# Algorithmic Game Theory and Applications

(Approximate) Mechanism Design on Restricted Domains

# The Gibbard-Satterthwaite Theorem

Theorem (Gibbard 73 - Satterthwaite 75): In the unrestricted domain, when there are $m \geq 3$ candidates, a voting rule is truthful and onto if and only if it is dictatorial.

# Escaping the Gibbard-Satterthwaite Theorem

Theorem (Gibbard 73 - Satterthwaite 75): In the unrestricted domain, when there are $m \geq 3$ candidates, a voting rule is truthful and onto if and only if it is dictatorial.

# Unrestricted Domain

A *social choice function*, or *voting rule*, or *mechanism* is a function $f:(\succ)^n \rightarrow A$ mapping preference profiles to candidates,

where $\succ^n$ is the space of all possible preference profiles.

The unrestricted domain: $\succ^n$ can contain any preference profile.

i.e., for any voter $i \in N$, $\succ_i$ is the set of *all permutations* of $\{1,\ldots,m\}$.

# Single-Peaked Preferences

Assume that we have a set of possible temperatures for the thermostat, e.g., $\{-10, -5, 0, 5, 10, 15, 20, 25, 30, 35, 40\}$.

Let's say that your ideal temperature would be $20$ degrees.

It is reasonable to assume that you would also prefer $25$ degrees to $30$ degrees, and likewise, $15$ degrees to $10$ degrees.

Generally, the "farther away" we move from your ideal temperature, the less happy you become.

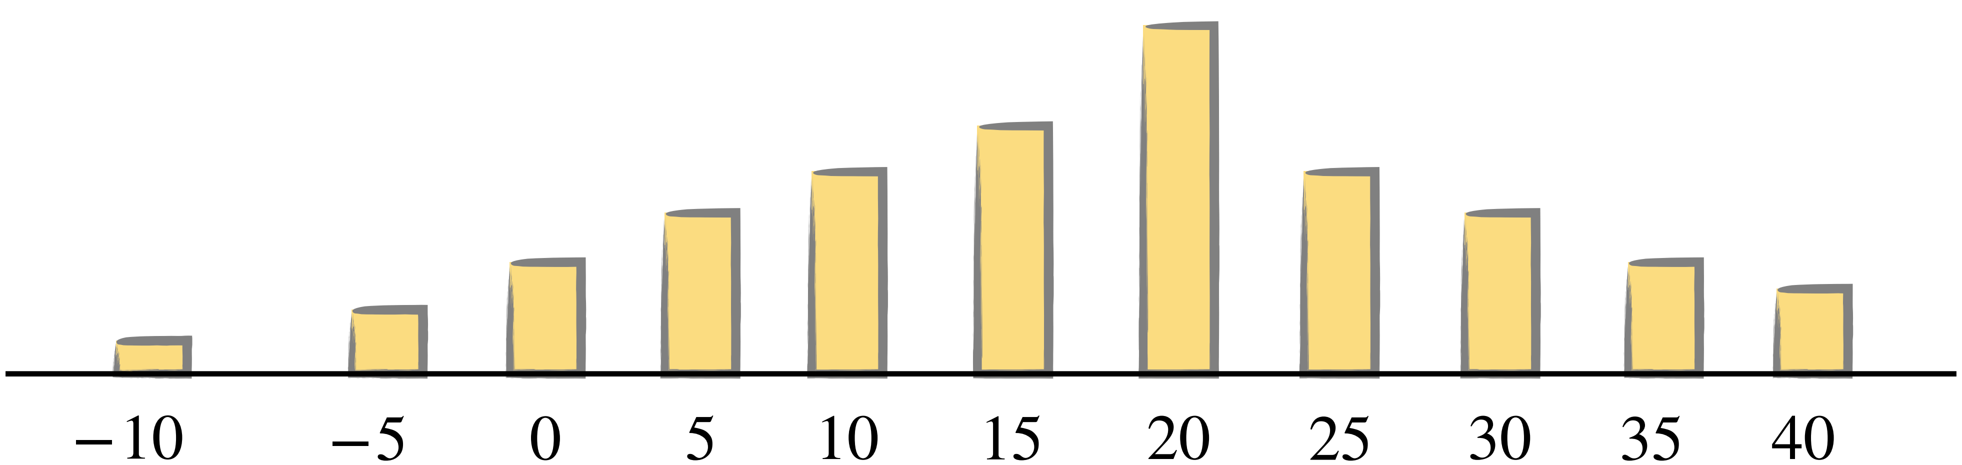

# Single-Peaked Preferences

Other applications:

Political spectrum (from left to right, from conservative to progressive etc).

Building a library on a street (facility location).

Introduced by Black in 1948, as a domain for which Condorcet winners always exist.

# Find the Condorcet winner

Recall: A Condorcet winner wins a pairwise majority election against any other candidate.

$\text{MAJ}(x_1, x_2) = x_2$

$\text{MAJ}(x_2, x_3) = x_3$

$\text{MAJ}(x_1, x_3) = x_3$

$\text{MAJ}(x_3, x_4) = x_4$

$\text{MAJ}(x_1, x_4) = x_4$

$\text{MAJ}(x_2, x_4) = x_4$

$\text{MAJ}(x_4, x_5) = x_4$

$\text{MAJ}(x_4, x_6) = x_4$

$\text{MAJ}(x_4, x_7) = x_4$

$x_1 \quad x_2 \quad x_3 \quad x_4 \quad x_5 \quad x_6 \quad x_7$

$x_4$ is a Condorcet winner among the peaks.

What else is $x_4$?

# Median Voter Rule

Consider a social choice setting in which the preferences $\succ_i$ of the voters are single-peaked, and let $x_i$ be the peak of voter $i$.

Median Voter Rule: Select the median of the (reported) peaks $x_i$, i.e.,

$$f(\succ) = \text{med}\{x_1, x_2, \ldots, x_n\}$$

Theorem: The median voter rule is truthful.

Proof: Easy, we've seen it before.

# Median Voter Rule

The proof follows from the same monotonicity argument: To affect the median, a voter with $x_i \leq f(\succ)$ must report something to the right of the median. But then the median will move farther away from $x_i$.

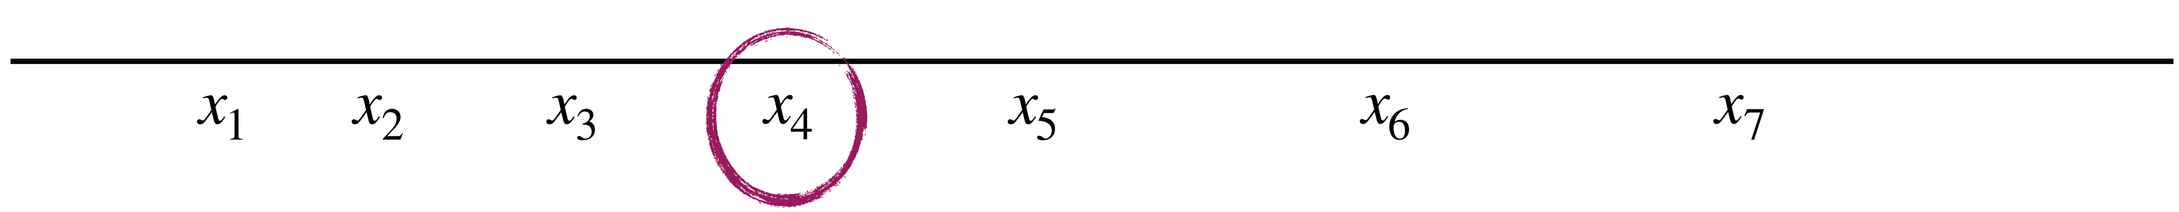

# Ordered Statistic Voter Rule

Consider a social choice setting in which the preferences $\succ_i$ of the voters are single-peaked, and let $x_i$ be the peak of voter $i$.

$k$-th Order Statistic Voter Rule: Select the $k$-th ordered statistic of the (reported) peaks $x_i$, i.e.,

$$f(\succ) = \{x_i : x_i \text{ is at least as large as exactly } k \text{ peaks.}\}$$

Theorem: For any $k$, the $k$-th order statistic voter rule is truthful.

Proof: Virtually identical to before, check at home.

# Let's pretend to be economists for a while

1. We want a voting rule that is "good", according to some definition of "good".

   In economics, a rule is "good" when it satisfies certain desirable properties (axioms): here truthfulness and onto.

   So, according to our economics interpretation, every $k$-th order statistic voter rule is good.

2. We want to identify (or "characterise") all *good* voting rules.

   Here, we would like to prove a theorem that says that a voting rule is truthful and onto if and only if it looks like *something*.

   e.g., *"A voting rule is truthful and onto if and only if it is a $k$-th order statistic voter rule."*

# Median Voter Rule

Here, $x_i$ are the peaks, $y_i$ are the other possible candidates, which are not peaks of any voter.

Here in fact, the median candidate is the Condorcet winner among all candidates (not just the peaks).

We can also have any $k$-th ordered statistic among all the candidates.

# A characterisation

We will need one more natural property.

Property (Anonymity): A voting rule $f$ is anonymous if renaming the voters does not change the outcome.

Formally, for any $\succ$ and any permutation $\succ'$ of $\succ$ (when $\succ$ is seen as a vector), we have $f(\succ) = f(\succ')$.

Theorem (Moulin 1980): A voting rule $f$ is truthful, onto, and anonymous if and only if there exist $y_1, y_2, \ldots, y_{n-1}$ such that for all $\succ$, it holds that

$$f(\succ) = \text{med}\{x_1, x_2, \ldots, x_n, y_1, \ldots, y_{n-1}\}$$

# Towards a characterisation

Theorem (Moulin 1980): A voting rule $f$ is truthful, onto, and anonymous if and only if there exist $y_1, y_2, \ldots, y_{n-1}$ such that for all $\succ$, it holds that

$$f(\succ) = \text{med}\{x_1, x_2, \ldots, x_n, y_1, \ldots, y_{n-1}\}$$

There is also a characterisation without the anonymity property, which is slightly more complicated ("Generalised Median Voter Schemes").

