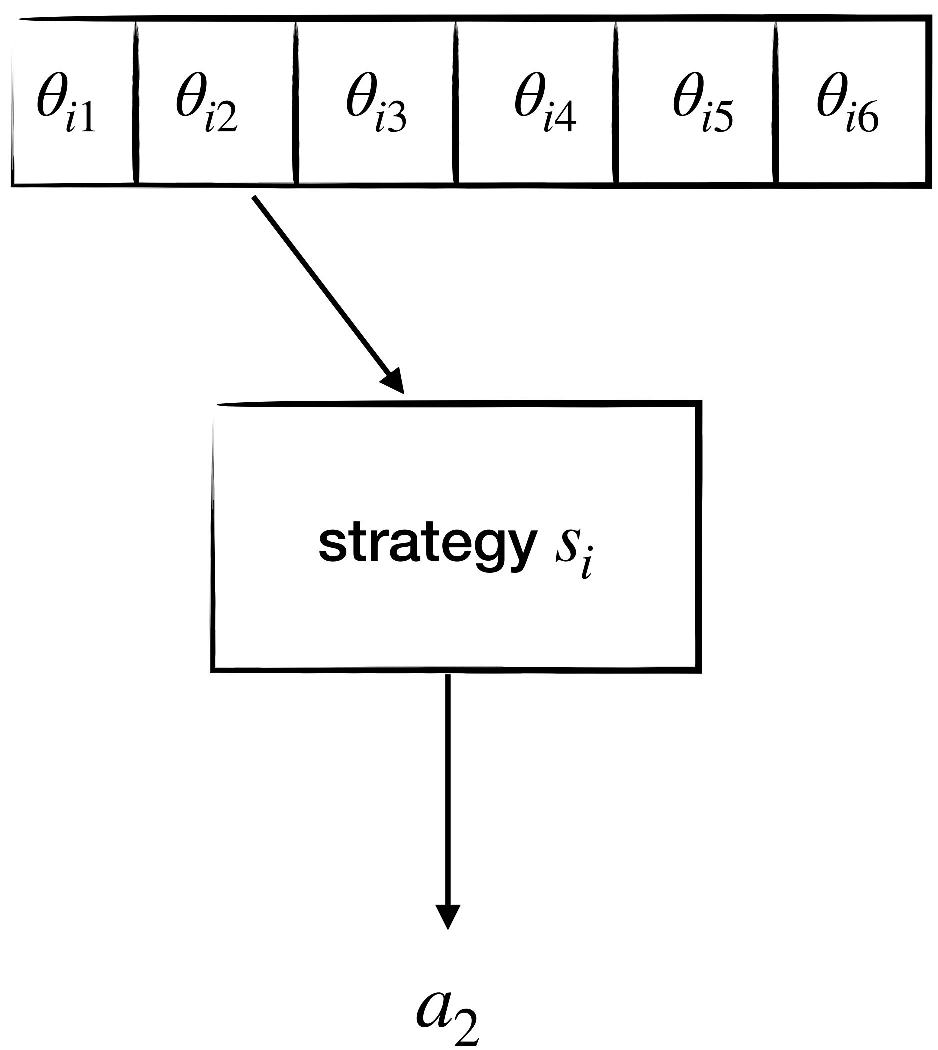

θ_{i1}
θ_{i2}
θ_{i3}
θ_{i4}
θ_{i5}
θ_{i6}
strategy s_i
a_2

By picture

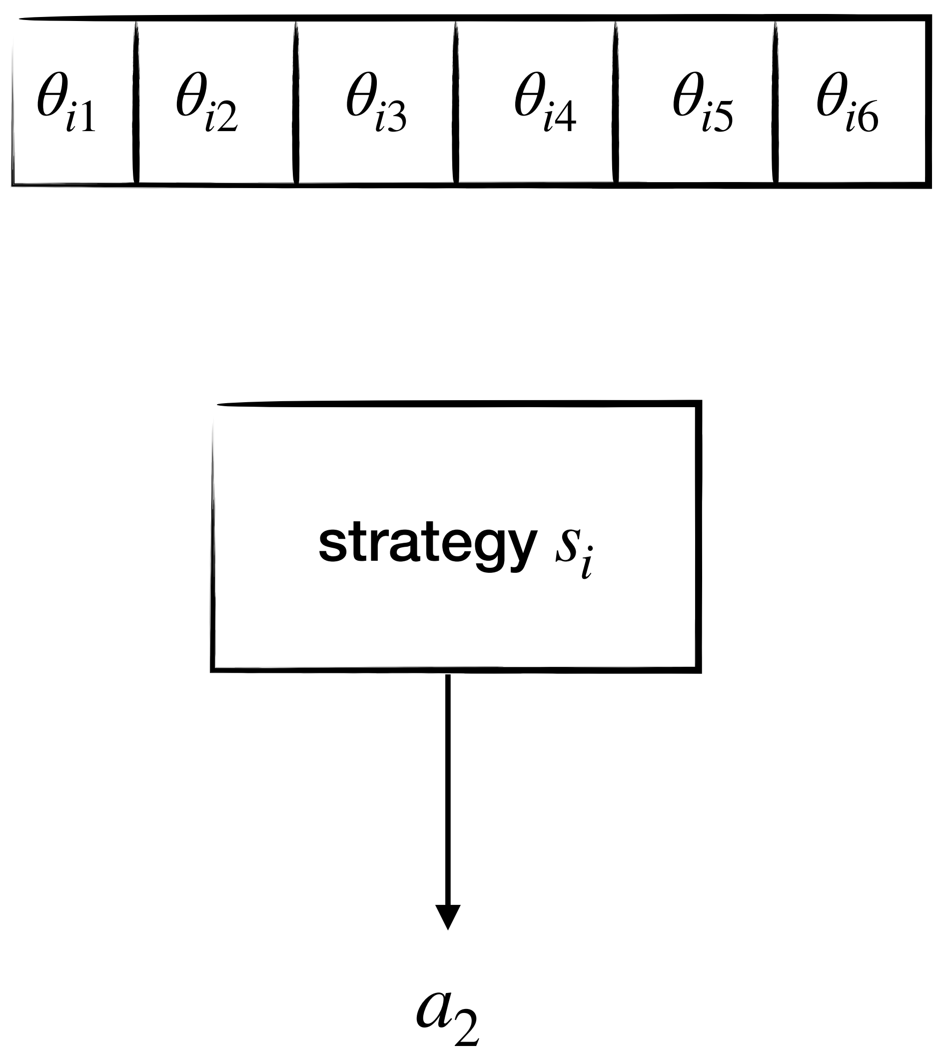

θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i
a_2

By picture

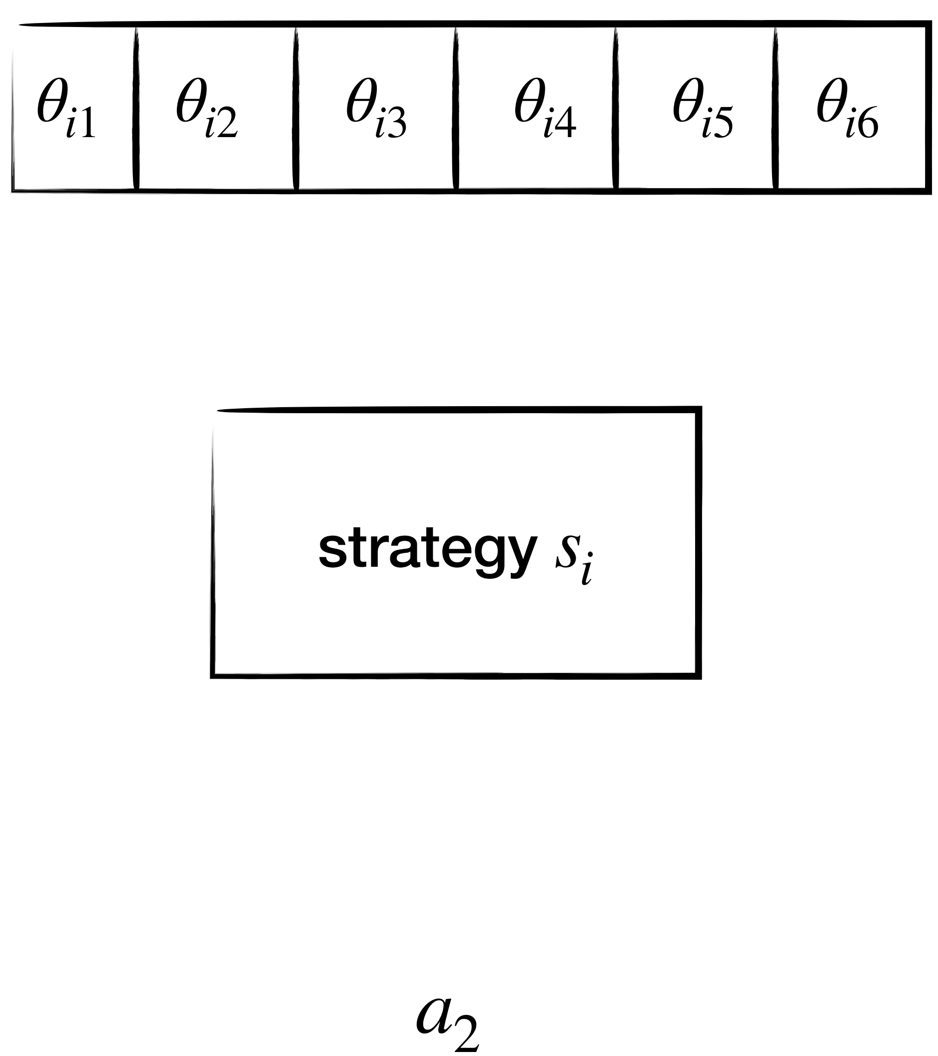

θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i
a_2

By picture

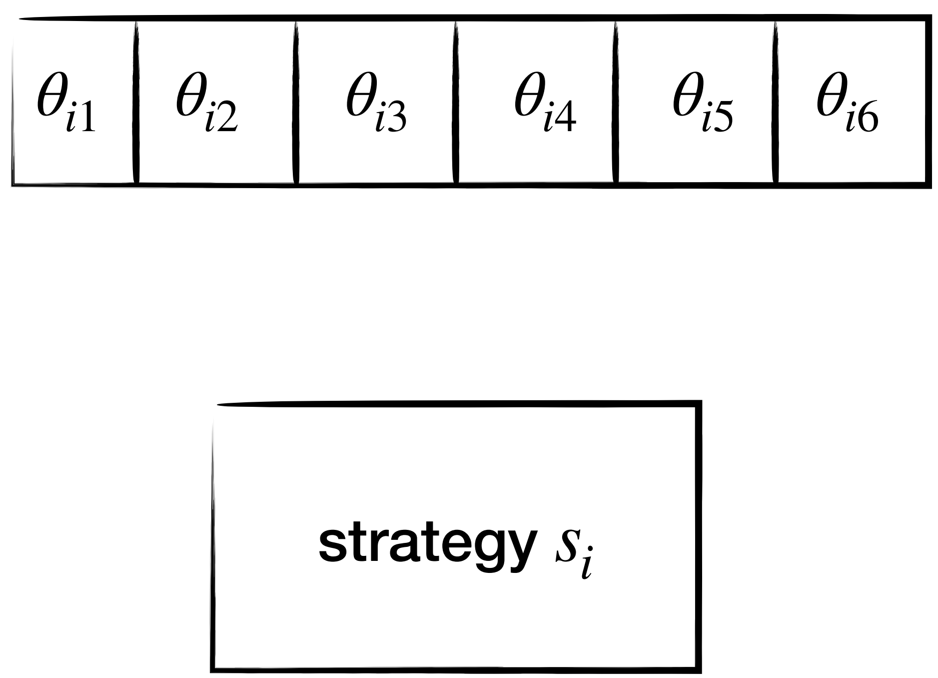
θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i

