Algorithmic Game Theory and Applications

Introduction to Mechanism Design: Social Choice Theory

Analytic vs Proactive

Analytic vs Proactive

Analytic Approach (Game Theory): So far we have mainly been analysing games, e.g., to find whether they have equilibria, to find algorithms for computing equilibria, to quantify the inefficiency of the equilibria etc.

Analytic vs Proactive

Analytic Approach (Game Theory): So far we have mainly been analysing games, e.g., to find whether they have equilibria, to find algorithms for computing equilibria, to quantify the inefficiency of the equilibria etc.

Proactive Approach (Mechanism Design): We can design the rules of the game, in a way that induces good properties, e.g., good equilibria.

The general mechanism design setup

The general mechanism design setup

At a very high level:

The general mechanism design setup

At a very high level:

We have a problem that we would like to solve, given as input information (or "preferences" or "values" or "utilities") controlled by self-interested agents (players).

The general mechanism design setup

At a very high level:

We have a problem that we would like to solve, given as input information (or “preferences” or “values” or “utilities”) controlled by self-interested agents (players).

A mechanism is a function that maps this input to an output.

The general mechanism design setup

At a very high level:

We have a problem that we would like to solve, given as input information (or “preferences” or “values” or “utilities”) controlled by self-interested agents (players).

A mechanism is a function that maps this input to an output.

The “mechanism” is like an “algorithm”, but it takes into account the fact that the agents will act strategically.

The general mechanism design setup

At a very high level:

We have a problem that we would like to solve, given as input information (or “preferences” or “values” or “utilities”) controlled by self-interested agents (players).

A mechanism is a function that maps this input to an output.

The “mechanism” is like an “algorithm”, but it takes into account the fact that the agents will act strategically.

The goal is to design mechanisms that *align the incentives* of the agents with those of society (or the designer).

The general mechanism design setup

At a very high level:

We have a problem that we would like to solve, given as input information (or “preferences” or “values” or “utilities”) controlled by self-interested agents (players).

A mechanism is a function that maps this input to an output.

The “mechanism” is like an “algorithm”, but it takes into account the fact that the agents will act strategically.

The goal is to design mechanisms that *align the incentives* of the agents with those of society (or the designer).

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

<u>Scenario:</u>

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?