Intuitively, some voters have more "power" than others.

If you are interested, check the AGT book Definition 10.3.

# Let's become Computer Scientists again

1. We want a voting rule that is "good", according to some definition of "good".

In economics, a rule is "good" when it satisfies certain desirable properties (axioms): here truthfulness and onto.

In computer science, we usually aim to optimise some global objective, e.g., maximise the social welfare, or minimise the social cost.

Obviously we still want the voting rule to be robust to incentives, so we are still interested in truthfulness.

# Let's become Computer Scientists again

General Question: Among all the truthful voting rules, which one is the best with respect to the global objective?

# A brief note about approximation in CS

Approximation algorithms for intractable problems:

- We are faced with an optimisation problem which is NP-hard (e.g., MAX-SAT), so we cannot solve it exactly in polynomial time, unless P=NP.
- We design polynomial-time algorithms which do not achieve an optimal solution, but an *approximation* to it.
- The approximation ratio measures the value of the optimal over the value of our algorithm (for maximisation problems) or the inverse of this ratio (for minimisation problems), taken worst-case over all the possible inputs to the problem.

# Approximation in mechanism design

Approximation algorithms for mechanism design problems:

- We are faced with an optimisation problem which we could solve optimally if the agents were being honest and not strategic.
- We design *truthful mechanisms* which do not achieve an optimal solution, but an *approximation* to it.
- The approximation ratio measures the value of the optimal over the value of our algorithm (for maximisation problems) or the inverse of this ratio (for minimisation problems), taken worst-case over all the possible inputs to the problem.

# Let's become Computer Scientists again

General Question: Among all the truthful voting rules, which one is the best with respect to the global objective?

Refined Question: Among all the truthful voting rules, or, in this context, *mechanisms*, what is the one with the smallest possible approximation ratio?

# Truthful Facility Location (Procaccia and Tennenholtz 2010)

There is a set $N = \{1,\ldots,n\}$ of agents (voters), each of which has an ideal location (the "peak") $x_i$ on the real line $\mathbb{R}$.

We want to place a facility at some location $y \in \mathbb{R}$. Any location $y \in \mathbb{R}$ is a possible candidate.

Given a location $y \in \mathbb{R}$, the cost of agent $i$ is defined as $|y - x_i|$, i.e., the distance between its peak and the location $y$.

A mechanism asks the agents to report their peaks $x_i$, and outputs a location $y = f(x_1, \ldots, x_n)$.

Each agent aims to minimise its cost and reports its peak as $\hat{x}_i$ accordingly.

We want to design a truthful mechanism (voting rule) for the problem that has the minimum possible approximation ratio for the social cost objective, i.e., the sum of agents' costs $\sum_{i \in N} |y - x_i|$

# Example 2:
# Setting the temperature

The reports shown in the picture are the peaks, but any temperature is a possible outcome.

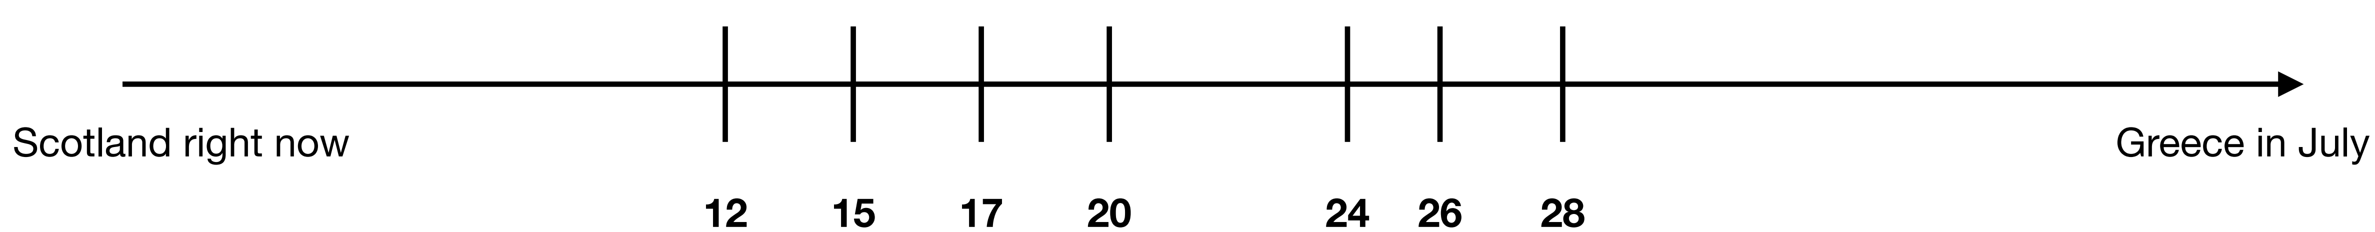

# Single-Peaked Preferences

How is the facility location setting different from this?

Assume that we have a set of possible temperatures for the thermostat, e.g., $\{-10, -5, 0, 5, 10, 15, 20, 25, 30, 35, 40\}$.

Let's say that your ideal temperature would be $20$ degrees.

It is reasonable to assume that you would also prefer $25$ degrees to $30$ degrees, and likewise, $15$ degrees to $10$ degrees.

Generally, the "farther away" we move from your ideal temperature, the less happy you become.

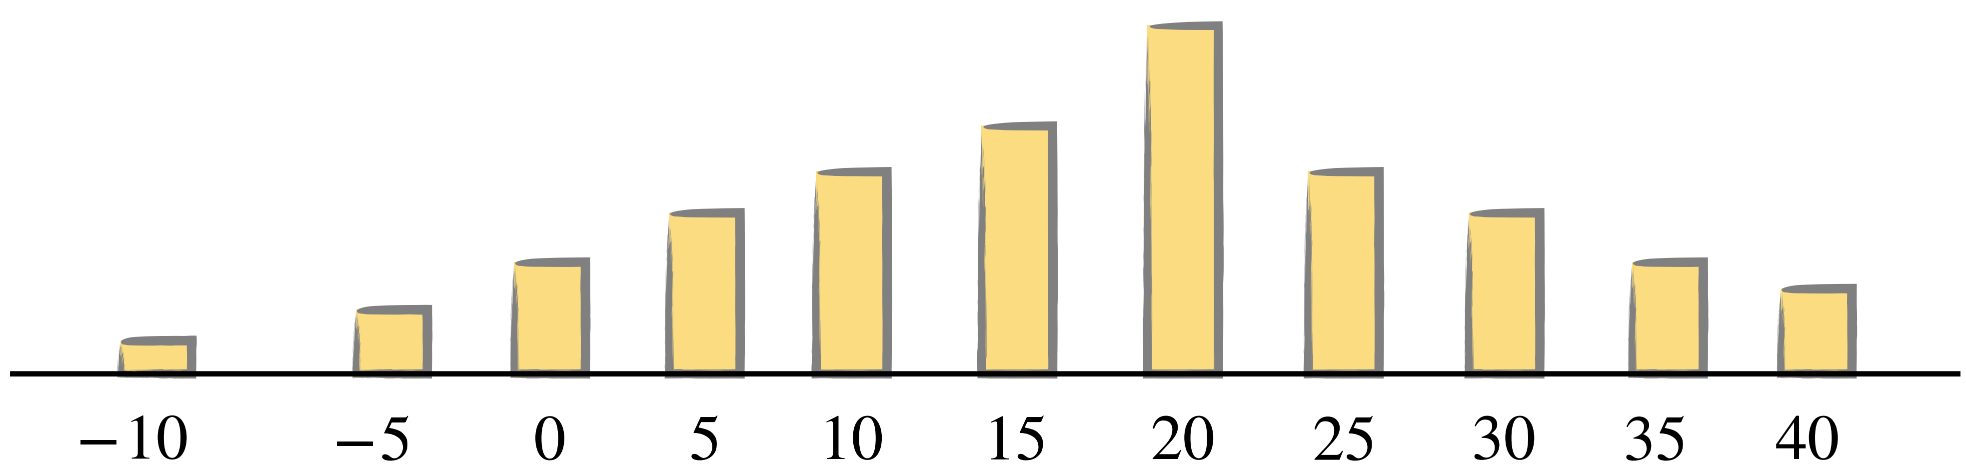

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let's consider any other location $z \in \mathbb{R}$

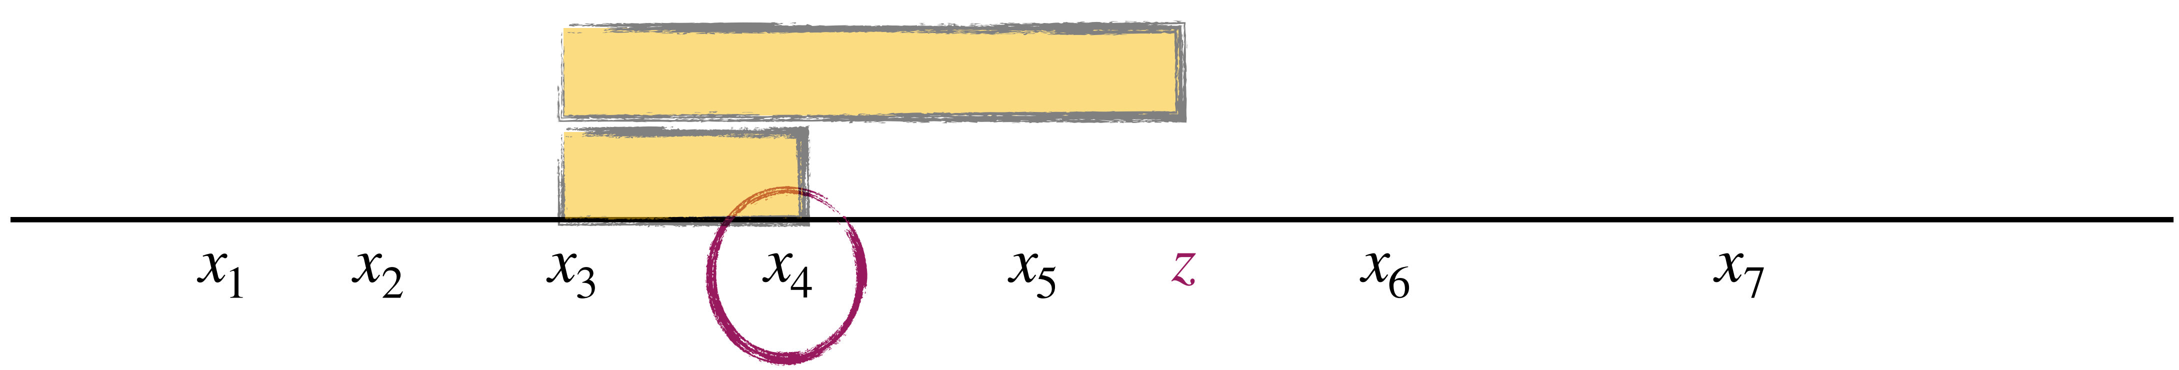

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let's consider any other location $z \in \mathbb{R}$