By picture

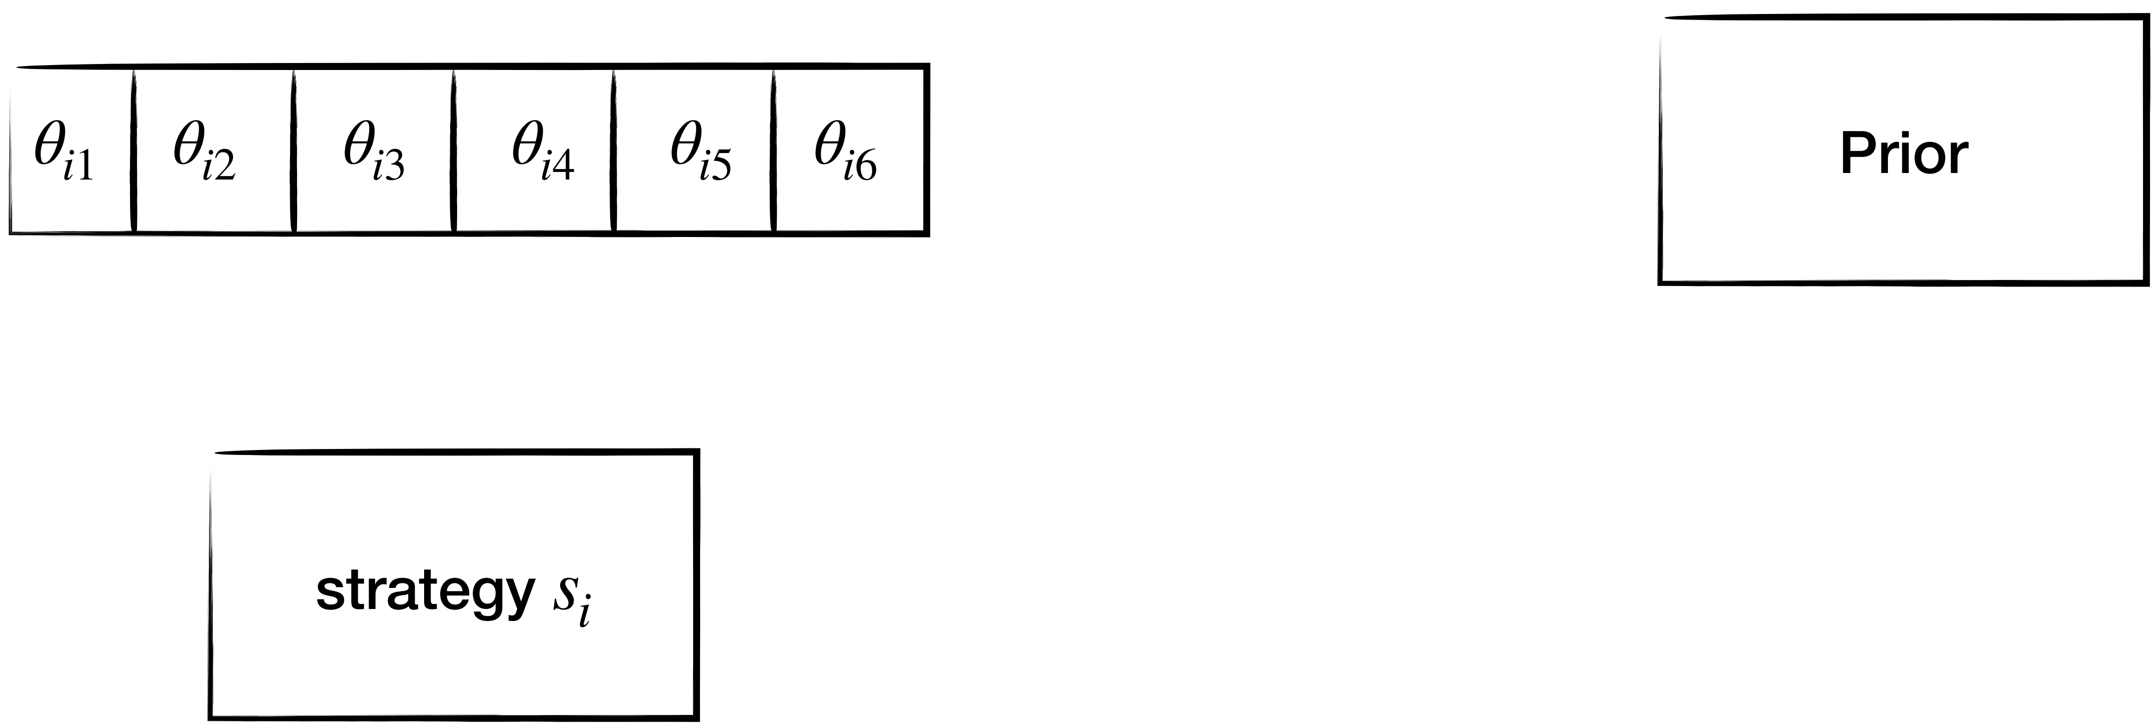

θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
Prior
strategy s_i

By picture

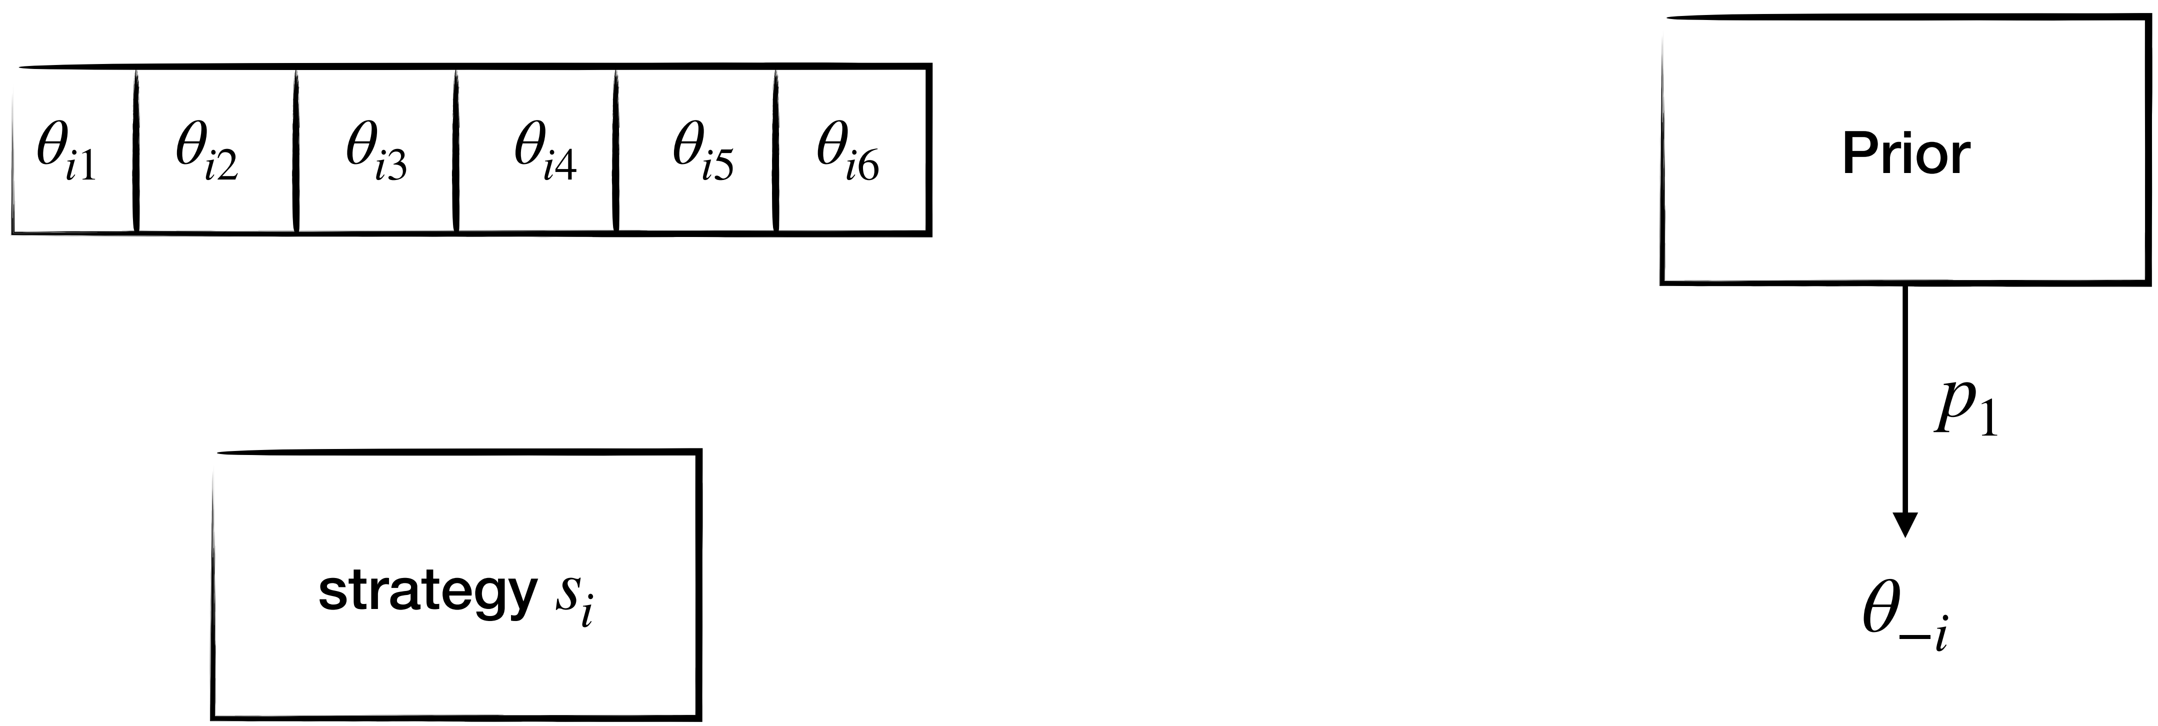

θ_{i1} θ_{i2} θ_{i3} θ_{i4} θ_{i5} θ_{i6}
strategy s_i
Prior
p_1
θ_{-i}

By picture

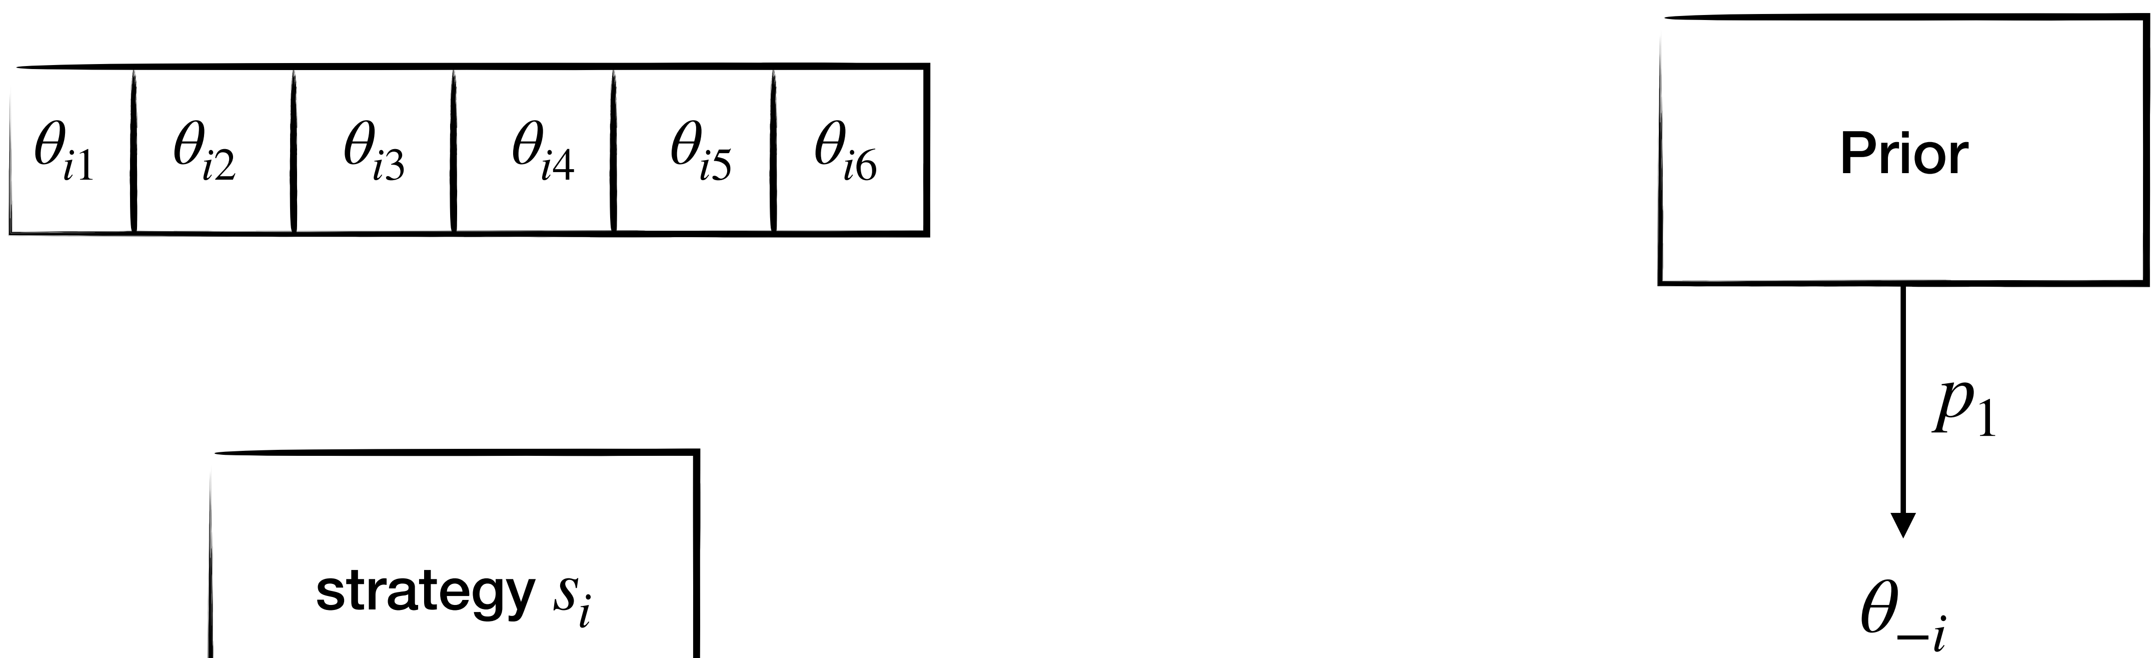

$$p_1 \cdot u_i \left(s_i(\theta_{i1}), s_{-i}(\theta_{-i}) \right)$$

By picture

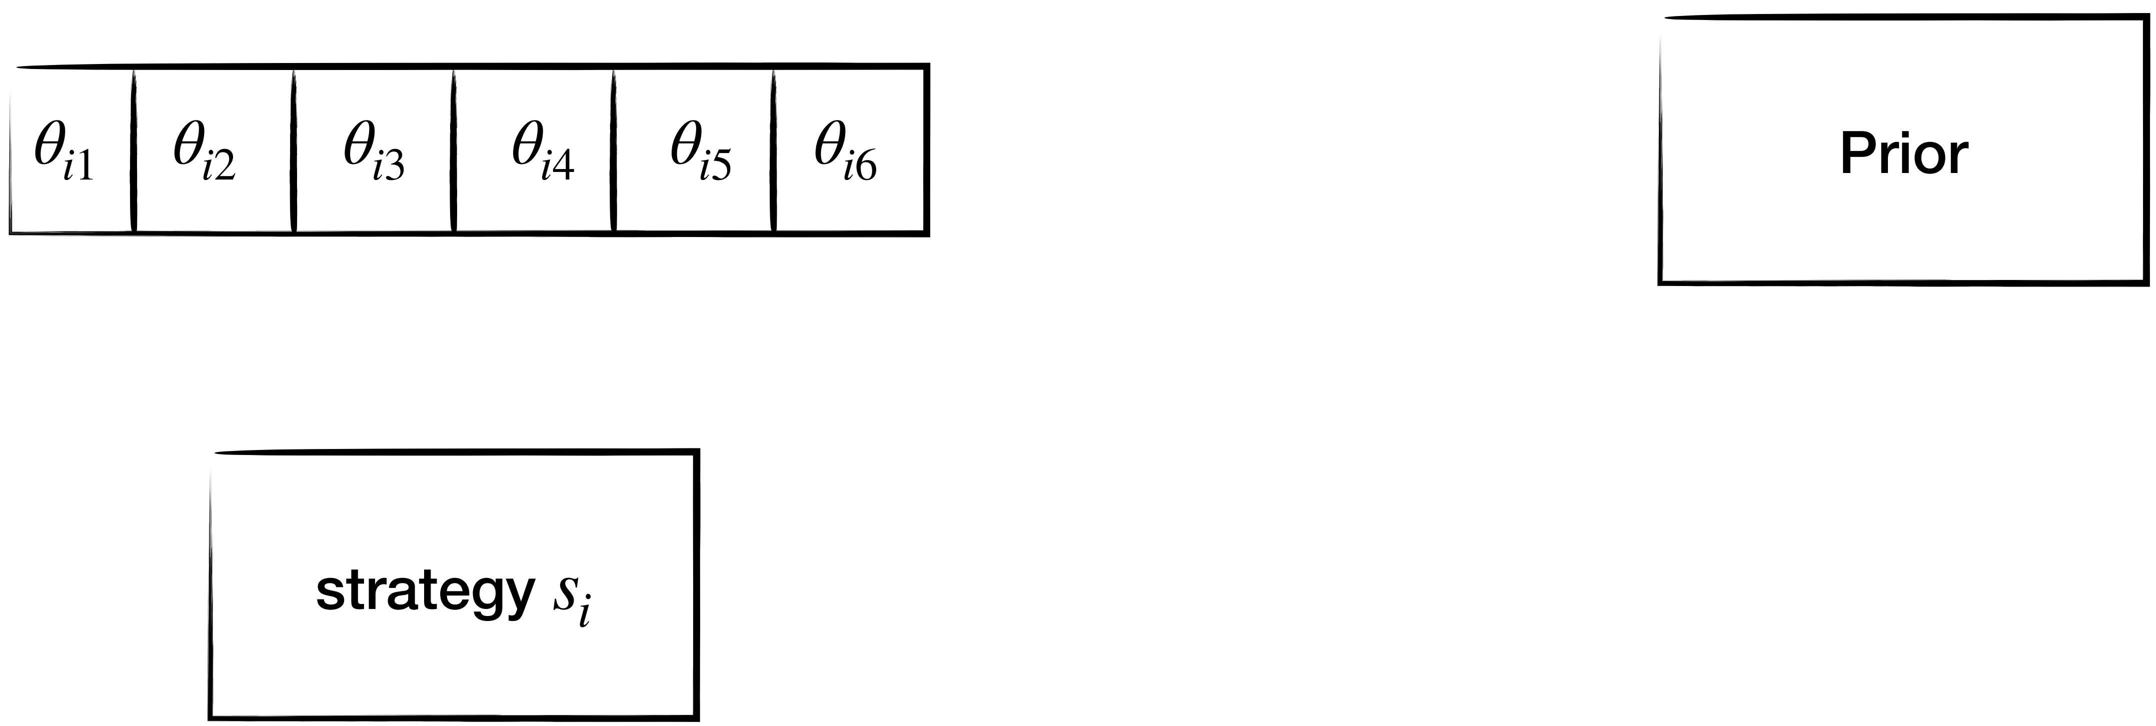

$$p_1 \cdot u_i \left(s_i(\theta_{i1}), s_{-i}(\theta_{-i}) \right)$$

By picture

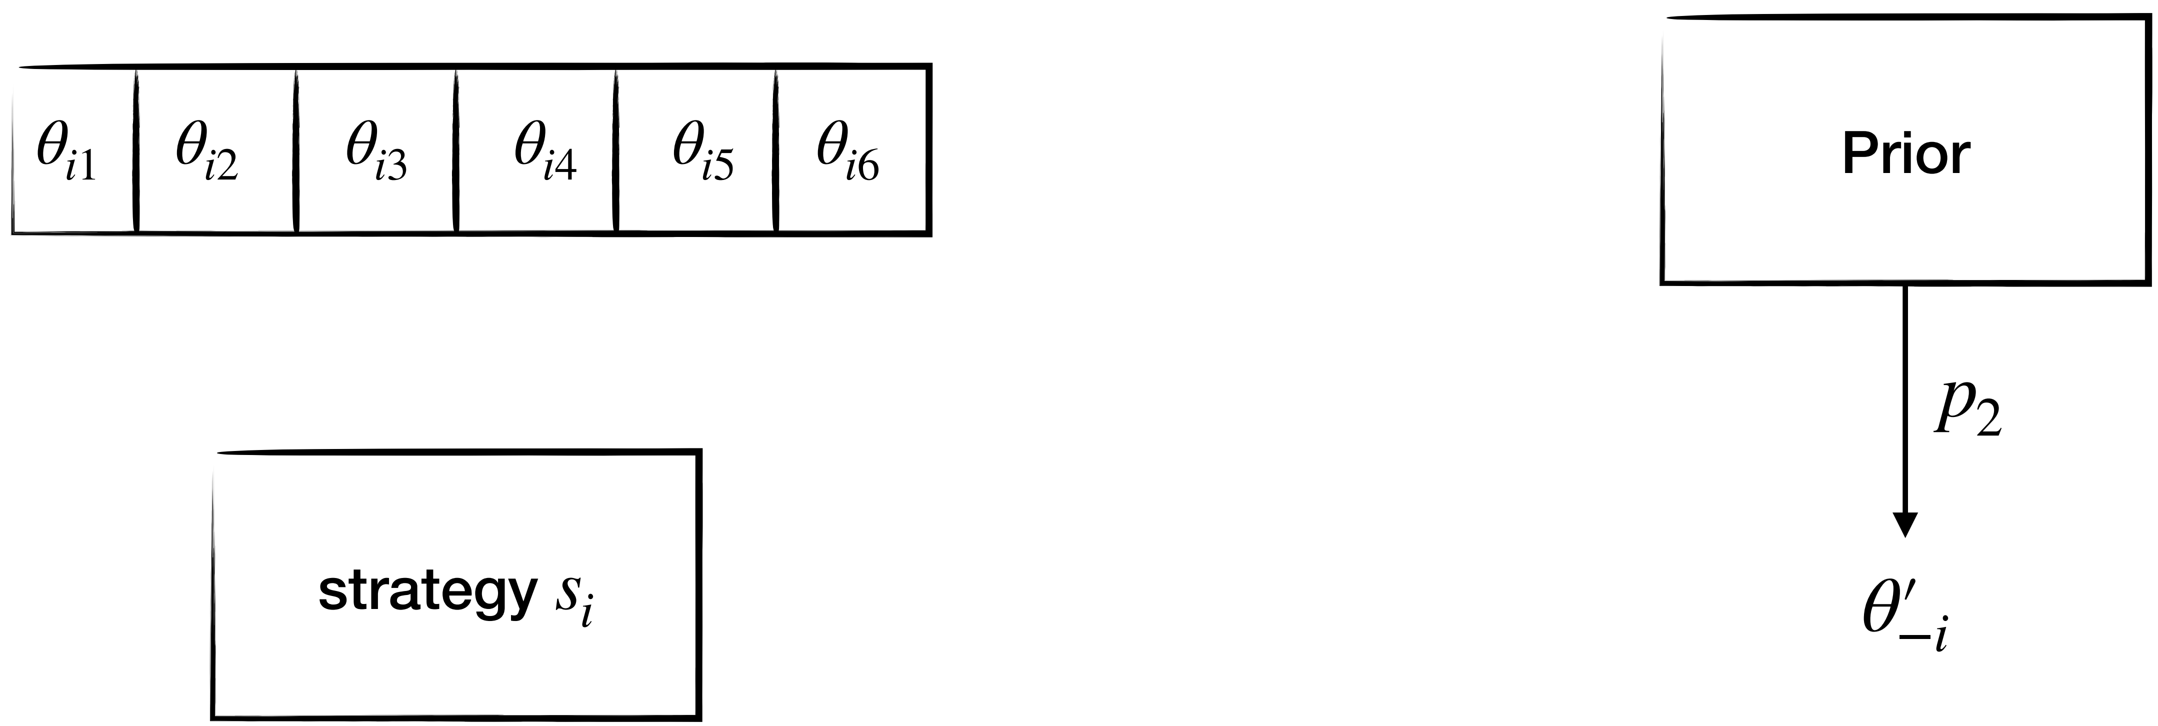

$$p_1 \cdot u_i \left(s_i(\theta_{i1}), s_{-i}(\theta_{-i}) \right)$$

By picture

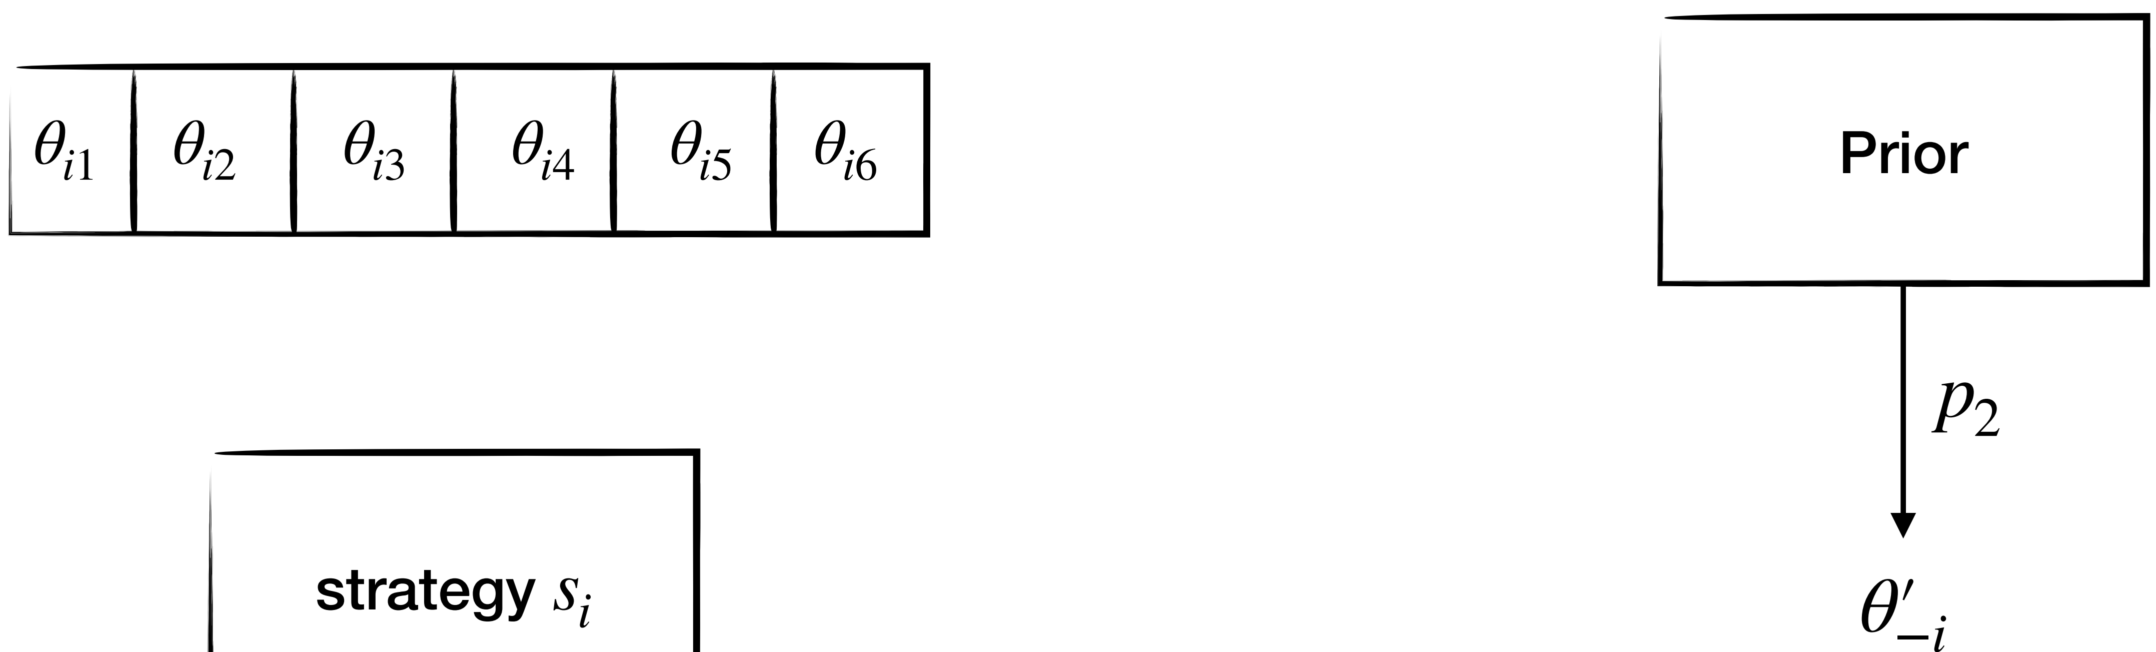

$$p_1 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta_{-i})\right) \; +$$
$$p_2 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta'_{-i})\right)$$

By picture

θ_{i1}	θ_{i2}	θ_{i3}	θ_{i4}	θ_{i5}	θ_{i6}

Prior

strategy s_i

$$p_1 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta_{-i})\right) \; +$$

$$p_2 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta'_{-i})\right)$$

By picture

θ_{i1}	θ_{i2}	θ_{i3}	θ_{i4}	θ_{i5}	θ_{i6}

Prior

strategy s_i

$$p_1 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta_{-i})\right) \;+$$

$$p_2 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta'_{-i})\right) \;+$$

By picture

θ_{i1}	θ_{i2}	θ_{i3}	θ_{i4}	θ_{i5}	θ_{i6}

Prior

strategy s_i

$$p_1 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta_{-i})\right) \;+$$

$$p_2 \cdot u_i\left(s_i(\theta_{i1}), s_{-i}(\theta'_{-i})\right) \;+\; \cdots$$

Solution Concept:
Bayes-Nash Equilibrium

Solution Concept:
Bayes-Nash Equilibrium

(Pure) Bayes-Nash Equilibrium: A (pure) strategy profile $s = (s_1, \ldots, s_n)$ such that every player maximises its expected utility.