Manipulating the average

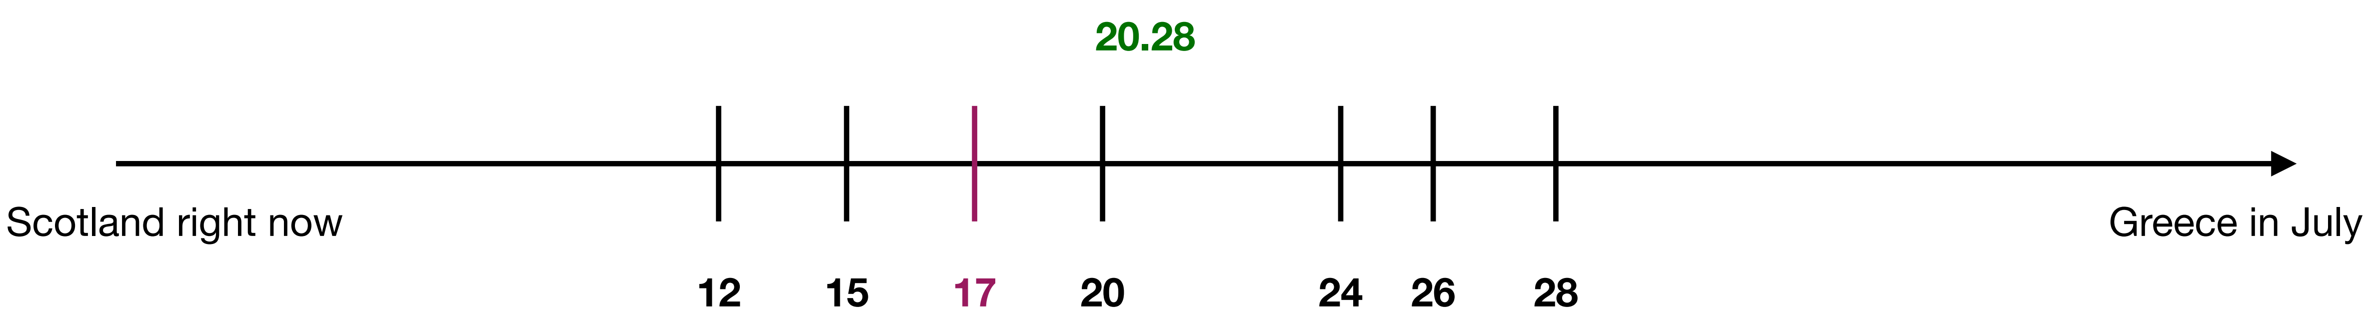

Manipulating the average

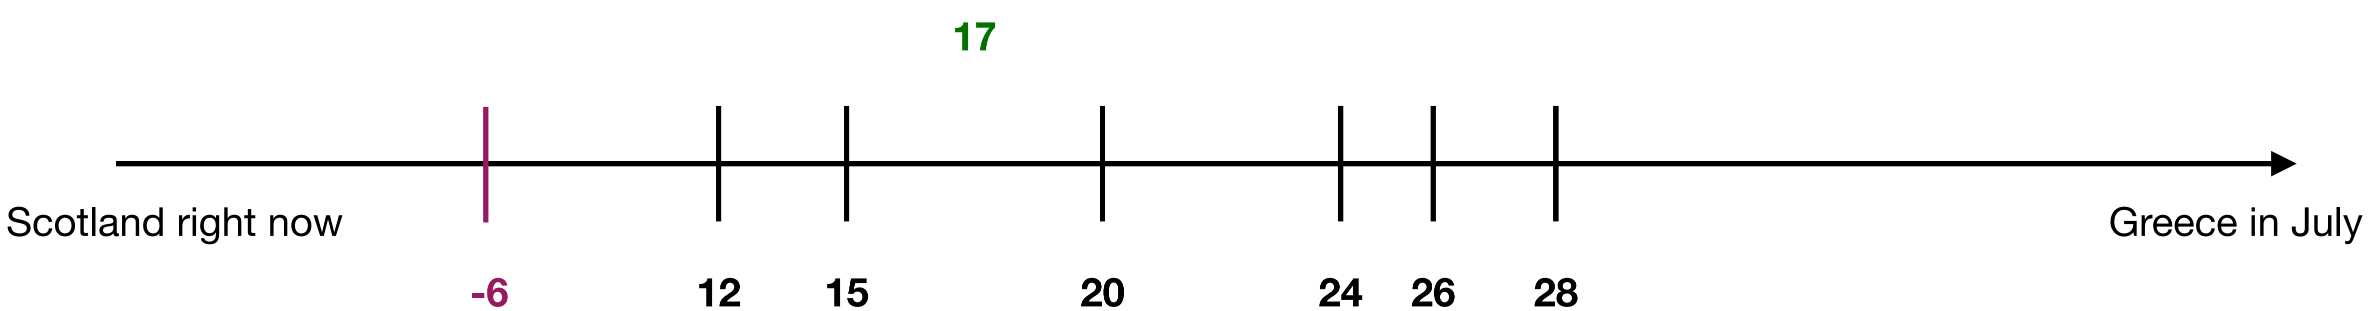

Manipulating the average

The "average choice" mechanism forces people to engage in strategic considerations.