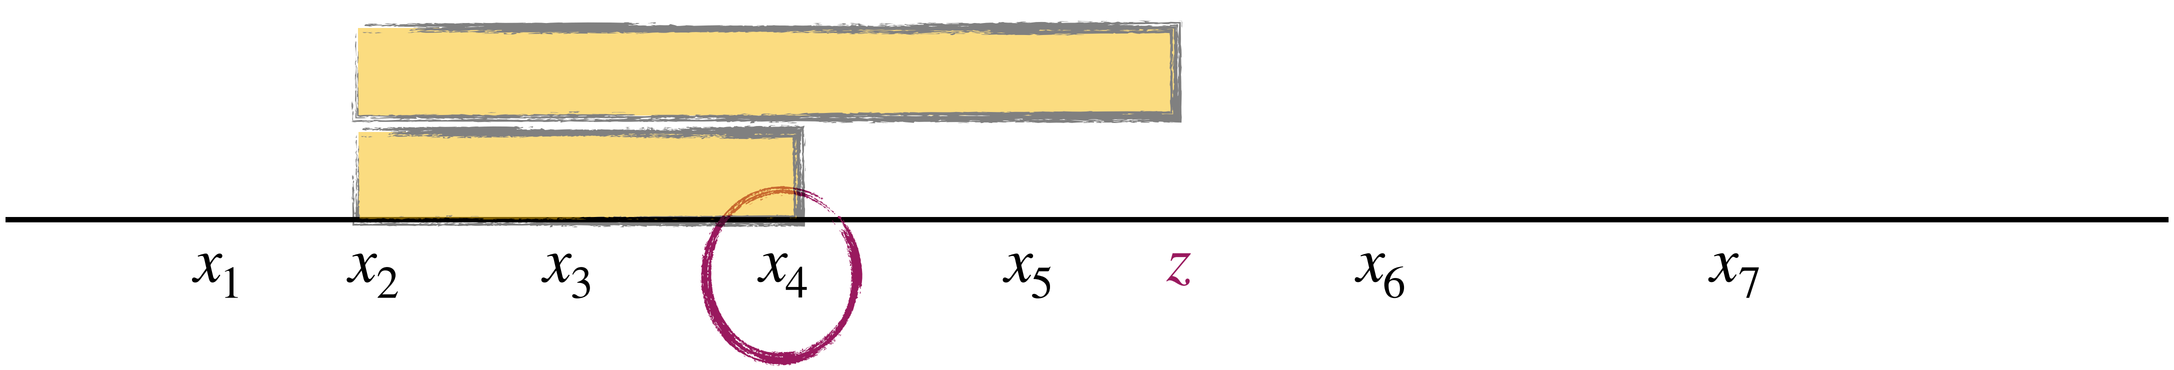

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let's consider any other location $z \in \mathbb{R}$

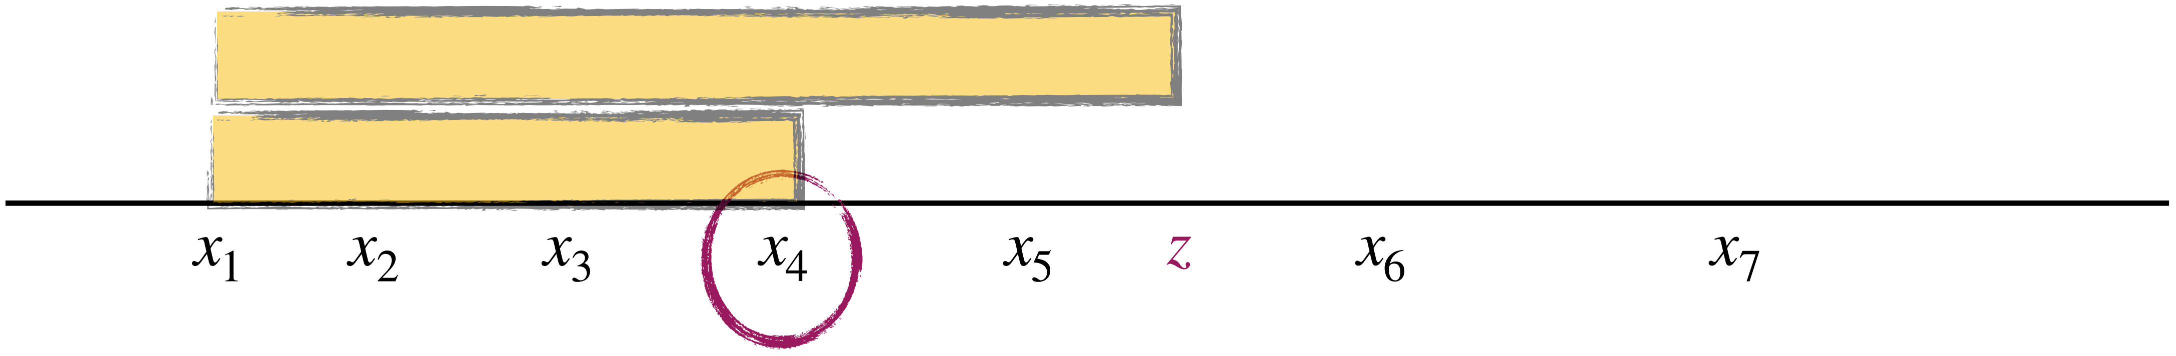

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let's consider any other location $z \in \mathbb{R}$

At least half of the agents have a cost surplus

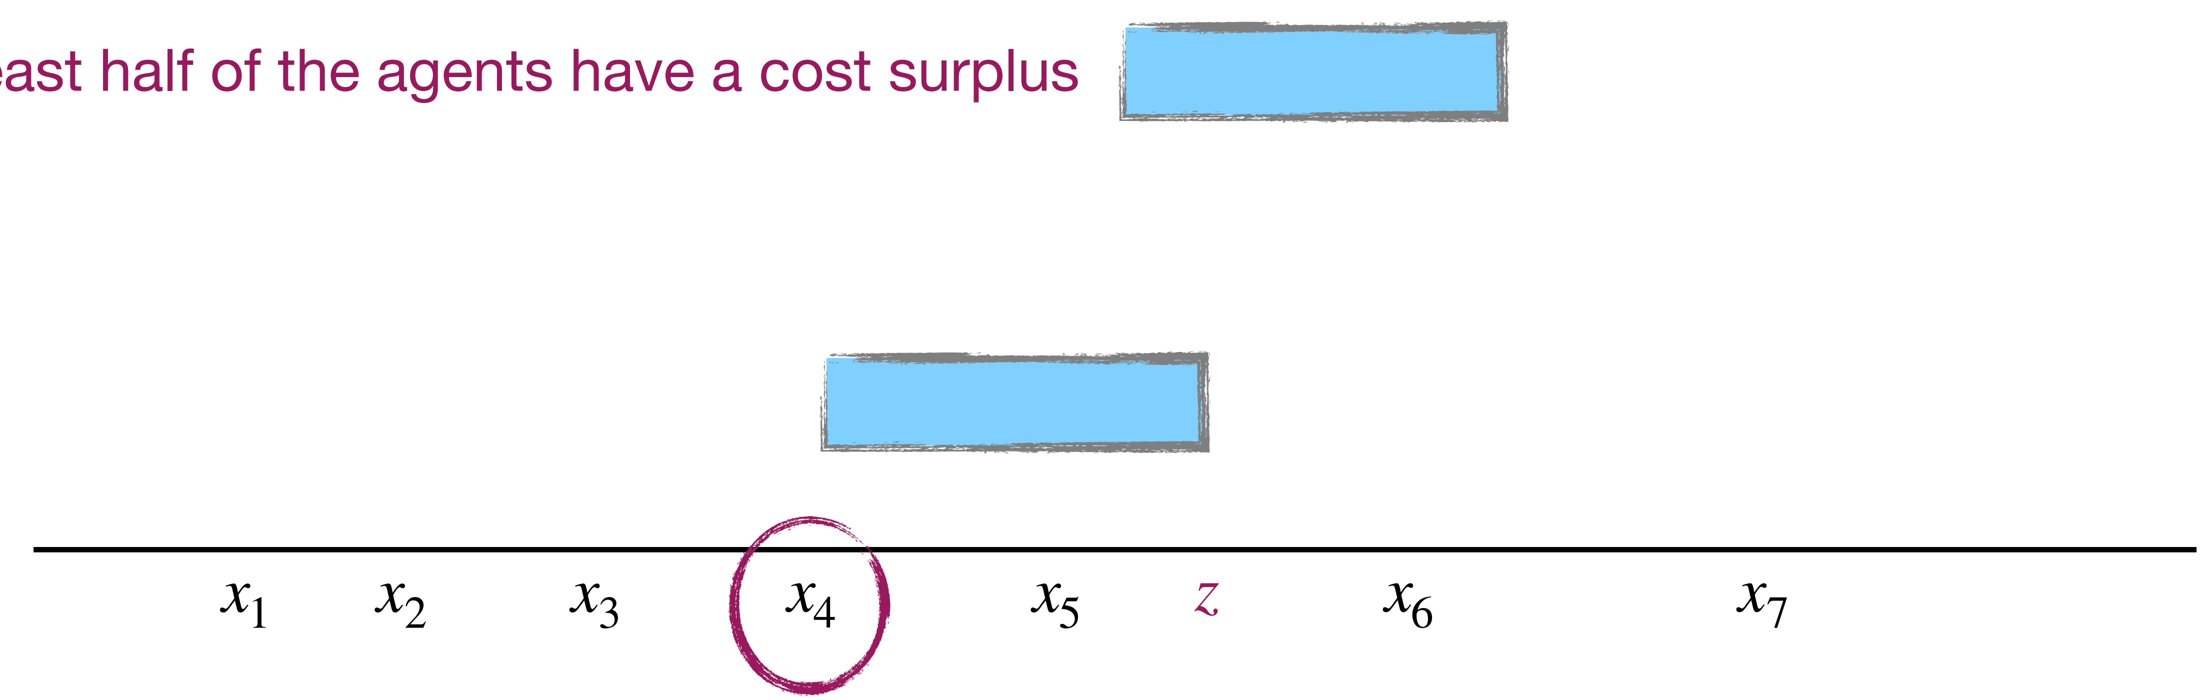

# Median Voter Mechanism

Let’s use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let’s consider any other location $z \in \mathbb{R}$