Solution Concept:
Bayes-Nash Equilibrium

(Pure) Bayes-Nash Equilibrium: A (pure) strategy profile $s = (s_1, \ldots, s_n)$ such that every player maximises its expected utility.

where the expectation is taken over the uncertainty about the types of the others.

Solution Concept: Bayes-Nash Equilibrium

(Pure) Bayes-Nash Equilibrium: A (pure) strategy profile $s = (s_1, \ldots, s_n)$ such that every player maximises its expected utility.

where the expectation is taken over the uncertainty about the types of the others.

One can also similarly define mixed Bayes-Nash equilibria.

Auctions as Bayesian Games

Auctions as Bayesian Games

We model the beliefs of bidders for other bidders via *probability distributions*.

Auctions as Bayesian Games

We model the beliefs of bidders for other bidders via *probability distributions*.

In particular, for each pair of bidders i and j, there is a probability distribution F_{ij} which captures the beliefs of bidder i for the values of bidder j.

Auctions as Bayesian Games

We model the beliefs of bidders for other bidders via *probability distributions*.

In particular, for each pair of bidders i and j, there is a probability distribution F_{ij} which captures the beliefs of bidder i for the values of bidder j.

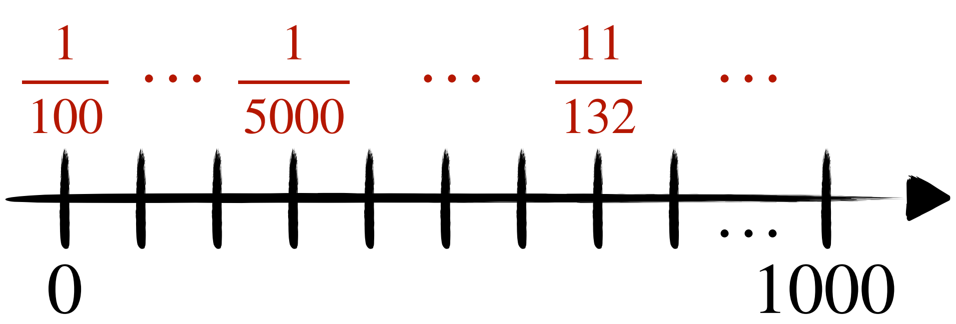

A discrete distribution

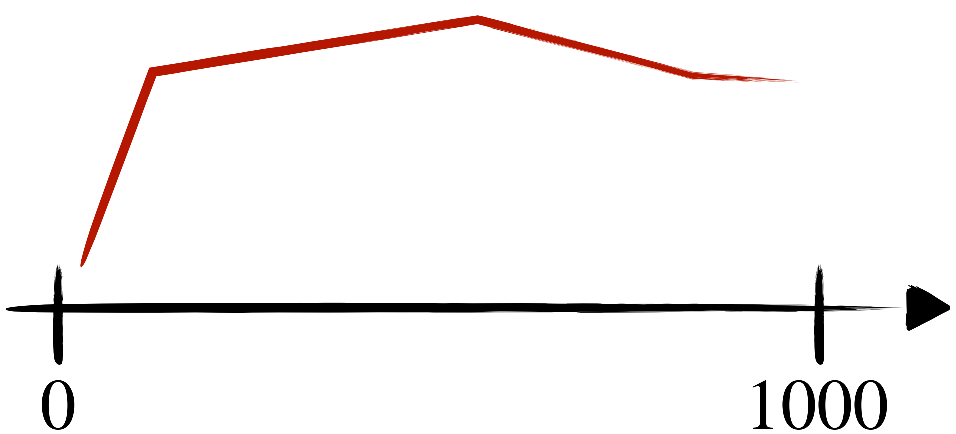

A continuous distribution

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players,

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i,

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i,

$p^i : \Theta \to [0,1]$ is a *joint prior* over types (from the perspective of i), and
$u = (u_1, \ldots, u_n)$ is the utility function for player i, with $u : A \times \Theta \to \mathbb{R}$.