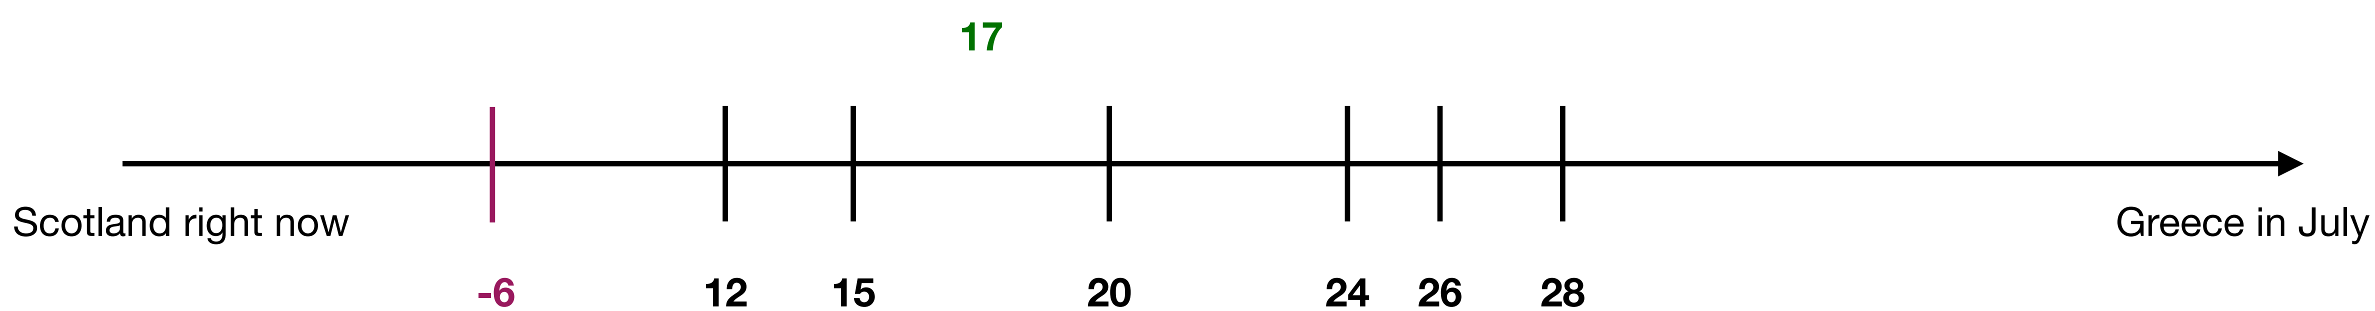

Manipulating the average

The “average choice” mechanism forces people to engage in strategic considerations.

The game that it induces might not even have pure Nash equilibria.

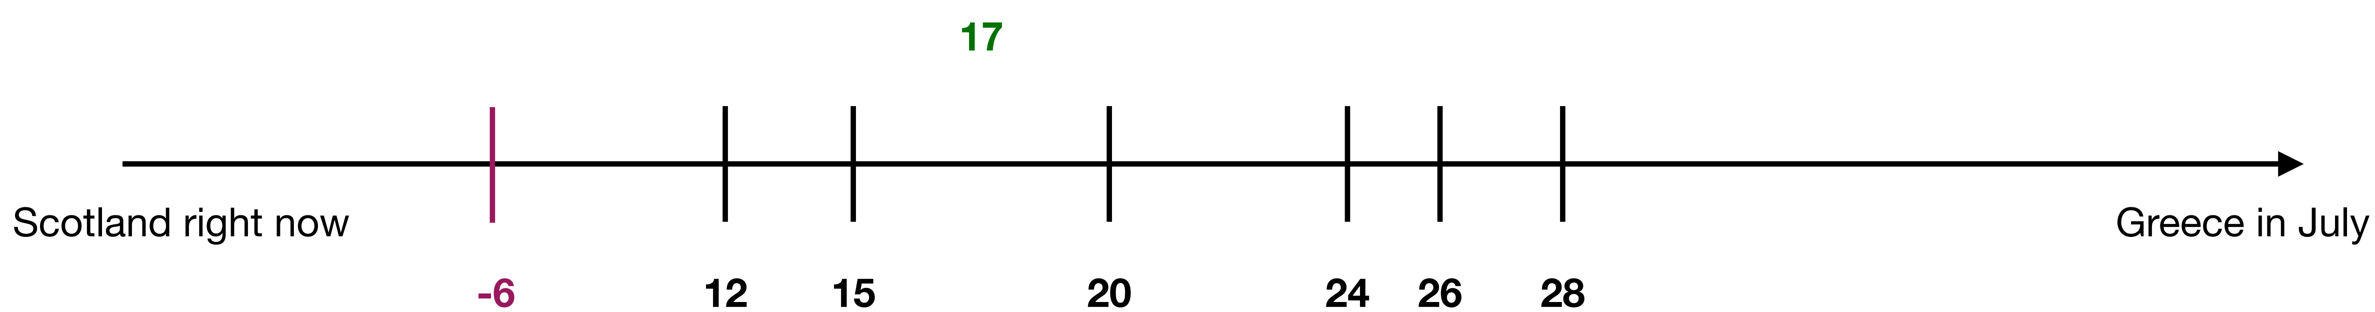

Manipulating the average

The “average choice” mechanism forces people to engage in strategic considerations.

The game that it induces might not even have pure Nash equilibria.