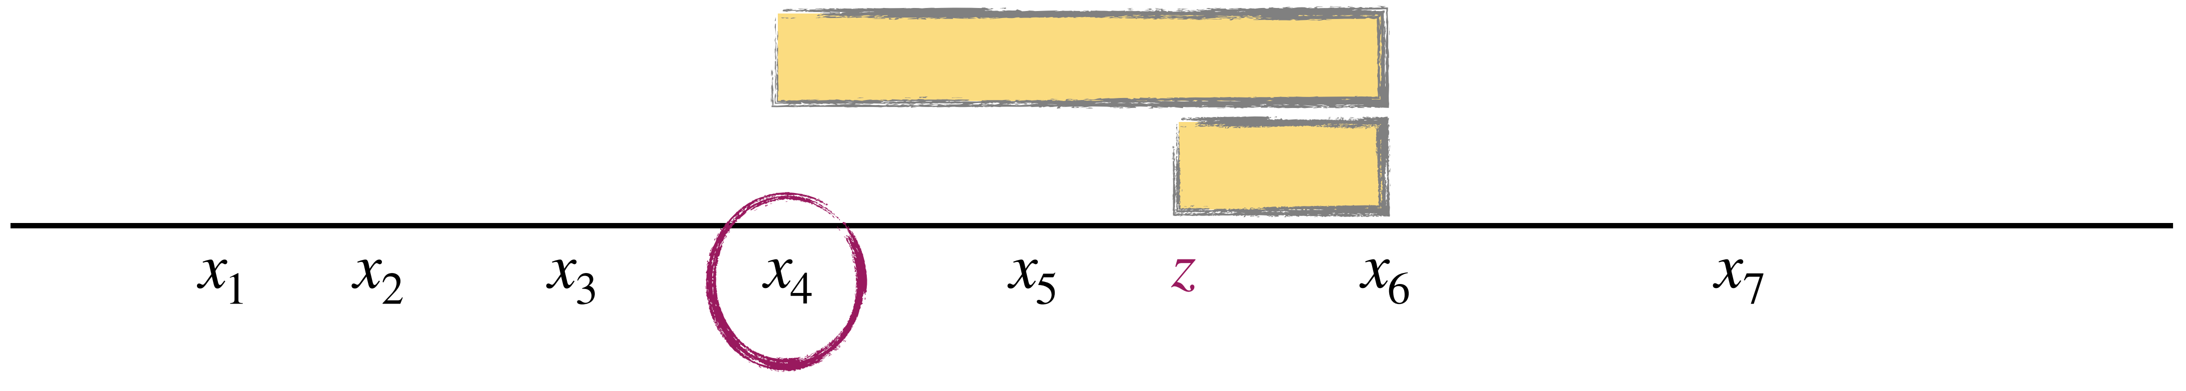

# Median Voter Mechanism

Let’s use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let’s consider any other location $z \in \mathbb{R}$

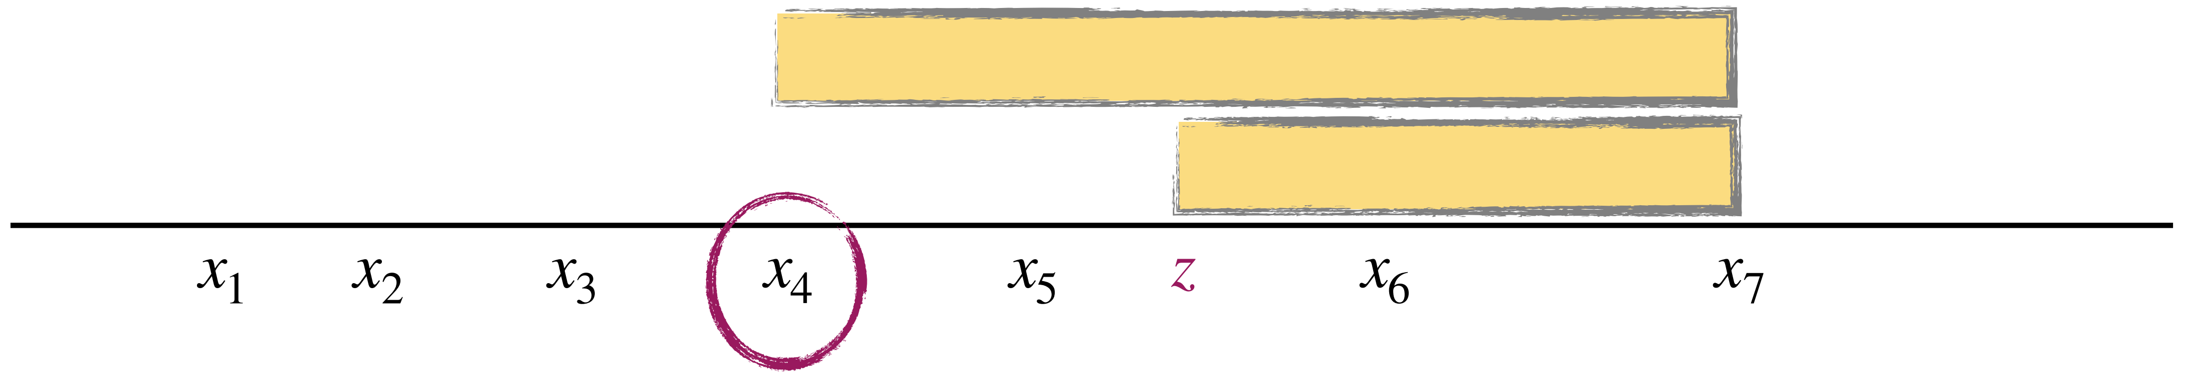

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let's consider any other location $z \in \mathbb{R}$