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players, **the bidders**

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i,

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i,

$p^i : \Theta \rightarrow [0,1]$ is a *joint prior* over types (from the perspective of i), and
$u = (u_1, \ldots, u_n)$ is the utility function for player i, with $u : A \times \Theta \rightarrow \mathbb{R}$.

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players, **the bidders**

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i, **the bids** $b_1, b_2, \ldots$

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i,

$p^i : \Theta \rightarrow [0,1]$ is a *joint prior* over types (from the perspective of i), and $u = (u_1, \ldots, u_n)$ is the utility function for player i, with $u : A \times \Theta \rightarrow \mathbb{R}$.

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players, **the bidders**

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i, **the bids** $b_1, b_2, \ldots$

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i, **the values** $v_1, v_2, \ldots, v_n$

$p^i : \Theta \to [0,1]$ is a *joint prior* over types (from the perspective of i), and $u = (u_1, \ldots, u_n)$ is the utility function for player i, with $u : A \times \Theta \to \mathbb{R}$.

Bayesian Games

Bayesian games of *incomplete information* (Harsanyi 1967) model uncertainty about the values of the other players.

A Bayesian game is a tuple (N, A, Θ, p, u) where

N is the set of players, **the bidders**

$A = A_1 \times \ldots \times A_n$ is the set of action profiles, where A_i is the set of available actions to player i, **the bids** $b_1, b_2, \ldots$

$\Theta = \Theta_1 \times \ldots \times \Theta_n$ is a set of *type profiles*, where Θ_i is the set of *types* of player i, **the values** $v_1, v_2, \ldots, v_n$

the distributions F_{ij}

$p^i : \Theta \to [0,1]$ is a *joint prior* over types (from the perspective of i), and
$u = (u_1, \ldots, u_n)$ is the utility function for player i, with $u : A \times \Theta \to \mathbb{R}$.

Bayesian Games

A strategy s_i for player i is a function $s_i : \Theta_i \rightarrow A_i$ which prescribes an action for every type.

We will use U_i to denote the *expected utility* of agent i when her type is θ_i, its strategy is s_i and the strategies of the other players are s_{-i}. The expectation is over the common prior p.

We write $U_i(s_i, s_{-i}; \theta_i)$.

Bayesian Games

A strategy s_i for player i is a function $s_i : \Theta_i \rightarrow A_i$ which prescribes an action for every type.

We will use U_i to denote the *expected utility* of agent i when her type is θ_i, its strategy is s_i and the strategies of the other players are s_{-i}. The expectation is over the common prior p.

We write $U_i(s_i, s_{-i}; \theta_i)$.

A function β mapping values to bids.

Auctions as Bayesian Games

We model the beliefs of bidders for other bidders via *probability distributions*.

In particular, for each pair of bidders i and j, there is a probability distribution F_{ij} which captures the beliefs of bidder i for the values of bidder j.

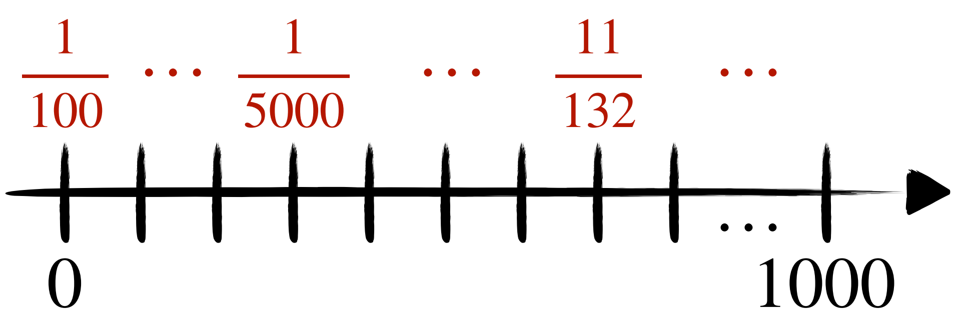

A discrete distribution

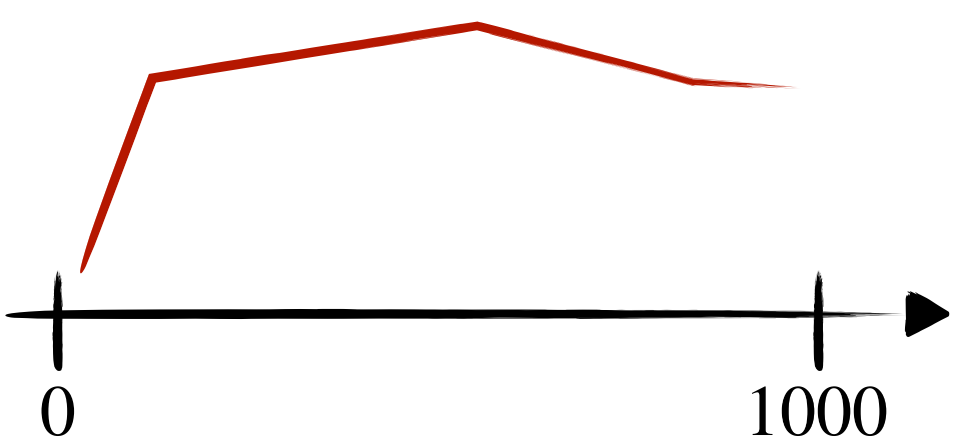

A continuous distribution

Some simpler interesting cases

We model the beliefs of bidders for other bidders via *probability distributions*.

In particular, for each pair of bidders i and j, there is a probability distribution F_{ij} which captures the beliefs of bidder i for the values of bidder j.

Some simpler interesting cases

We model the beliefs of bidders for other bidders via *probability distributions*.

In particular, for each pair of bidders i and j, there is a probability distribution F_{ij} which captures the beliefs of bidder i for the values of bidder j.

Given a bidder j, every bidder i has the same beliefs about j, i.e.,

$F_{ij} = F_j \quad \forall i$ (objective beliefs)

Some simpler interesting cases

We model the beliefs of bidders for other bidders via *probability distributions*.

In particular, for each pair of bidders i and j, there is a probability distribution F_{ij} which captures the beliefs of bidder i for the values of bidder j.

Given a bidder j, every bidder i has the same beliefs about j, i.e.,

$F_{ij} = F_j \quad \forall i$ (objective beliefs)

The values of all bidders come from the same distribution, i.e.,

$F_i = F_j \quad \forall i, j$ (symmetric beliefs)

The bidder’s optimisation problem

The bidder's optimisation problem

A bidder i would like to come up with a function $\beta_i : V \to B$ (a function which maps values to bids), which

The bidder's optimisation problem

A bidder i would like to come up with a function $\beta_i : V \to B$ (a function which maps values to bids), which

given the bidding functions $\beta_1, \beta_2, \ldots, \beta_{i-1}, \beta_{i+1}, \ldots, \beta_n$ of the other bidders,

The bidder's optimisation problem

A bidder i would like to come up with a function $\beta_i : V \rightarrow B$ (a function which maps values to bids), which

given the bidding functions $\beta_1, \beta_2, \ldots, \beta_{i-1}, \beta_{i+1}, \ldots, \beta_n$ of the other bidders,

and given the beliefs F_{ij} of the bidder for the values of any other bidder j,

The bidder's optimisation problem

A bidder i would like to come up with a function $\beta_i : V \to B$ (a function which maps values to bids), which

given the bidding functions $\beta_1, \beta_2, \ldots, \beta_{i-1}, \beta_{i+1}, \ldots, \beta_n$ of the other bidders,

and given the beliefs F_{ij} of the bidder for the values of any other bidder j,

maximises the expected utility of the bidder

$$\mathbb{E}_{v_j \sim F_{ij},\ \forall j \neq i} \left[(v_i - \beta(v_i)) \cdot \frac{1}{W(\beta_1(v_1), \ldots, \beta_n(v_n))} \right]$$

Nash Equilibrium Existence

Does a mixed Nash equilibrium always exist?

Solution Concept:
Mixed Nash Equilibrium

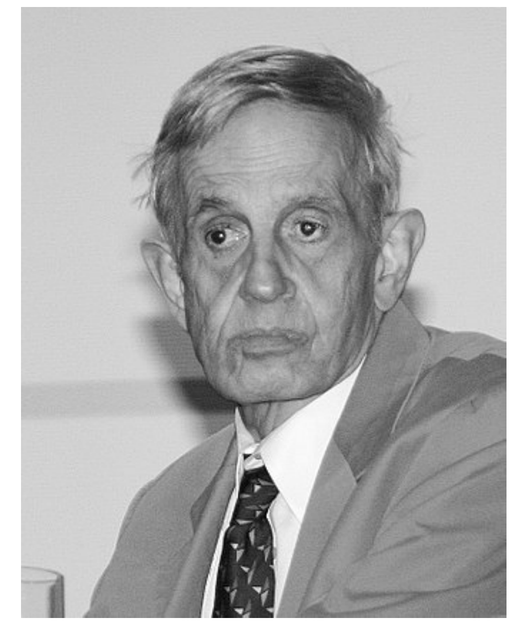

Introduced by Nash in 1951 (in his PhD dissertation).

MNE: "I won't change unless the others change", hence a *stable* outcome.

It is universal!

Theorem (Nash 1951): Every (finite normal-form) game has at least one mixed Nash equilibrium.

Nash Equilibrium Existence

Does a mixed Bayes-Nash equilibrium always exist?

Nash Equilibrium Existence

Does a mixed Bayes-Nash equilibrium always exist?

If the bids are infinite and/or the distributions are continuous, then Nash's theorem does not apply.

Nash Equilibrium Existence

Does a mixed Bayes-Nash equilibrium always exist?

If the bids are infinite and/or the distributions are continuous, then Nash's theorem does not apply.

When the bids are finite and distributions are discrete, then we have a finite game. Nash's theorem should apply…

Nash Equilibrium Existence

Does a mixed Bayes-Nash equilibrium always exist?

If the bids are infinite and/or the distributions are continuous, then Nash's theorem does not apply.

When the bids are finite and distributions are discrete, then we have a finite game. Nash's theorem should apply…

… except this game is Bayesian, not a normal form game.

Nash Equilibrium Existence

Does a mixed Bayes-Nash equilibrium always exist?

If the bids are infinite and/or the distributions are continuous, then Nash's theorem does not apply.

When the bids are finite and distributions are discrete, then we have a finite game. Nash's theorem should apply…

… except this game is Bayesian, not a normal form game.

Still, it can be proven that for a Bayesian game with finite type space and finite action space, a Bayes-Nash equilibrium always exists.

Nash Equilibrium Existence

Does a mixed Bayes-Nash equilibrium always exist?

If the bids are infinite and/or the distributions are continuous, then Nash's theorem does not apply.

When the bids are finite and distributions are discrete, then we have a finite game. Nash's theorem should apply…

… except this game is Bayesian, not a normal form game.

Still, it can be proven that for a Bayesian game with finite type space and finite action space, a Bayes-Nash equilibrium always exists.

> Idea: Transform the Bayesian game into a (standard) extensive form or normal form game.

Nash Equilibrium Existence

Does a pure Nash equilibrium always exist?

Nash Equilibrium Existence

Does a pure Nash equilibrium always exist?

Not necessarily, even when we have finite type and action spaces.

Nash Equilibrium Existence

Does a PNE exist for the first-price auction game?

Nash Equilibrium Existence

Does a PNE exist for the first-price auction game?

- If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

Nash Equilibrium Existence

Does a PNE exist for the first-price auction game?

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

If the distributions are continuous, then a Nash equilibrium always exists (Vickrey 1961, Athey 2001).

Nash Equilibrium Existence

Does a PNE exist for the first-price auction game?

> If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).
>
> If the distributions are continuous, then a Nash equilibrium always exists (Vickrey 1961, Athey 2001).

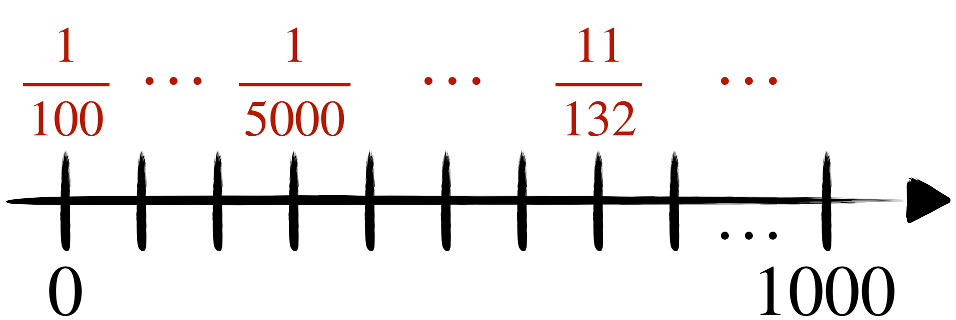

A discrete distribution

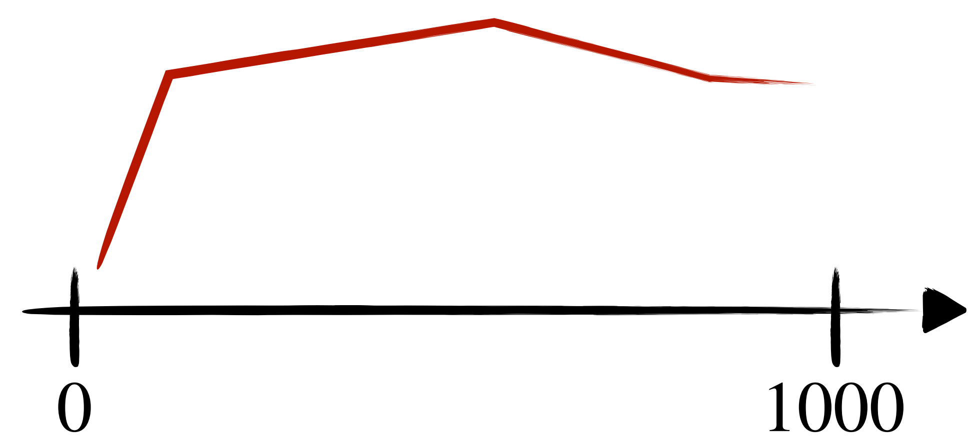

A continuous distribution

How would one prove these statements?

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

How would one prove these statements?

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

> To prove this, it suffices to construct an example of an auction (with appropriately chosen parameters - values, distributions, bids) such that an equilibrium does not exist.

How would one prove these statements?

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

> To prove this, it suffices to construct an example of an auction (with appropriately chosen parameters - values, distributions, bids) such that an equilibrium does not exist.

If the distributions are continuous, then a Nash equilibrium always exists (Vickrey 1961, Athey 2001).

How would one prove these statements?

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

> To prove this, it suffices to construct an example of an auction (with appropriately chosen parameters - values, distributions, bids) such that an equilibrium does not exist.