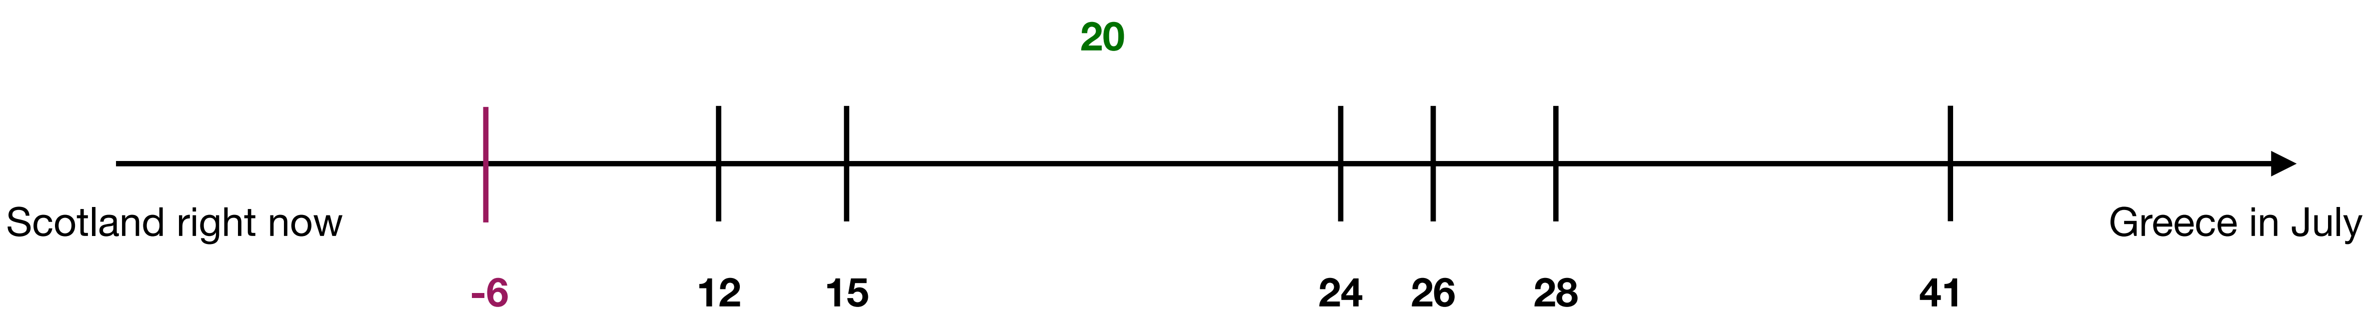

Example 2:
Setting the temperature

Scenario:

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

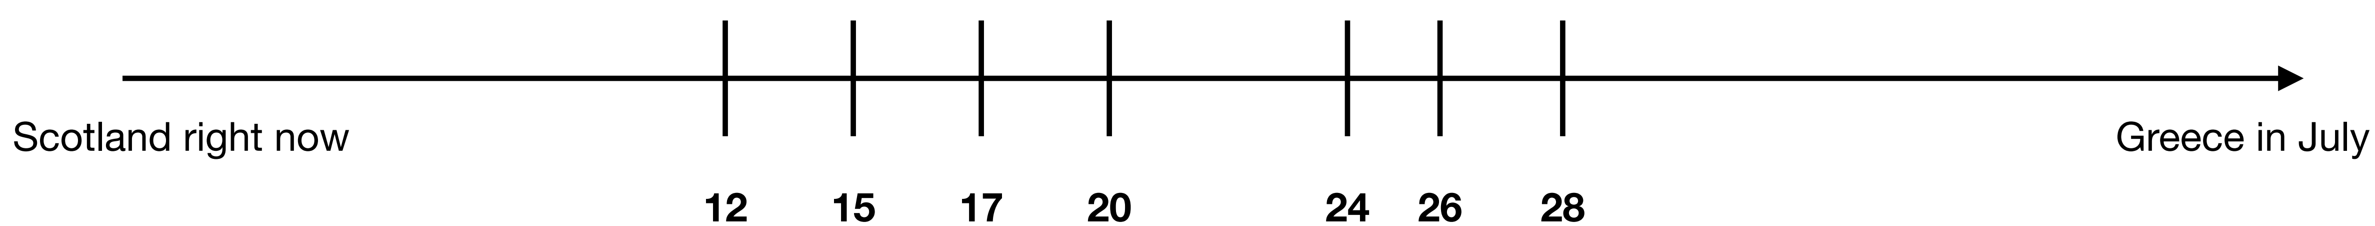

Example 2:
Setting the temperature

<u>Scenario:</u>

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

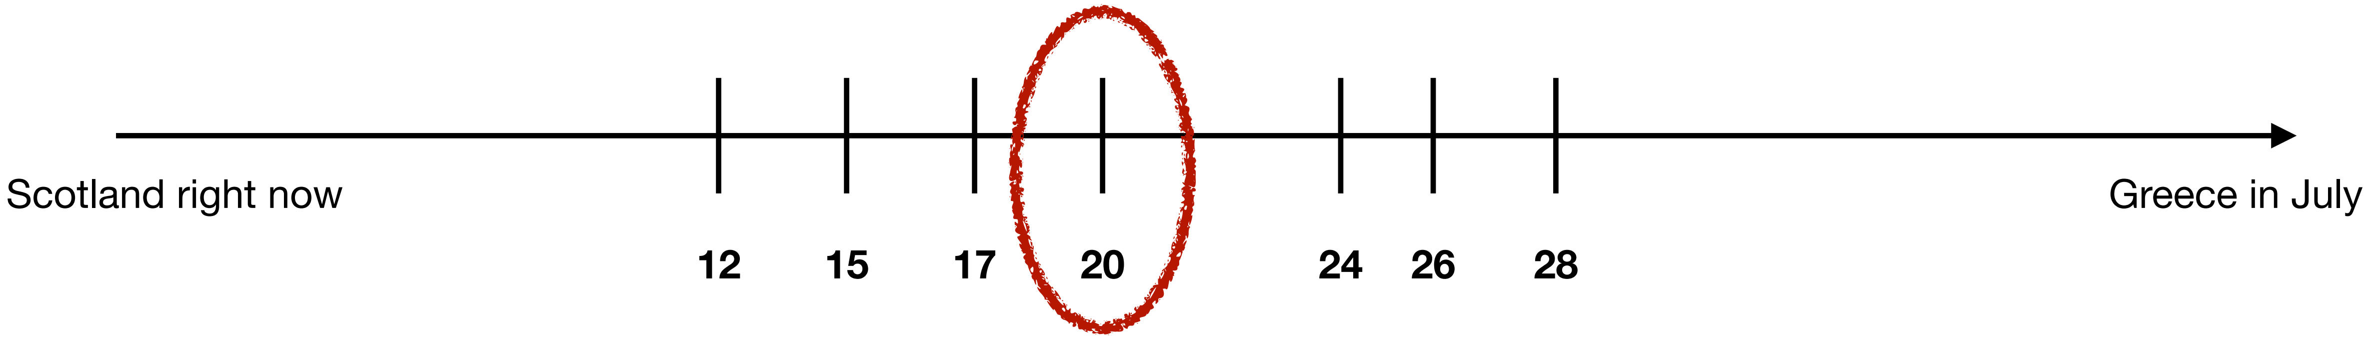

Example 2:
Setting the temperature

<u>Scenario:</u>

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