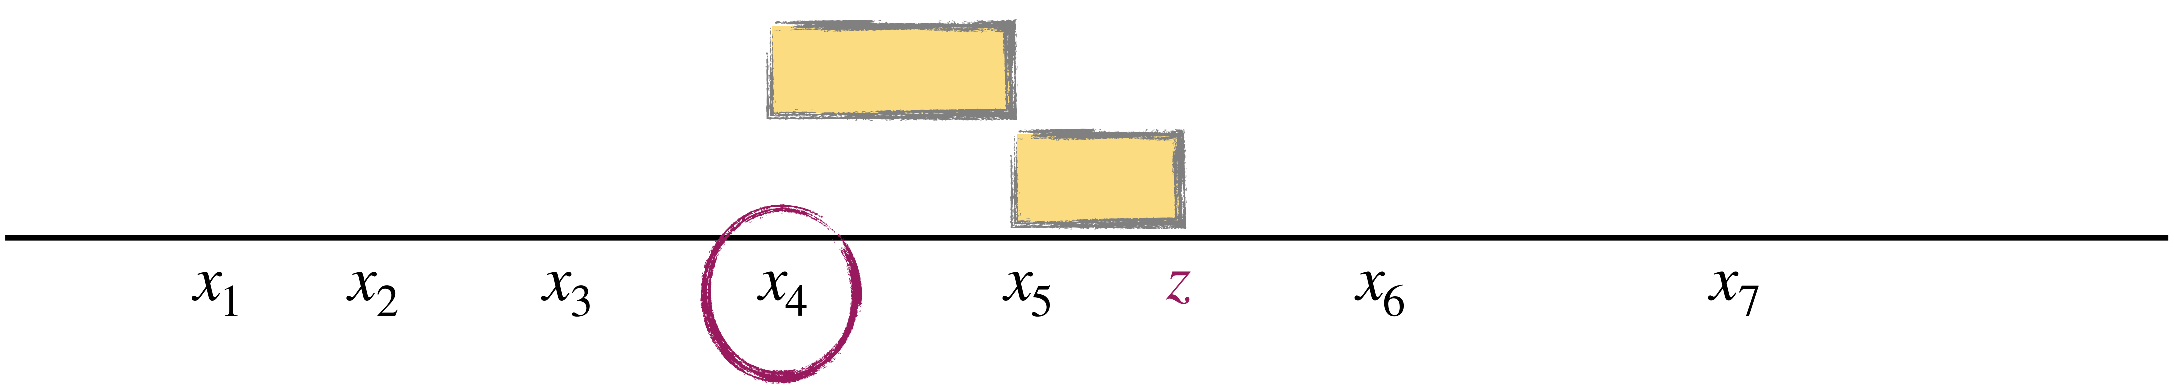

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

Let's consider any other location $z \in \mathbb{R}$

At least half of the agents have a cost surplus

At most half of the agents have a cost deficit

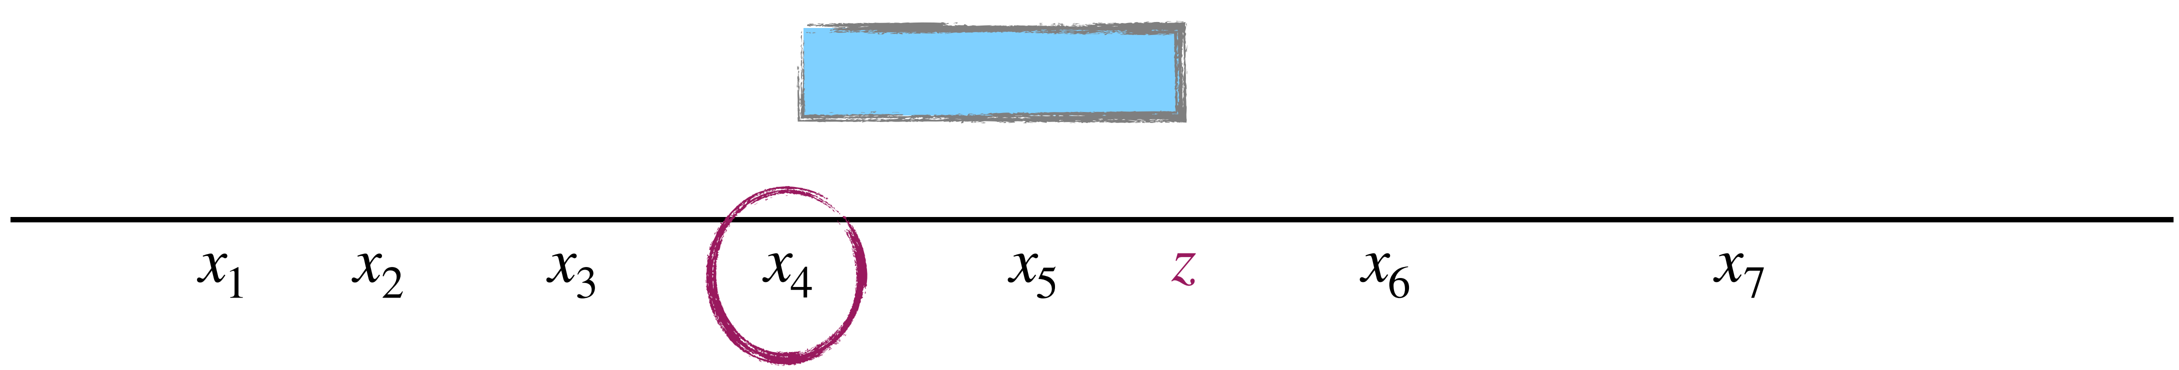


The social cost of the median is at most the social cost of any other location.

What is the approximation ratio of the median voter mechanism?

# Truthful Facility Location (Procaccia and Tennenholtz 2010)

There is a set $N = \{1, \ldots, n\}$ agents (voters), each of which has an ideal location (the "peak") $x_i$ on the real line $\mathbb{R}$.

We want to place a facility at some location $y \in \mathbb{R}$. Any location $y \in \mathbb{R}$ is a possible candidate.

Given a location $y \in \mathbb{R}$, the cost of agent $i$ is defined as $|y - x_i|$, i.e., the distance between its peak and the location $y$.

A mechanism asks the agents to report their peaks $x_i$, and outputs a location $y = f(x_1, \ldots, x_n)$.

Each agent aims to minimise its cost and reports its peak as $\hat{x}_i$ accordingly.

We want to design a truthful mechanism (voting rule) for the problem that has the minimum possible approximation ratio for the social cost objective, i.e., the sum of agents' costs $\sum_{i \in N} |y - x_i|$

# Truthful Facility Location (Procaccia and Tennenholtz 2010)

There is a set $N = \{1, \ldots, n\}$ agents (voters), each of which has an ideal location (the "peak") $x_i$ on the real line $\mathbb{R}$.

We want to place a facility at some location $y \in \mathbb{R}$. Any location $y \in \mathbb{R}$ is a possible candidate.

Given a location $y \in \mathbb{R}$, the cost of agent $i$ is defined as $|y - x_i|$, i.e., the distance between its peak and the location $y$.

A mechanism asks the agents to report their peaks $x_i$, and outputs a location $y = f(x_1, \ldots, x_n)$.

Each agent aims to minimise its cost and reports its peak as $\hat{x}_i$ accordingly.

We want to design a truthful mechanism (voting rule) for the problem that has the minimum possible approximation ratio for the maximum cost objective, i.e., the maximum of agents' costs $\max_{i \in N} |y - x_i|$

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

What is the optimal location (the one that minimises the maximum cost)?

How far can the median be from the optimal location?

It looks fairly close in the example!

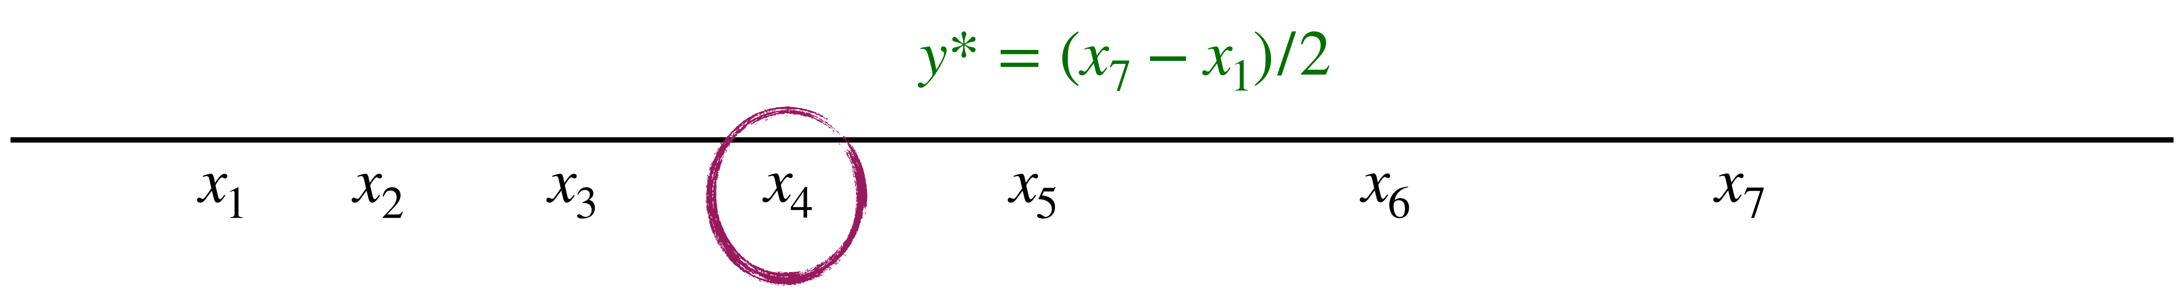

# Median Voter Mechanism

Let's use the median voter rule for the TFL problem.

The mechanism is truthful for the same reason as before.

What is the optimal location (the one that minimises the maximum cost)?

How far can the median be from the optimal location?

It looks fairly close in the example!

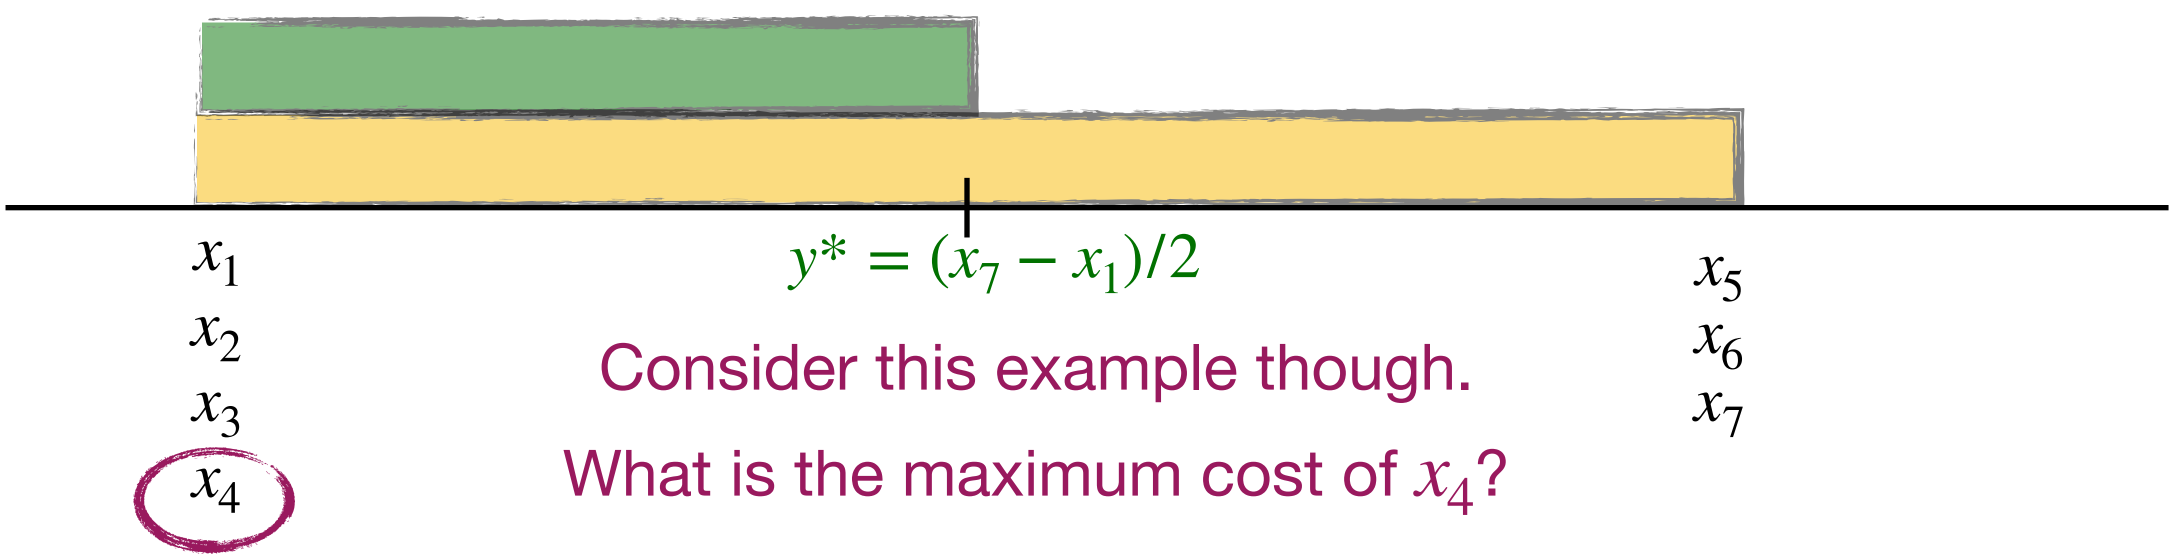


Consider this example though.