If the distributions are continuous, then a Nash equilibrium always exists (Vickrey 1961, Athey 2001).

> This is more intricate. The known proofs go via fixed-point theorems and/or topological lemmas.

How would one prove these statements?

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

> To prove this, it suffices to construct an example of an auction (with appropriately chosen parameters - values, distributions, bids) such that an equilibrium does not exist.

If the distributions are continuous, then a Nash equilibrium always exists (Vickrey 1961, Athey 2001).

> This is more intricate. The known proofs go via fixed-point theorems and/or topological lemmas.
>
> Athey's proof using Kakutani, [F., Giannakopoulos, Hollender, Lazos, and Poças 2023] provide a proof that used Brouwer's fixed point theorem.

A caveat of these existence proofs

A caveat of these existence proofs

This existence proofs ensure that an equilibrium exists.

A caveat of these existence proofs

This existence proofs ensure that an equilibrium exists.

So we can hope that the bidders are going to “find it” by iteratively adjusting their bids while maximising their utilities against the bids of the others.

A caveat of these existence proofs

This existence proofs ensure that an equilibrium exists.

So we can hope that the bidders are going to “find it” by iteratively adjusting their bids while maximising their utilities against the bids of the others.

But we would like to know more.

A caveat of these existence proofs

This existence proofs ensure that an equilibrium exists.

So we can hope that the bidders are going to “find it” by iteratively adjusting their bids while maximising their utilities against the bids of the others.

But we would like to know more.

How do these equilibria look like? Can we describe them?

Back to the beginning

Back to the beginning

Vickrey in 1961 did not only define the SPA…

Back to the beginning

Vickrey in 1961 did not only define the SPA…

… he also described the equilibria of the FPA …

Back to the beginning

Vickrey in 1961 did not only define the SPA…

… he also described the equilibria of the FPA …

… when all of the bidders values come from the uniform distribution.

Back to the beginning

Vickrey in 1961 did not only define the SPA…

… he also described the equilibria of the FPA …

… when all of the bidders values come from the uniform distribution.

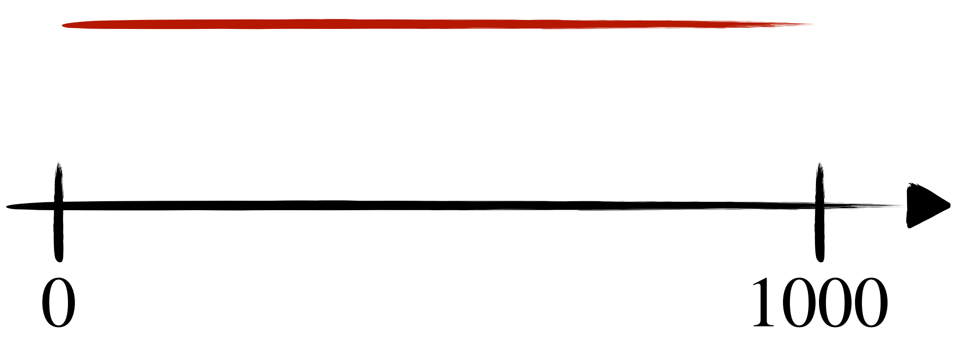

The uniform distribution

PBNE in the FPA

PBNE in the FPA

Theorem (Vickrey 1961): Consider a first-price auction with n bidders whose values are drawn independently from the uniform distribution on $[0,1]$. Then the unique symmetric equilibrium is for each bidder to bid $\frac{n-1}{n} \cdot v_i$.

Proof for $n = 2$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$\int_0^{2s_1} (v_1 - s_1)\ dv_2 + \int_{2s_1}^{1} 0\ dv_2$$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$\int_0^{2s_1} \underbrace{(v_1 - s_1)}\, dv_2 + \int_{2s_1}^{1} 0\, dv_2$$

As long as $s_1 > v_2/2 \Rightarrow v_2 < 2s_1$
bidder 1 wins and has utility $v_1 - s_1$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$\underbrace{\int_0^{2s_1} (v_1 - s_1)}\, dv_2 + \int_{2s_1}^{1} 0 \, dv_2$$

In the other case, bidder 1 loses and has utility zero.

As long as $s_1 > v_2/2 \Rightarrow v_2 < 2s_1$
bidder 1 wins and has utility $v_1 - s_1$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$U_1(s_1, s_2; v_1) = \int_0^{2s_1} (v_1 - s_1)\ dv_2$$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$U_1(s_1, s_2; v_1) = \int_0^{2s_1} (v_1 - s_1)\ dv_2 = (v_1 - s_1) \cdot (2s_1 - 0)$$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$U_1(s_1, s_2; v_1) = \int_0^{2s_1} (v_1 - s_1)\ dv_2 = (v_1 - s_1) \cdot (2s_1 - 0)$$

$$= 2v_1 \cdot s_1 - s_1^2$$

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$U_1(s_1, s_2; v_1) = \int_0^{2s_1} (v_1 - s_1)\ dv_2 = (v_1 - s_1) \cdot (2s_1 - 0)$$

$$= 2v_1 \cdot s_1 - s_1^2$$

We would like to see where this is maximised.

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$U_1(s_1, s_2; v_1) = \int_0^{2s_1} (v_1 - s_1)\ dv_2 = (v_1 - s_1) \cdot (2s_1 - 0)$$

$$= 2v_1 \cdot s_1 - s_1^2$$

We would like to see where this is maximised.

i.e., where its derivative (for s_1) is 0.

Proof for $n = 2$

Assume that bidder 2 bids $\frac{1}{2}v_2$.

The expected utility of bidder 1 can be written as:

$$U_1(s_1, s_2; v_1) = \int_0^{2s_1} (v_1 - s_1)\ dv_2 = (v_1 - s_1) \cdot (2s_1 - 0)$$

$$= 2v_1 \cdot s_1 - s_1^2$$

We would like to see where this is maximised.

i.e., where its derivative (for s_1) is 0.

Derivative: $2v_1 - s_1 = 0 \Rightarrow v_1 = s_1/2$

Revenue of the second price auction

Identical Bidders

Revenue of the second price auction

Identical Bidders

COROLLARY (OF THE BULOW-KLEMPERER THEOREM)

For n bidders with uniform iid priors, the second-price auction achieves at least a $(n-1)/n$-fraction of the optimal expected revenue (in equilibrium).

Revenue of the second price auction

Identical Bidders

COROLLARY (OF THE BULOW-KLEMPERER THEOREM)

For n bidders with uniform iid priors, the second-price auction achieves at least a $(n-1)/n$-fraction of the optimal expected revenue (in equilibrium).

THEOREM

For n bidders with uniform iid priors, the first-price auction achieves at least a $(n-1)/n$-fraction of the optimal expected revenue in equilibrium.

Revenue of the second price auction

Identical Bidders

COROLLARY (OF THE BULOW-KLEMPERER THEOREM)

For n bidders with uniform iid priors, the second-price auction achieves at least a $(n-1)/n$-fraction of the optimal expected revenue (in equilibrium).

Is this a coincidence?

THEOREM

For n bidders with uniform iid priors, the first-price auction achieves at least a $(n-1)/n$-fraction of the optimal expected revenue in equilibrium.

Revenue Equivalence

Revenue Equivalence

Theorem (revenue equivalence, informal): Any auction mechanism that has the same outcome has the same expected revenue in equilibrium.

Revenue Equivalence

Theorem (revenue equivalence, informal): Any auction mechanism that has the same outcome has the same expected revenue in equilibrium.

Theorem (revenue equivalence, more formal): Suppose that the values of the bidders are drawn iid from a distribution. Then any *symmetric* and *increasing* equilibrium of any direct revelation (auction) mechanism such that

Revenue Equivalence

Theorem (revenue equivalence, informal): Any auction mechanism that has the same outcome has the same expected revenue in equilibrium.

Theorem (revenue equivalence, more formal): Suppose that the values of the bidders are drawn iid from a distribution. Then any *symmetric* and *increasing* equilibrium of any direct revelation (auction) mechanism such that

1. the winner is the bidder with the highest bid,

Revenue Equivalence

Theorem (revenue equivalence, informal): Any auction mechanism that has the same outcome has the same expected revenue in equilibrium.

Theorem (revenue equivalence, more formal): Suppose that the values of the bidders are drawn iid from a distribution. Then any *symmetric* and *increasing* equilibrium of any direct revelation (auction) mechanism such that

1. the winner is the bidder with the highest bid,

2. the expected payment of any bidder with value 0 is 0,

Revenue Equivalence

Theorem (revenue equivalence, informal): Any auction mechanism that has the same outcome has the same expected revenue in equilibrium.

Theorem (revenue equivalence, more formal): Suppose that the values of the bidders are drawn iid from a distribution. Then any *symmetric* and *increasing* equilibrium of any direct revelation (auction) mechanism such that

1. the winner is the bidder with the highest bid,

2. the expected payment of any bidder with value 0 is 0,

has the *same expected revenue*.

For example

For example

<u>All-pay auction:</u> The winner is the highest bidder but every bidder pays their bid.

For example

All-pay auction: The winner is the highest bidder but every bidder pays their bid.

You would expect that this auction could make more money, since losers also pay.

For example

All-pay auction: The winner is the highest bidder but every bidder pays their bid.

You would expect that this auction could make more money, since losers also pay.

But *theoretically*, in equilibrium, it does not!

Back to Vickrey

This existence proofs ensure that an equilibrium exists.

So we can hope that the bidders are going to “find it” by iteratively adjusting their bids while maximising their utilities against the bids of the others.

But we would like to know more.

How do these equilibria look like? Can we describe them?

Back to Vickrey

This existence proofs ensure that an equilibrium exists.

So we can hope that the bidders are going to “find it” by iteratively adjusting their bids while maximising their utilities against the bids of the others.

But we would like to know more.

How do these equilibria look like? Can we describe them?

What Vickrey provided is called a “*closed form solution*”.

Beyond Uniform Distributions

Beyond Uniform Distributions

The literature has managed to produce (more complicated) closed form solutions for settings with symmetric beliefs.

Beyond Uniform Distributions

The literature has managed to produce (more complicated) closed form solutions for settings with symmetric beliefs.

For more general settings, we need to be more modest, as such closed form solutions may not exist.

Beyond Uniform Distributions

The literature has managed to produce (more complicated) closed form solutions for settings with symmetric beliefs.

For more general settings, we need to be more modest, as such closed form solutions may not exist.

- An alternative is describing an equilibrium via e.g., a set of differential equations.

Beyond Uniform Distributions

The literature has managed to produce (more complicated) closed form solutions for settings with symmetric beliefs.

For more general settings, we need to be more modest, as such closed form solutions may not exist.

> An alternative is describing an equilibrium via e.g., a set of differential equations.
>
> These are however not easy to solve…

Beyond Uniform Distributions

The literature has managed to produce (more complicated) closed form solutions for settings with symmetric beliefs.

For more general settings, we need to be more modest, as such closed form solutions may not exist.

An alternative is describing an equilibrium via e.g., a set of differential equations.

These are however not easy to solve…

… and often we cannot even get those!

Finding Equilibria in Auctions

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

This question is more intricate than it initially looks like.

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

This question is more intricate than it initially looks like.

Vickrey, Myerson, Milgrom, Wilson, and all the other great economists were not thinking about computation.

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

This question is more intricate than it initially looks like.

Vickrey, Myerson, Milgrom, Wilson, and all the other great economists were not thinking about computation.

Before we even attempt such an algorithm, we need to think about how to *represent* the *inputs* and *outputs* of our problem.

Discrete FPA

In the discrete FPA, the representation is easy.

Discrete FPA

In the discrete FPA, the representation is easy.

Representation of the input:

Discrete FPA

In the discrete FPA, the representation is easy.

Representation of the input:

Each v_i is given as a number in binary.

Discrete FPA

In the discrete FPA, the representation is easy.

Representation of the input:

Each v_i is given as a number in binary.

Each F_{ij} is given explicitly as pairs (v^i_{jk}, p^i_{jk}) (in binary) for every possible value v^i_{jk} (from the perspective of bidder i).

Discrete FPA

In the discrete FPA, the representation is easy.

Representation of the input:

Each v_i is given as a number in binary.

Each F_{ij} is given explicitly as pairs (v^i_{jk}, p^i_{jk}) (in binary) for every possible value v^i_{jk} (from the perspective of bidder i).

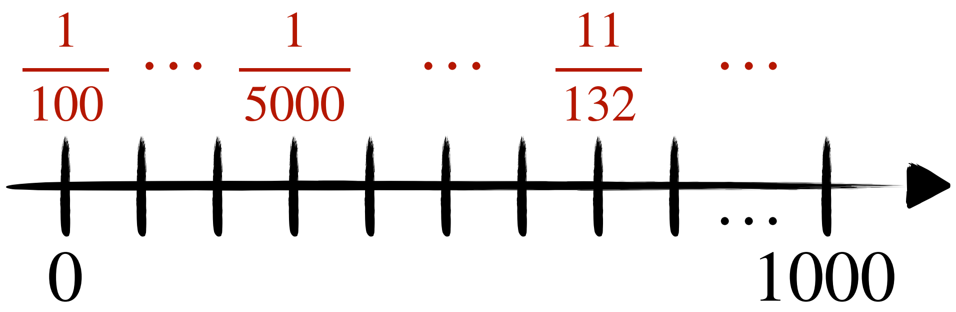

A discrete distribution

Discrete FPA

In the discrete FPA, the representation is easy.

Discrete FPA

In the discrete FPA, the representation is easy.

<u>Representation of the output:</u>

Discrete FPA

In the discrete FPA, the representation is easy.

Representation of the output:

The output is the β_i function for each bidder i.

Discrete FPA

In the discrete FPA, the representation is easy.

Representation of the output:

The output is the β_i function for each bidder i.

This is given explicitly as a vector $(\beta_i(v_{i1}), \beta_i(v_{i2}), \ldots, \beta_i(v_{in}))$.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the input:

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the input:

How do we represent the distributions F_{ij}? These are now continuous functions.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the input:

How do we represent the distributions F_{ij}? These are now continuous functions.

We will restrict attention to distributions for which there is a natural representation.