What is the maximum cost of $x_4$?

# Median Voter Mechanism

The approximation ratio of the MVM is at least 2.

The approximation ratio of the MVM is at most 2.

The approximation ratio of any $k$-th order statistic is exactly 2.

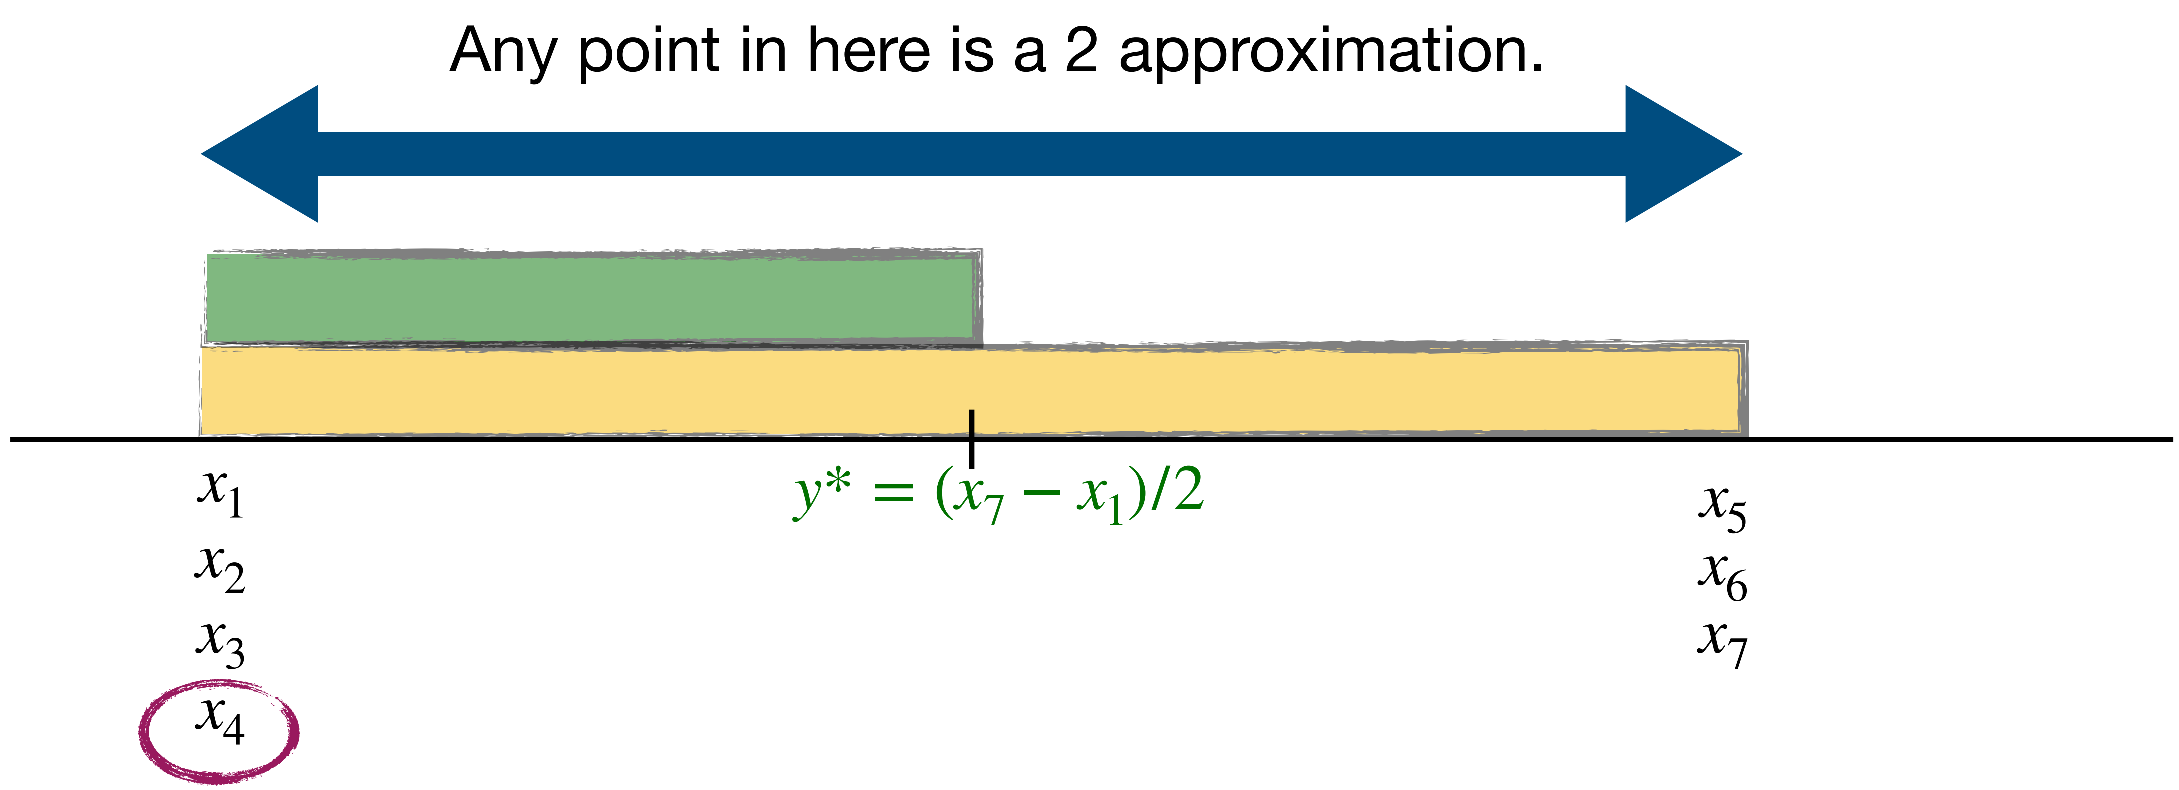

# Truthful Facility Location (Procaccia and Tennenholtz 2010)

There is a set $N = \{1, \ldots, n\}$ agents (voters), each of which has an ideal location (the "peak") $x_i$ on the real line $\mathbb{R}$.

We want to place a facility at some location $y \in \mathbb{R}$. Any location $y \in \mathbb{R}$ is a possible candidate.

Given a location $y \in \mathbb{R}$, the cost of agent $i$ is defined as $|y - x_i|$, i.e., the distance between its peak and the location $y$.

A mechanism asks the agents to report their peaks $x_i$, and outputs a location $y = f(x_1, \ldots, x_n)$.

Each agent aims to minimise its cost and reports its peak as $\hat{x}_i$ accordingly.

We want to design a truthful mechanism (voting rule) for the problem that has the minimum possible approximation ratio for the maximum cost objective, i.e., the maximum of agents' costs $\max_{i \in N} |y - x_i|$

# A lower bound for all truthful mechanisms

Consider the instance shown below.

Assume by contradiction that there exists some truthful mechanism $M$ with approximation ratio $< 2$.

$\Rightarrow$ the facility needs to be placed in the interior of the interval, wlog, closer to the right endpoint.

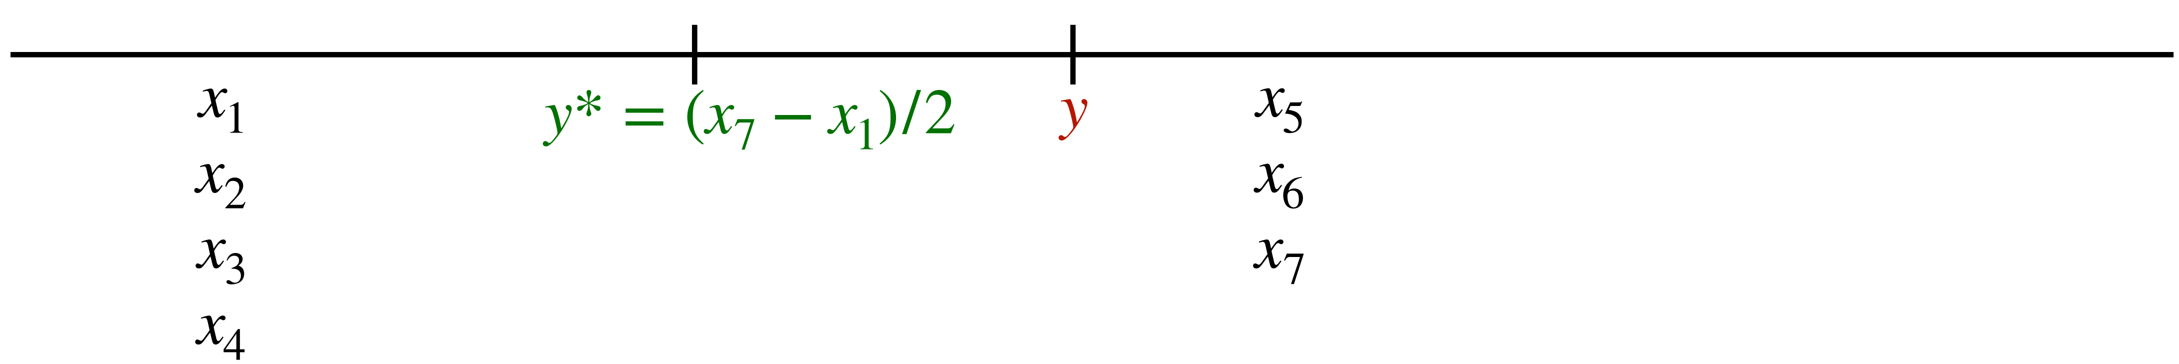

# A lower bound for all truthful mechanisms

Assume now that agent with peak $x_5$ reports $x_5' = y$

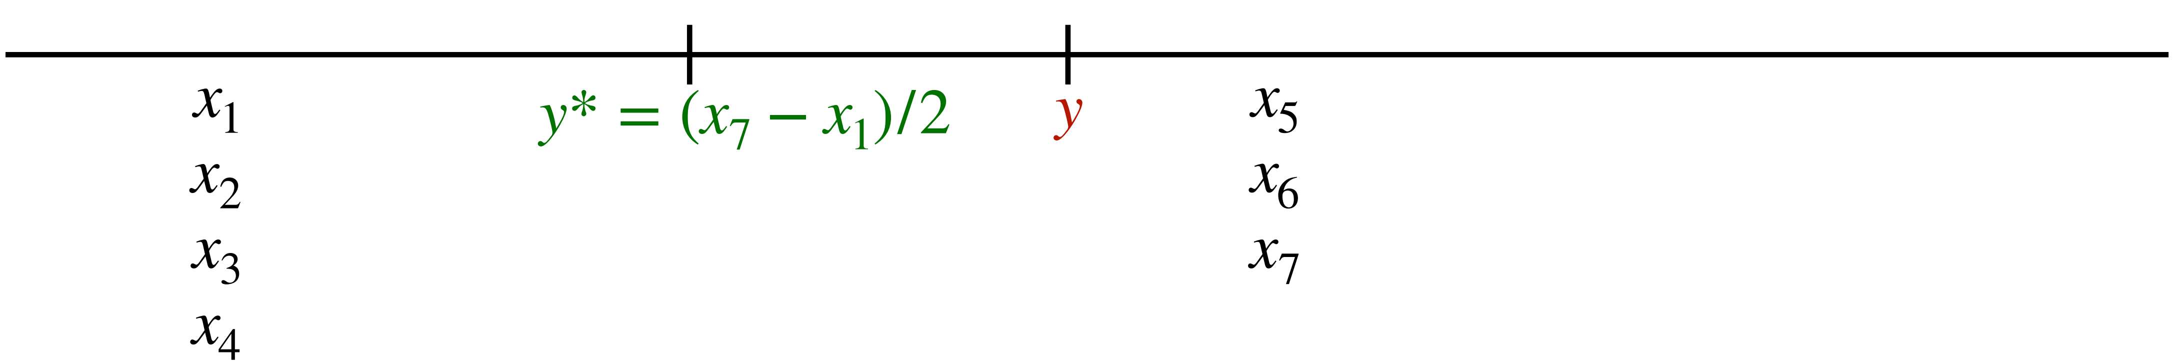

# A lower bound for all truthful mechanisms

Assume now that agent with peak $x_5$ reports $x_5' = y$

The facility cannot change position, let's see why.

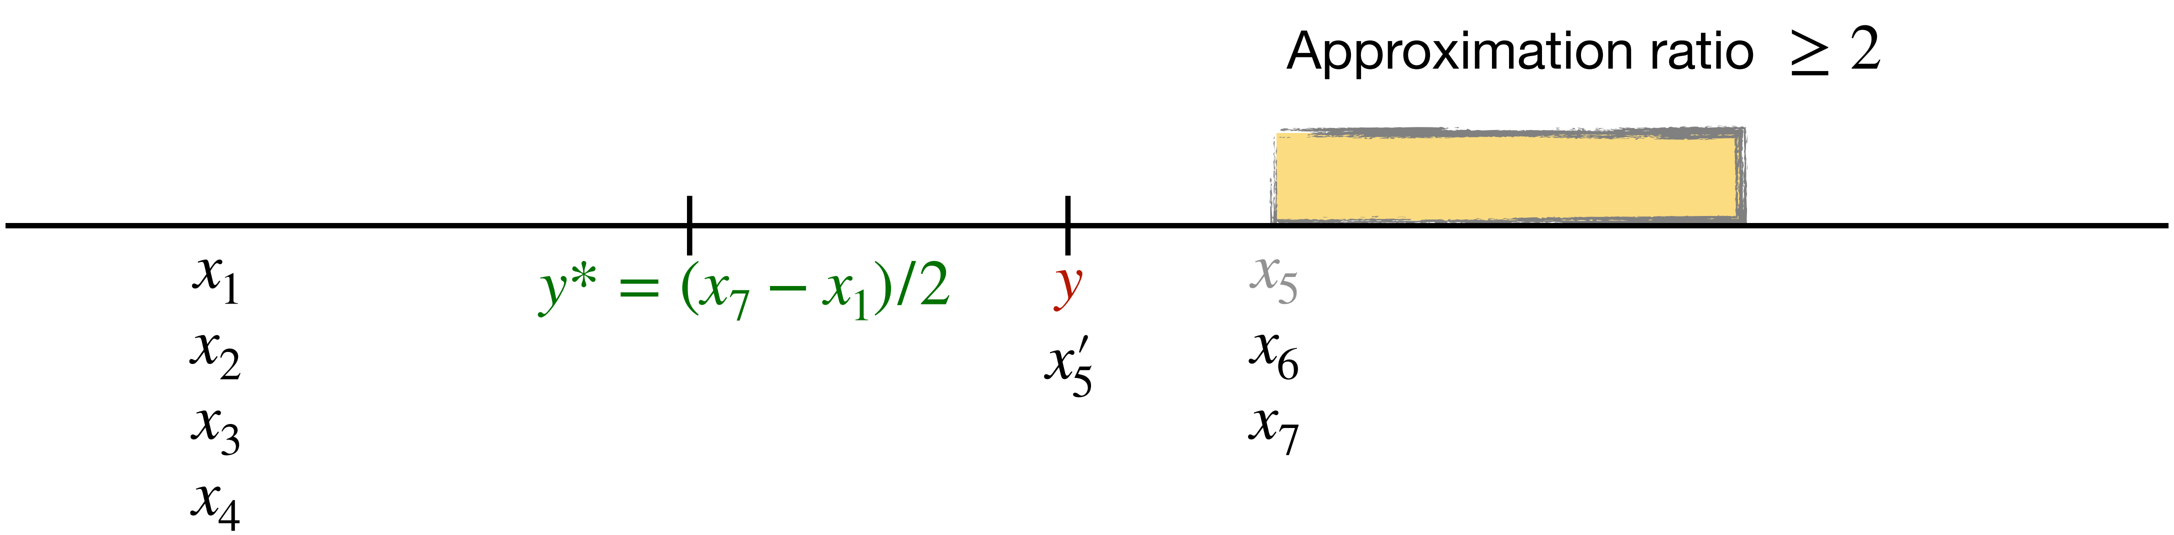

# A lower bound for all truthful mechanisms

Assume now that agent with peak $x_5$ reports $x_5' = y$

The facility cannot change position, let's see why.

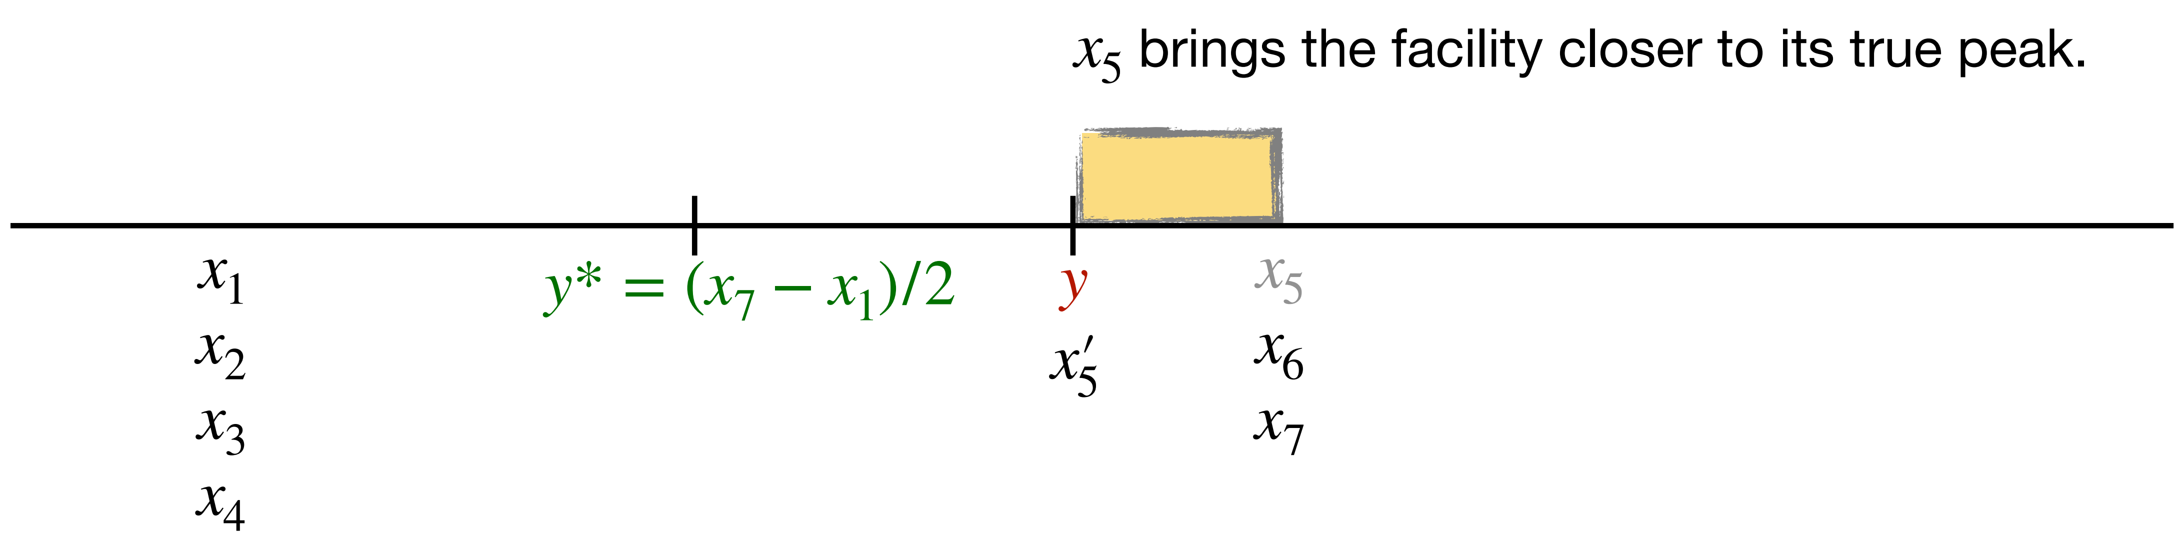

# A lower bound for all truthful mechanisms

Assume now that agent with peak $x_5$ reports $x_5' = y$

The facility cannot change position, let's see why.