Representable Distributions

Representable Distributions

Piece-wise constant densities:

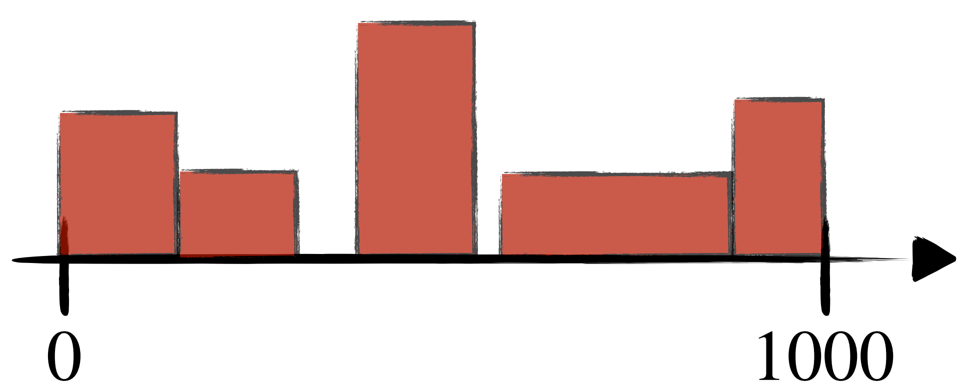

Representable Distributions

Piece-wise constant densities:

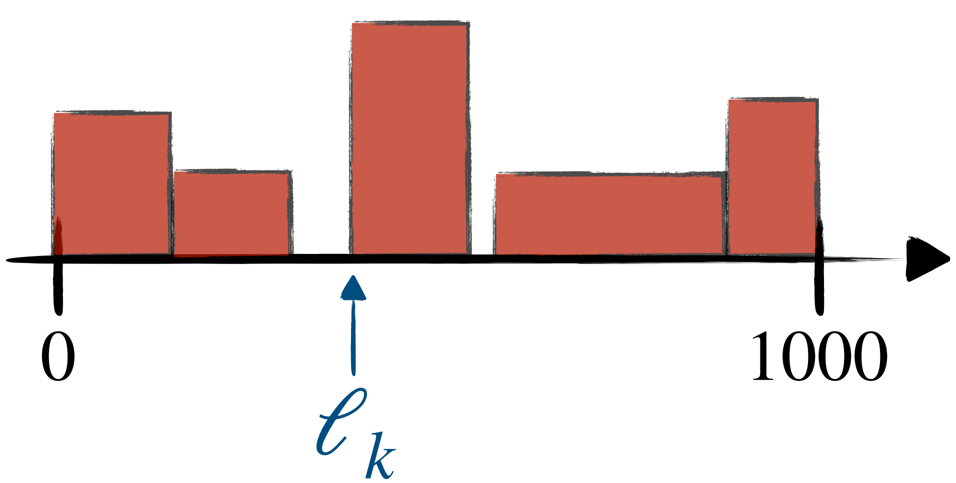

Representable Distributions

Piece-wise constant densities:

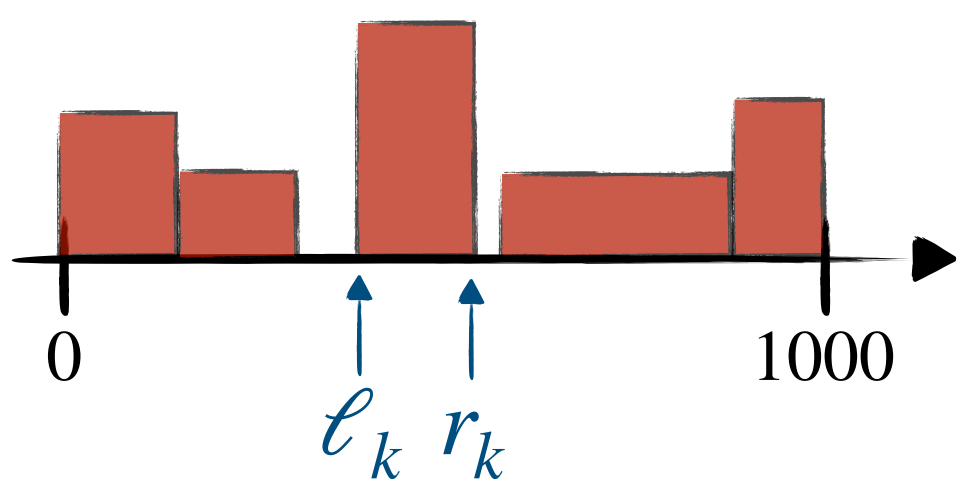

Representable Distributions

Piece-wise constant densities:

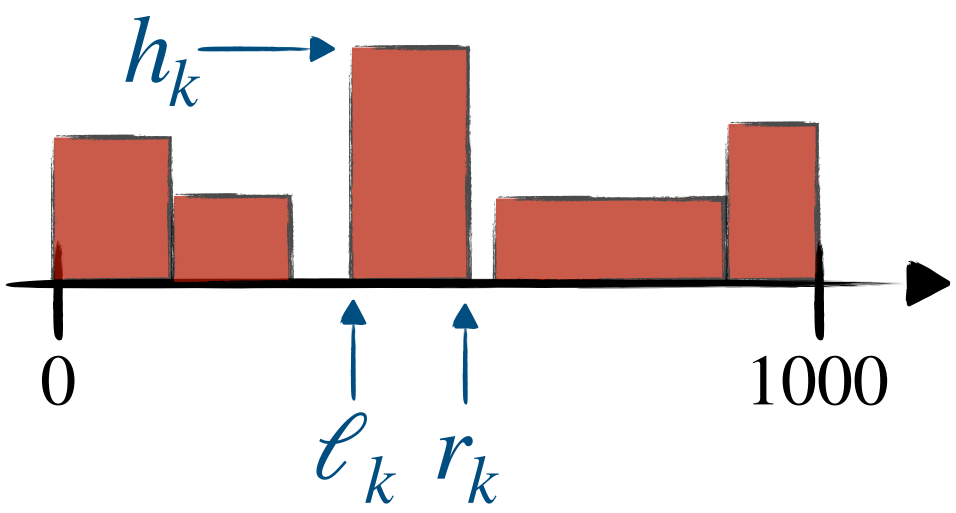

Representable Distributions

Piece-wise constant densities:

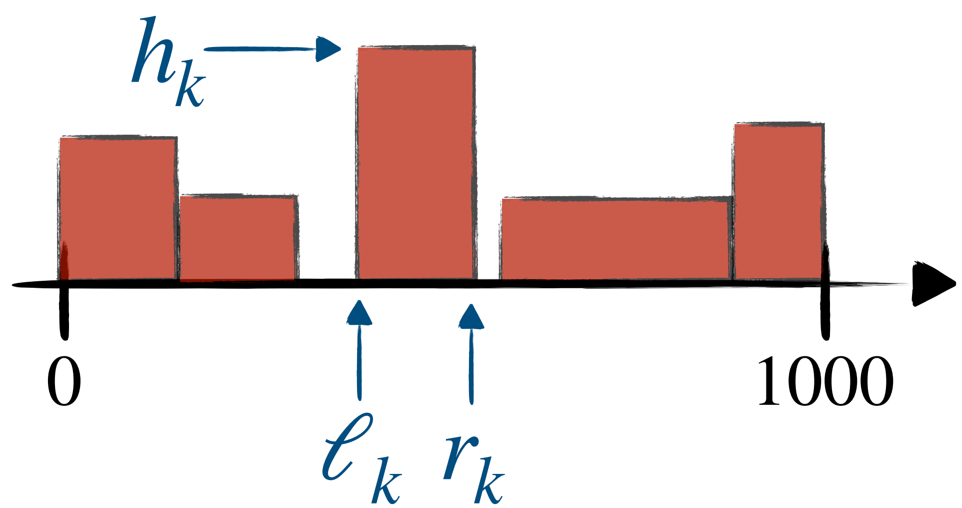

Piece-wise linear densities:

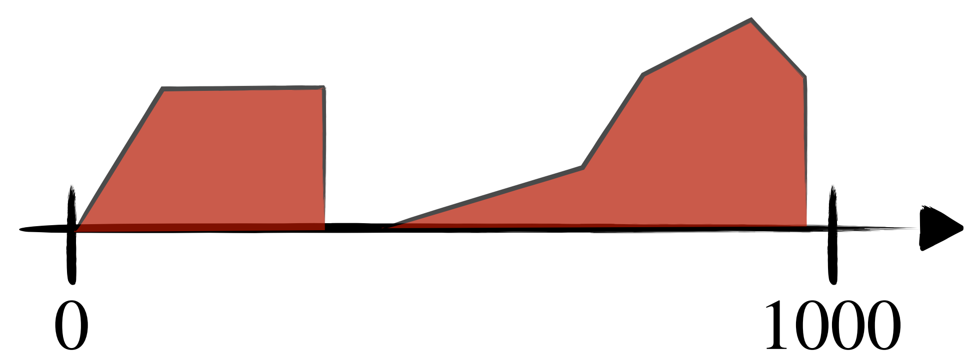

Representable Distributions

Piece-wise constant densities:

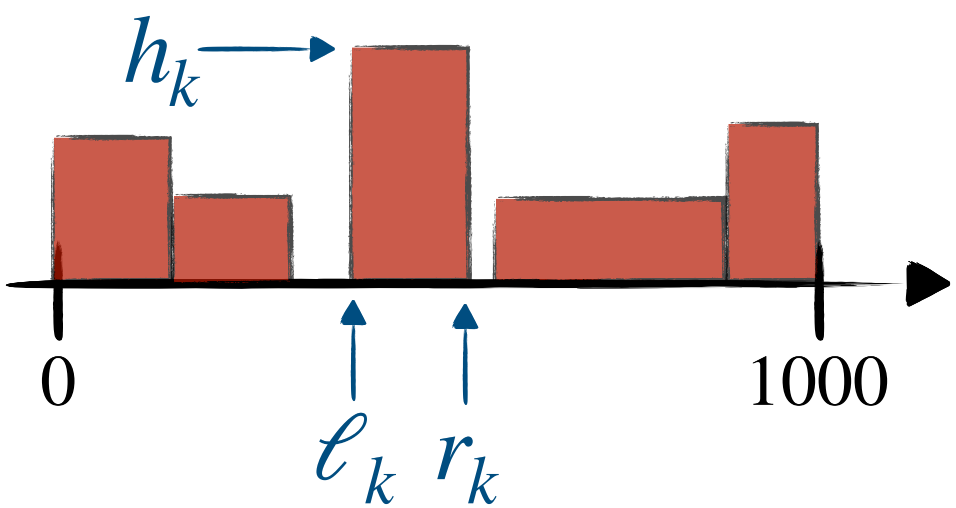

Piece-wise linear densities:

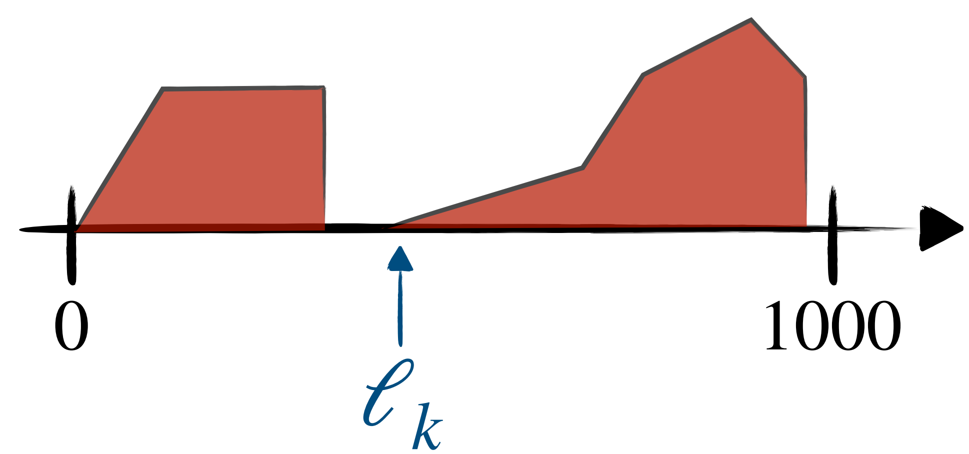

Representable Distributions

Piece-wise constant densities:

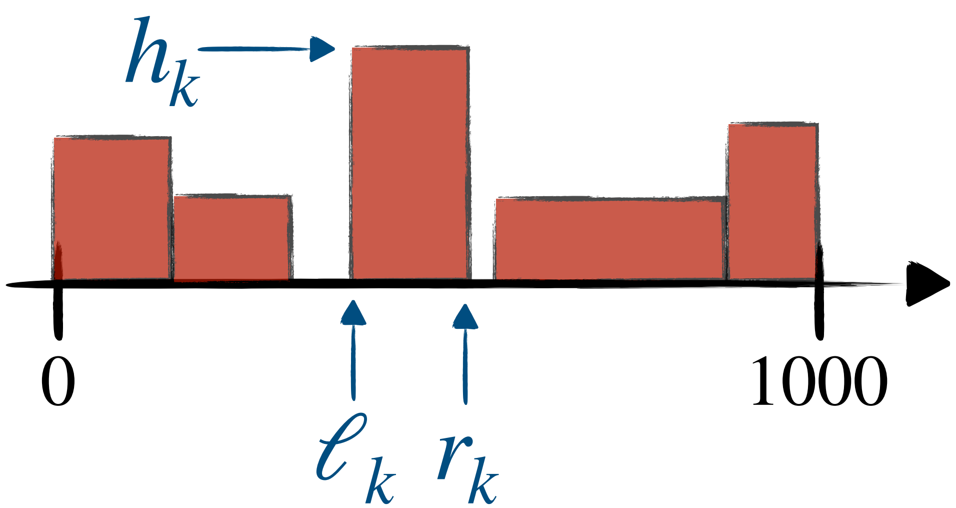

Piece-wise linear densities:

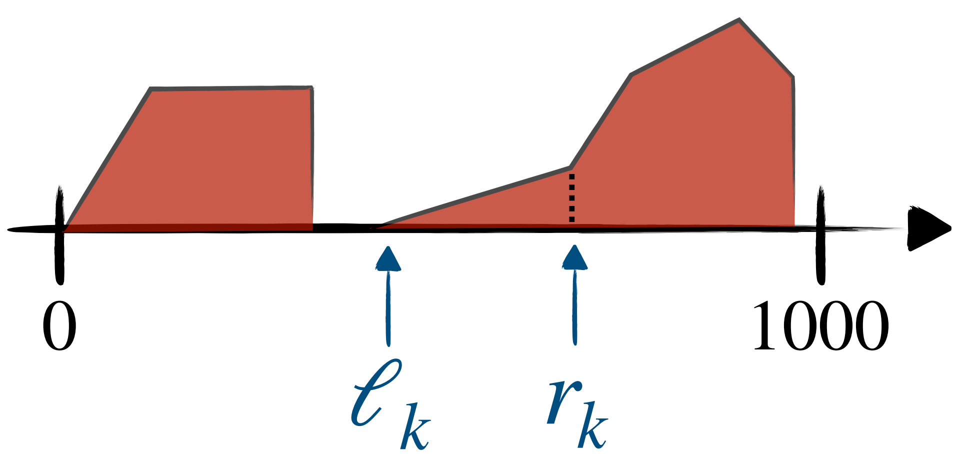

Representable Distributions

Piece-wise constant densities:

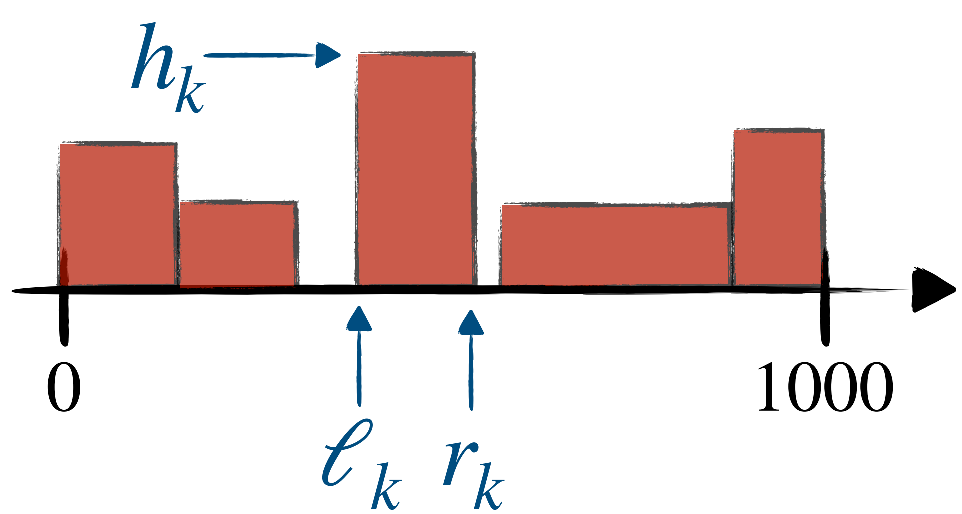

Piece-wise linear densities:

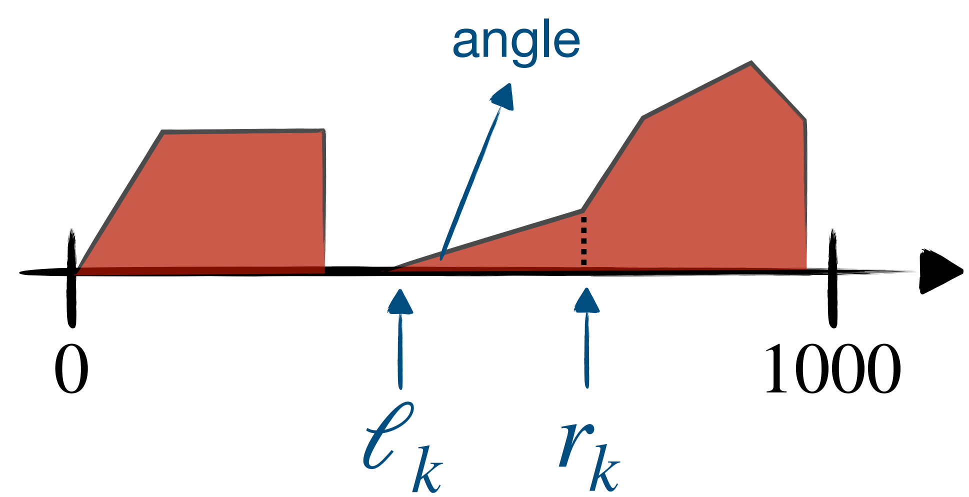

Representable Distributions

Piece-wise constant densities:

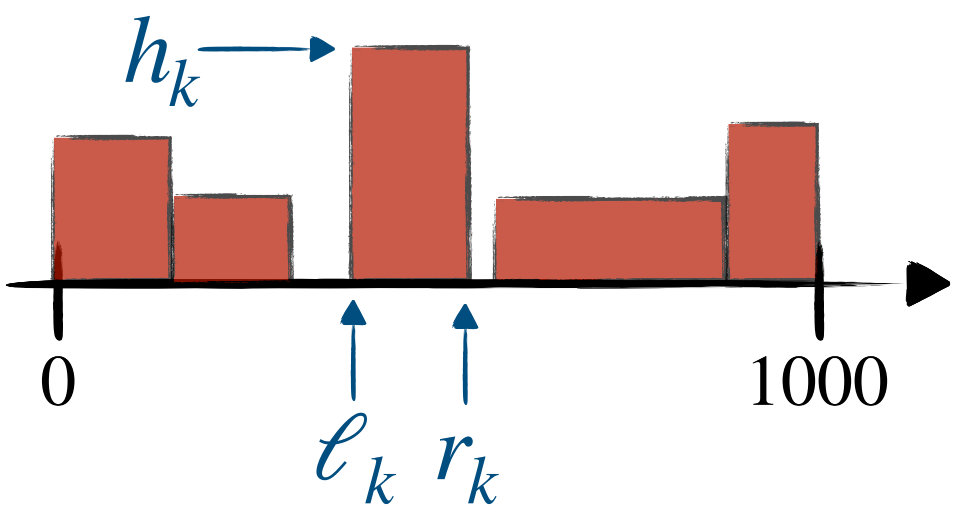

Piece-wise linear densities:

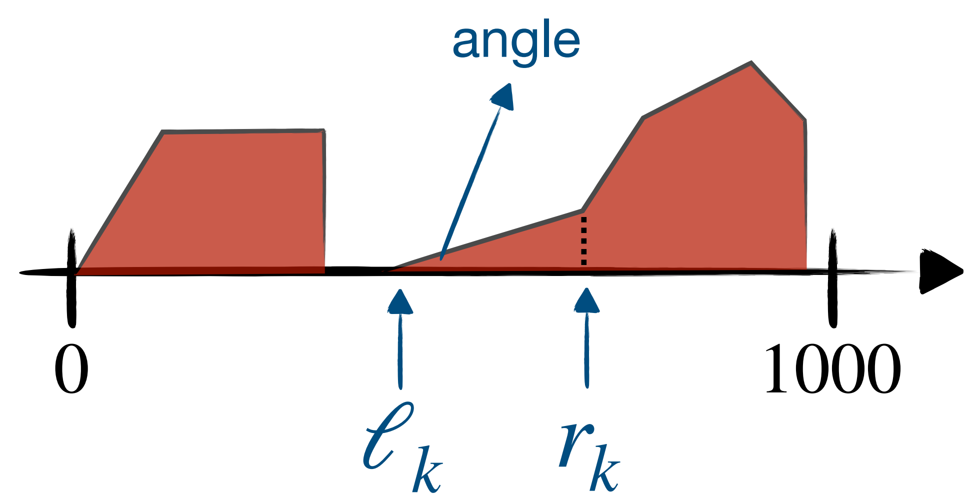

Piece-wise polynomial densities:

Representable Distributions

Piece-wise constant densities:

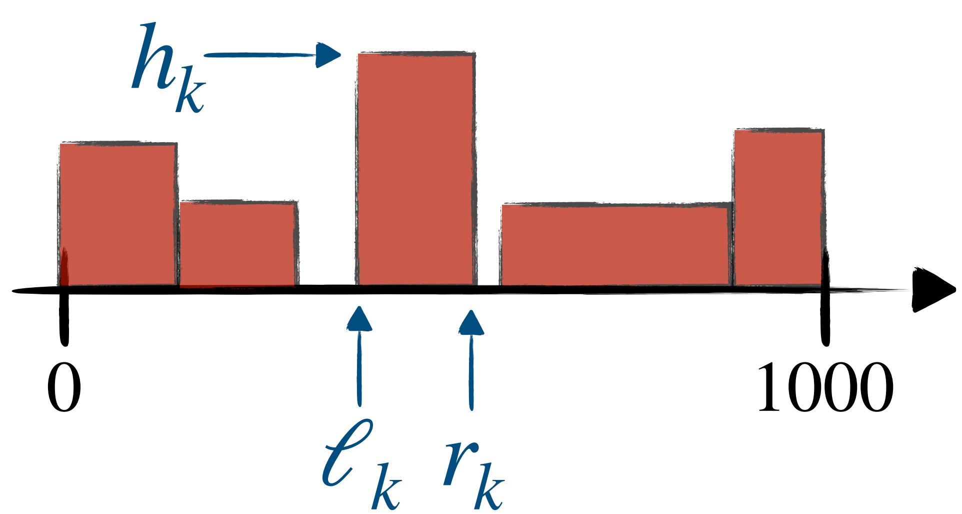

Piece-wise linear densities:

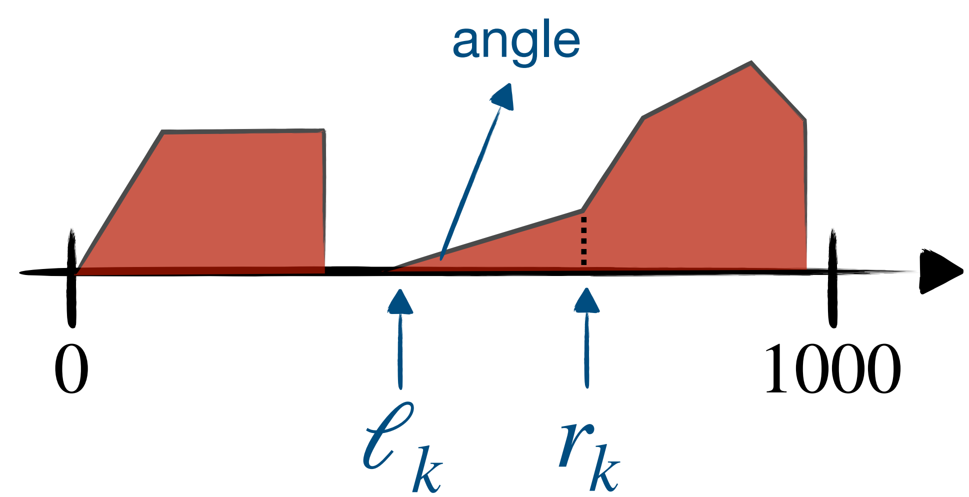

Piece-wise polynomial densities:

For the k-th subinterval $[\ell_k, r_k]$, a list of (rational) coefficients of the polynomial.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the output:

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the output:

Our functions β_i now have a continuous domain.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the output:

Our functions β_i now have a continuous domain.

We cannot simply represent them as a finite map, because we have infinitely many values v_i for which we need to specify $\beta_i(v_i)$.

Useful observations

Useful observations

Discrete bidding space: In all applications, the bids that a bidder can submit come from a finite set, e.g., multiples of one 1p.

Useful observations

Discrete bidding space: In all applications, the bids that a bidder can submit come from a finite set, e.g., multiples of one 1p.

Increasing bidding function: A bidding function is increasing if

higher values ⇒ higher bids

Useful observations

Discrete bidding space: In all applications, the bids that a bidder can submit come from a finite set, e.g., multiples of one 1p.

Increasing bidding function: A bidding function is increasing if higher values ⇒ higher bids

formally, $v_{ij} \geq v_{ik} \Rightarrow \beta_i(v_{ij}) \geq \beta_i(v_{ik})$

Useful observations

Discrete bidding space: In all applications, the bids that a bidder can submit come from a finite set, e.g., multiples of one 1p.

Increasing bidding function: A bidding function is increasing if higher values ⇒ higher bids

formally, $v_{ij} \geq v_{ik} \Rightarrow \beta_i(v_{ij}) \geq \beta_i(v_{ik})$

Fact (Athey 2001): The continuous FPA has Nash equilibria where all the bidding functions are increasing.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the output:

Our functions β_i now have a continuous domain.

We cannot simply represent them as a finite map, because we have infinitely many values v_i for which we need to specify $\beta_i(v_i)$.

Continuous FPA

In the continuous FPA, the representation requires more thought.

Representation of the output:

Our functions β_i now have a continuous domain.

We cannot simply represent them as a finite map, because we have infinitely many values v_i for which we need to specify $\beta_i(v_i)$.

We can specify for each bid b_j, the value $v_{i\ell}$ of the bidder for which $\beta_i(v_{i\ell}) = b_j$.

Continuous FPA

We can specify for each bid b_j, the value $v_{i\ell}$ of the bidder for which $\beta_i(v_{i\ell}) = b_j$.

Continuous FPA

We can specify for each bid b_j, the value $v_{i\ell}$ of the bidder for which $\beta_i(v_{i\ell}) = b_j$.

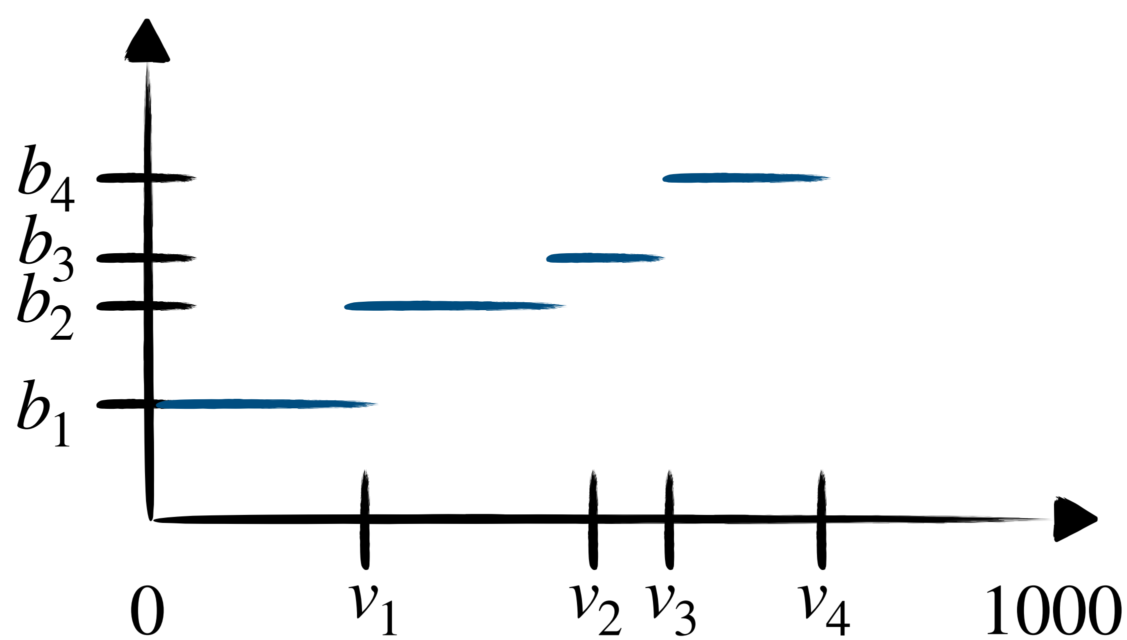

Continuous FPA

We can specify for each bid b_j, the value $v_{i\ell}$ of the bidder for which $\beta_i(v_{i\ell}) = b_j$.

We call this a "jump point" representation.

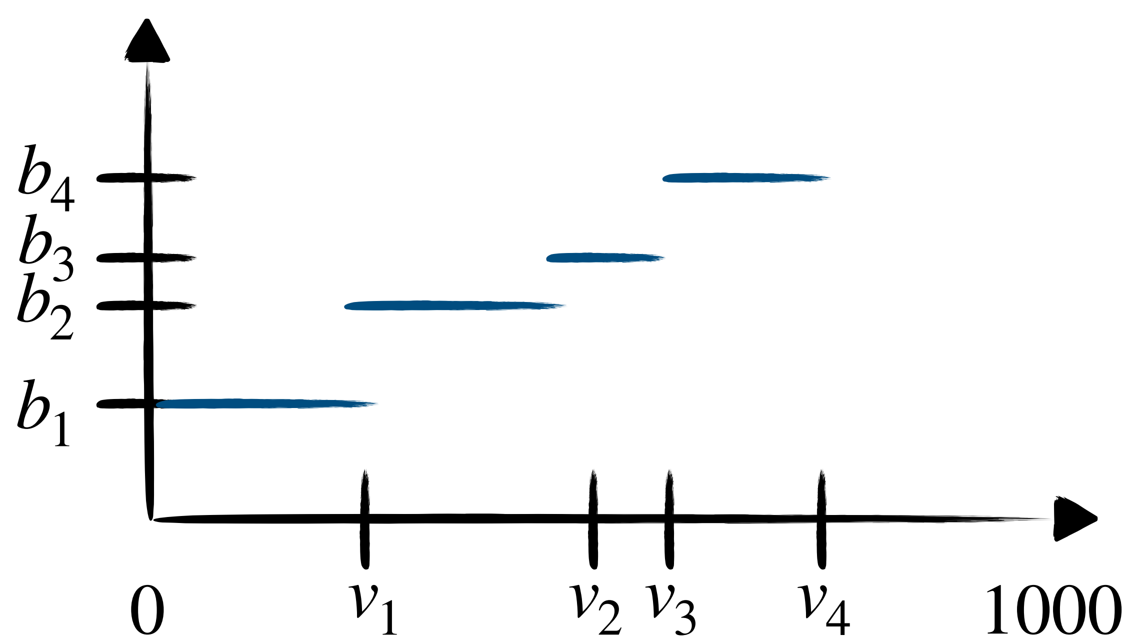

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

This question is more intricate than it initially looks like.