It could be the case that $x_5'$ is the true peak and $x_5$ is the misreport.

In that case the misreport would bring the facility exactly on the true peak.

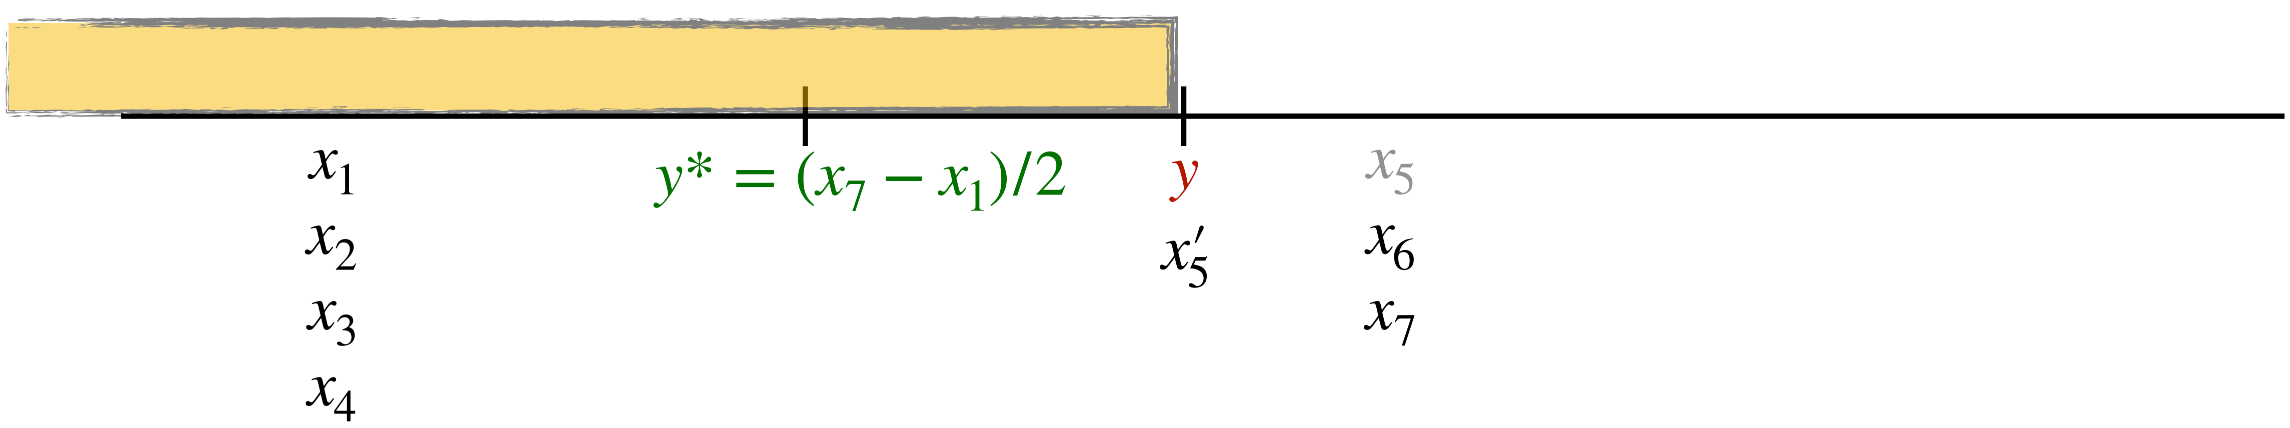

# A lower bound for all truthful mechanisms

Assume now that agent with peak $x_5$ reports $x'_5 = y$

The facility cannot change position, let's see why.

We can use the same argument for $x_6$ and $x_7$.

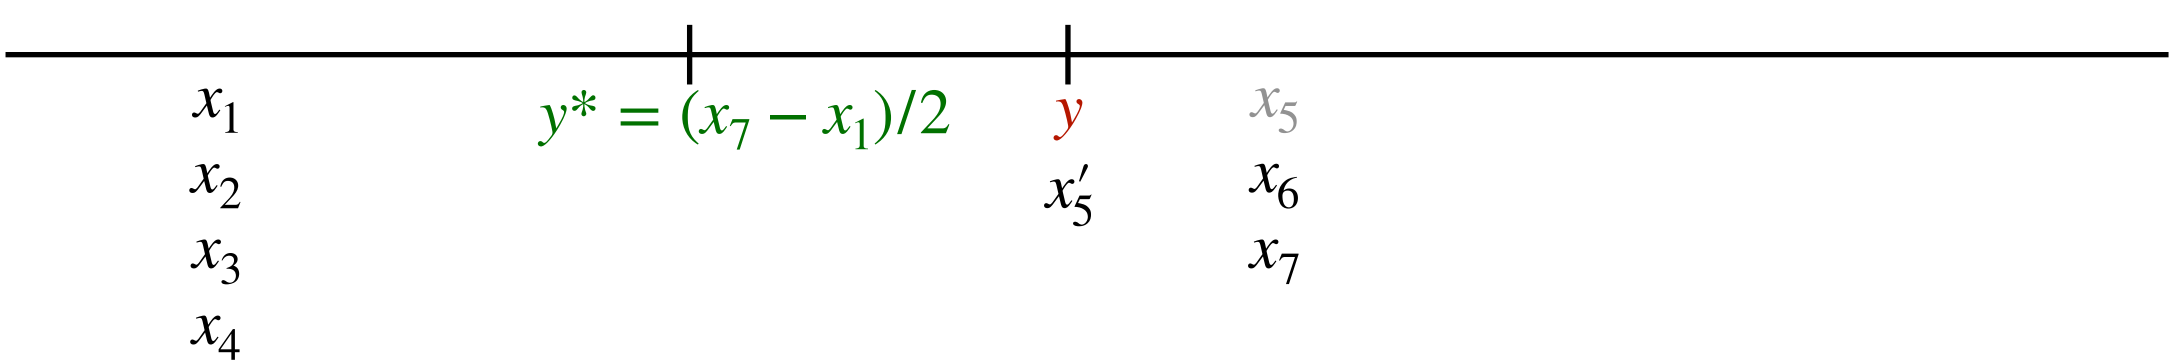

# A lower bound for all truthful mechanisms

Assume now that agent with peak $x_5$ reports $x_5' = y$

The facility cannot change position, let's see why.

We can use the same argument for $x_6$ and $x_7$.

What is the ratio on this instance?

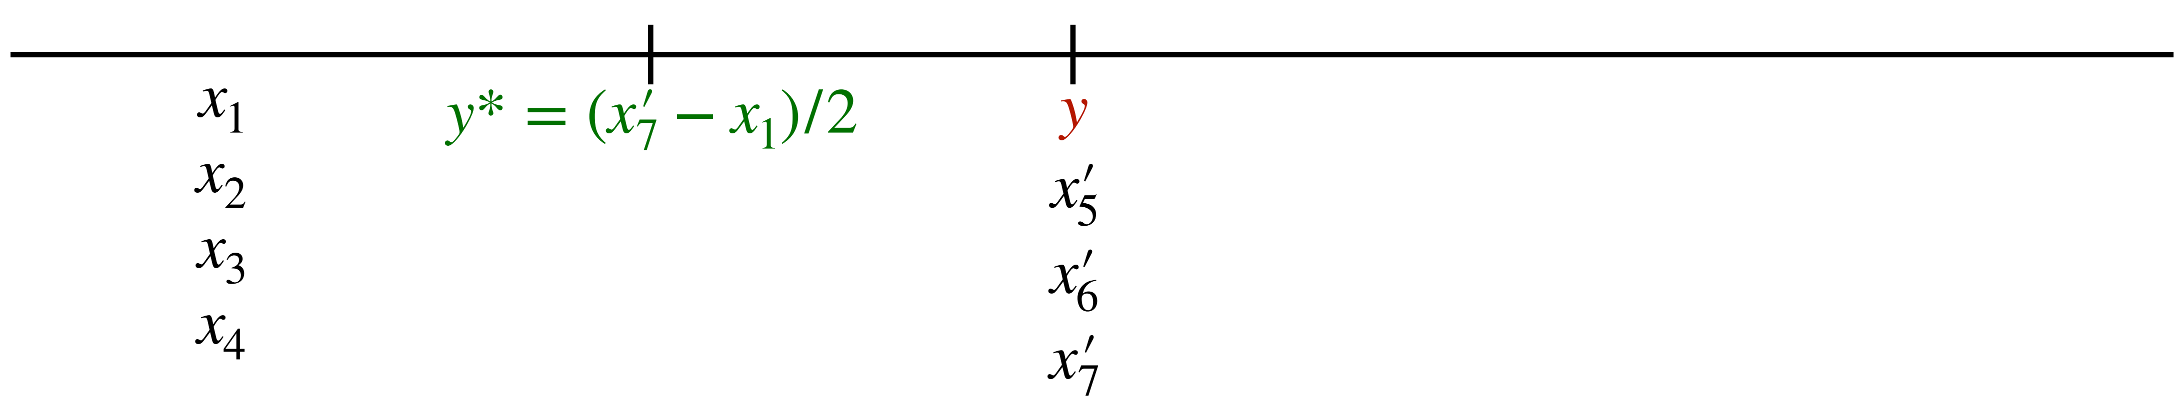

# Truthful Facility Location, max cost objective

Theorem (Procaccia and Tennenholtz 2010): The best possible approximation ratio achieved by any truthful mechanism for the maximum cost objective is 2. This is achieved by any $k$-th ordered statistic mechanism.