Vickrey, Myerson, Milgrom, Wilson, and all the other great economists were not thinking about computation.

Before we even attempt such an algorithm, we need to think about how to *represent* the *inputs* and *outputs* of our problem.

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

This question is more intricate than it initially looks like.

Vickrey, Myerson, Milgrom, Wilson, and all the other great economists were not thinking about computation.

Before we even attempt such an algorithm, we need to think about how to *represent* the *inputs* and *outputs* of our problem.

Finding Equilibria in Auctions

Finding Equilibria in Auctions

We are not done with the intricacies…

Finding Equilibria in Auctions

We are not done with the intricacies…

What if in every Nash equilibrium, there is some number that is irrational?

Finding Equilibria in Auctions

We are not done with the intricacies…

What if in every Nash equilibrium, there is some number that is irrational?

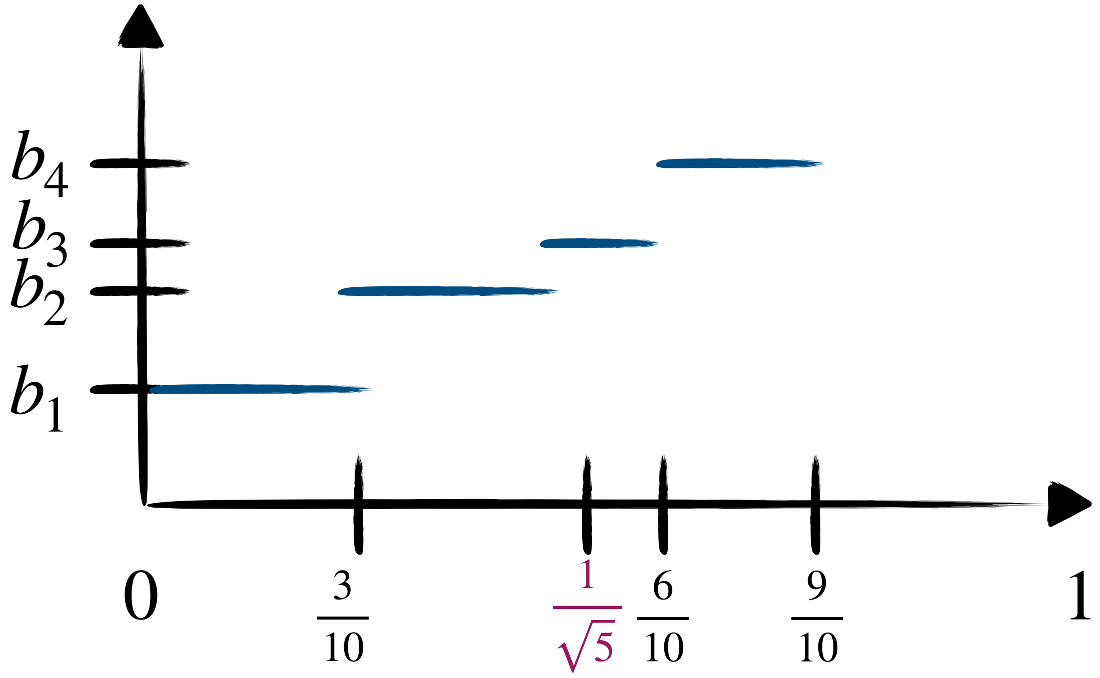

Finding Equilibria in Auctions

We are not done with the intricacies…

What if in every Nash equilibrium, there is some number that is irrational?

Could this happen?

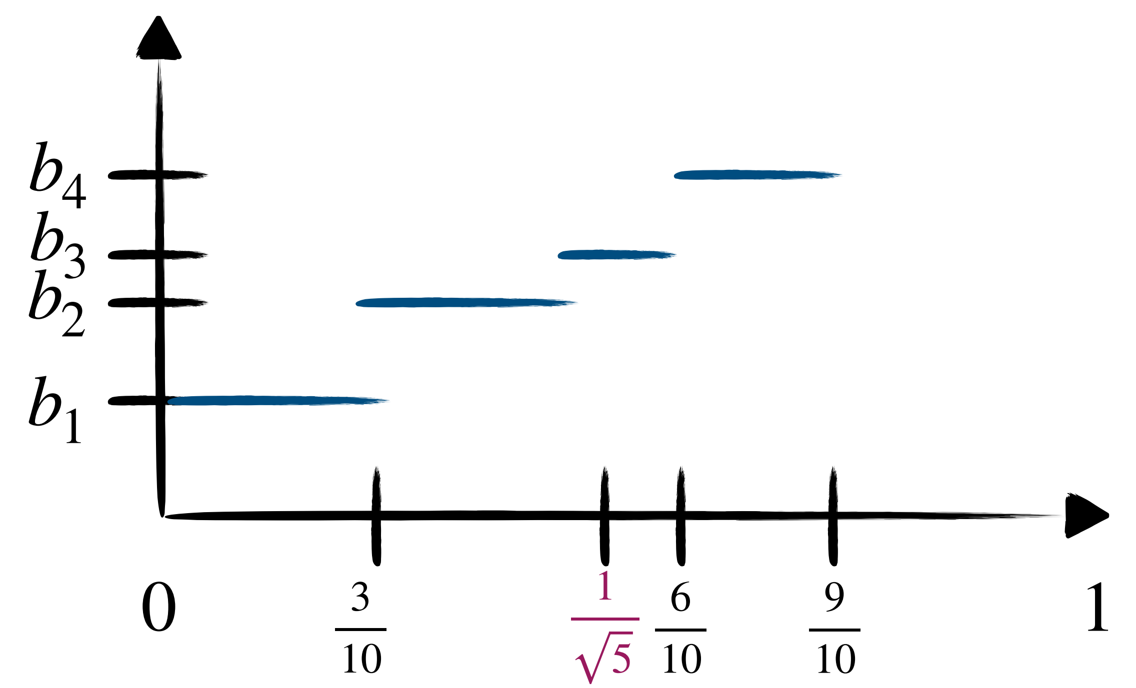

Finding Equilibria in Auctions

We are not done with the intricacies…

What if in every Nash equilibrium, there is some number that is irrational?

Could this happen?

It could!

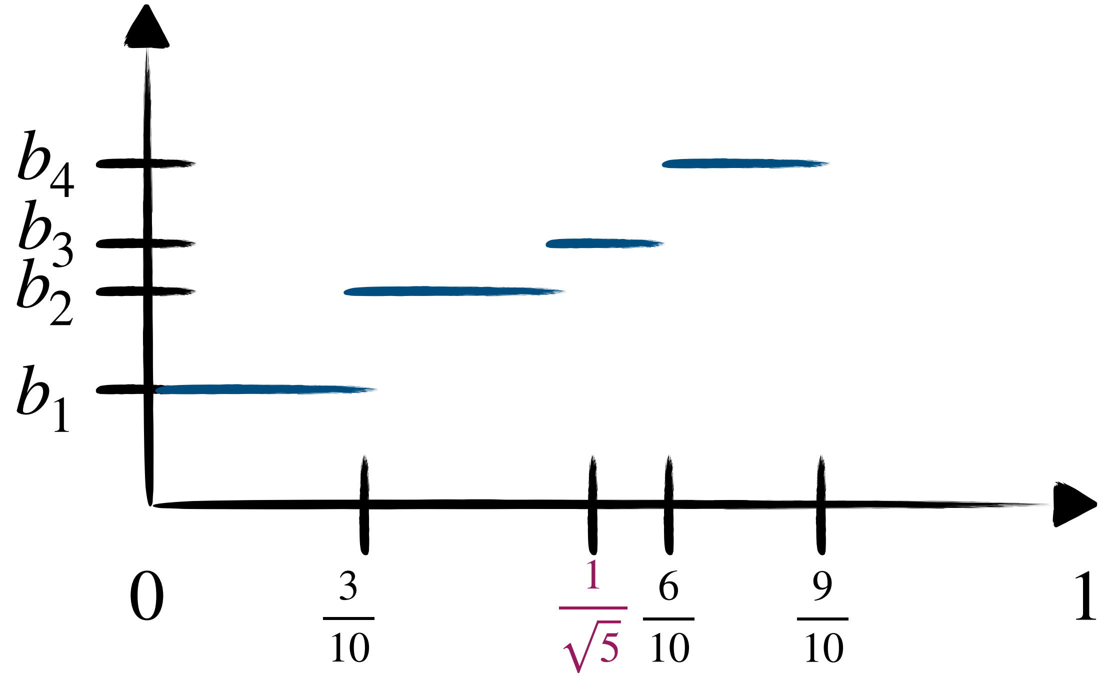

Finding Equilibria in Auctions

We are not done with the intricacies…

What if in every Nash equilibrium, there is some number that is irrational?

Could this happen?

It could!

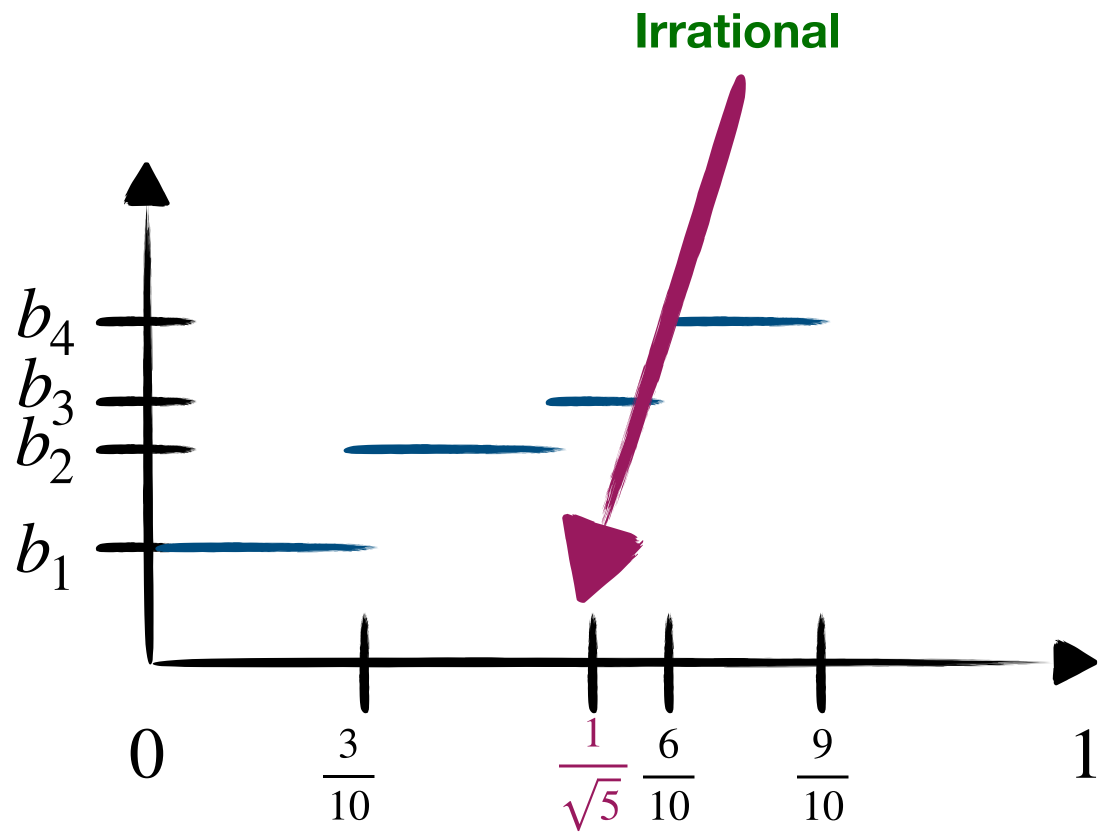

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

Now let's make this question more precise. Recall:

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

Now let's make this question more precise. Recall:

> If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

Now let's make this question more precise. Recall:

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

Discrete FPA problem: Given a discrete FPA and an ε as input, decide if the auction has an ε-approximate Nash equilibrium or not. If it does, return it.

Theorem (F., Giannakopoulos, Hollender, and Kokkalis 2024): The Discrete FPA problem is NP-complete.

Finding Equilibria in Auctions

Concrete CS problem: Given a FPA as an input, can we design an efficient algorithm for finding a Nash equilibrium of the auction?

Now let's make this question more precise. Recall:

If the distributions are discrete, then a Nash equilibrium might not exist (Maskin and Riley 1985).

Discrete FPA problem: Given a discrete FPA and an ε as input, decide if the auction has an ε-approximate Nash equilibrium or not. If it does, return it.

Finding Equilibria in Auctions

Continuous FPA problem: Given a continuous FPA and an $\mathcal{E}$ as input, find an $\mathcal{E}$-approximate Nash equilibrium of the auction.

Finding Equilibria in Auctions

Continuous FPA problem: Given a continuous FPA and an $\mathcal{E}$ as input, find an $\mathcal{E}$-approximate Nash equilibrium of the auction.

Theorem (F., Giannakopoulos, Hollender, Lazos, and Poças 2023):
The Discrete FPA problem is PPAD-complete.

Any hope for efficient algorithms?

Any hope for efficient algorithms?

The aforementioned (NP- and PPAD-) hardness results are evidence that we cannot design efficient algorithms for finding Nash equilibria in the FPA *in general*.

Any hope for efficient algorithms?

The aforementioned (NP- and PPAD-) hardness results are evidence that we cannot design efficient algorithms for finding Nash equilibria in the FPA *in general*.

<u>But we are actually far from fully understanding the complexity of the problem:</u>

Any hope for efficient algorithms?

The aforementioned (NP- and PPAD-) hardness results are evidence that we cannot design efficient algorithms for finding Nash equilibria in the FPA *in general*.

But we are actually far from fully understanding the complexity of the problem:

The hardness results require subjective beliefs $(F_{ij} \neq F_{\ell j})$.

Any hope for efficient algorithms?

The aforementioned (NP- and PPAD-) hardness results are evidence that we cannot design efficient algorithms for finding Nash equilibria in the FPA *in general*.

But we are actually far from fully understanding the complexity of the problem:

The hardness results require subjective beliefs $(F_{ij} \neq F_{\ell j})$.

For objective beliefs $(F_{ij} = F_{\ell j})$, we don't know similar hardness results or polynomial-time algorithms.

Any hope for efficient algorithms?

The aforementioned (NP- and PPAD-) hardness results are evidence that we cannot design efficient algorithms for finding Nash equilibria in the FPA *in general*.

<u>But we are actually far from fully understanding the complexity of the problem:</u>

The hardness results require subjective beliefs $(F_{ij} \neq F_{\ell j})$.

For objective beliefs $(F_{ij} = F_{\ell j})$, we don't know similar hardness results or polynomial-time algorithms.

For symmetric beliefs $(F_i = F_j)$, we have polynomial-time algorithms.
(F. et al. 2023, 2024, 2025, 2026).