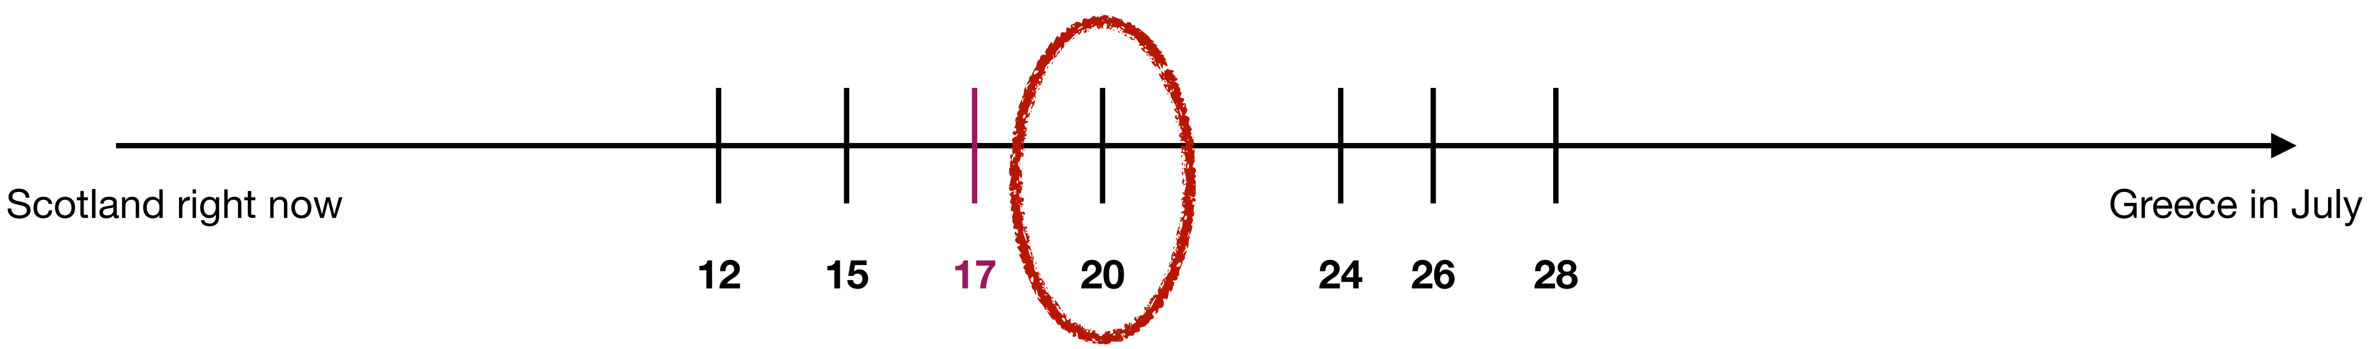

Example 2:
Setting the temperature

<u>Scenario:</u>

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

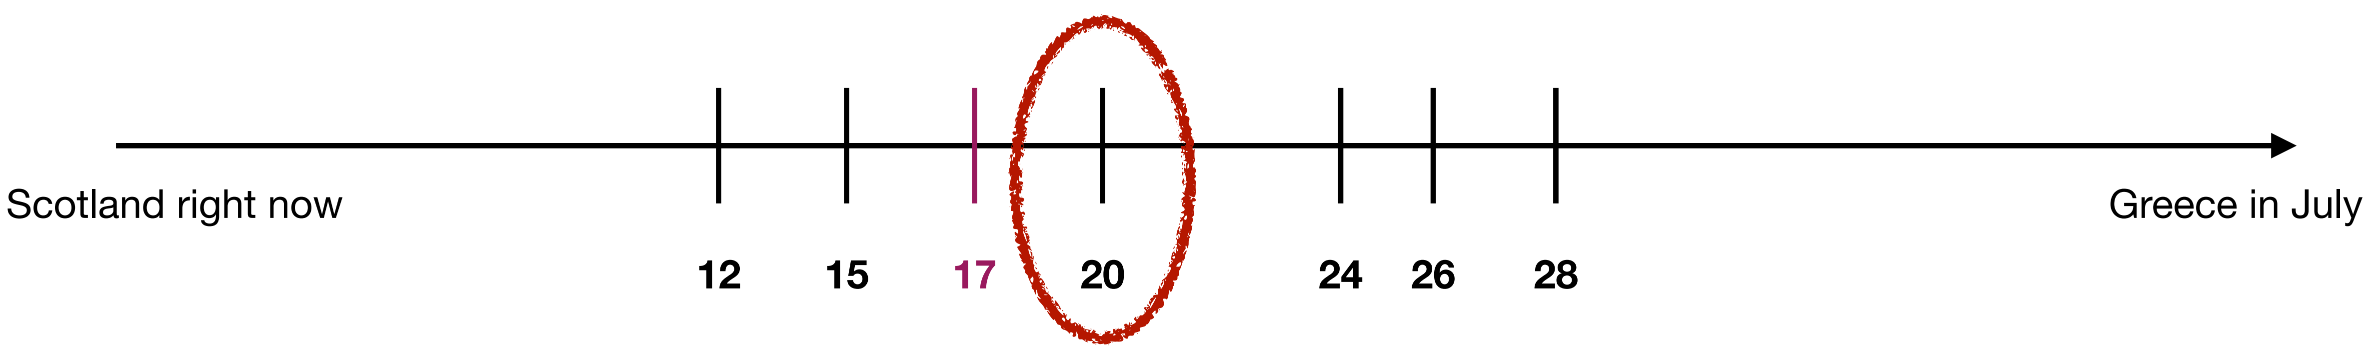

To affect the median, the agent needs to report something here

Example 2:
Setting the temperature

<u>Scenario:</u>

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

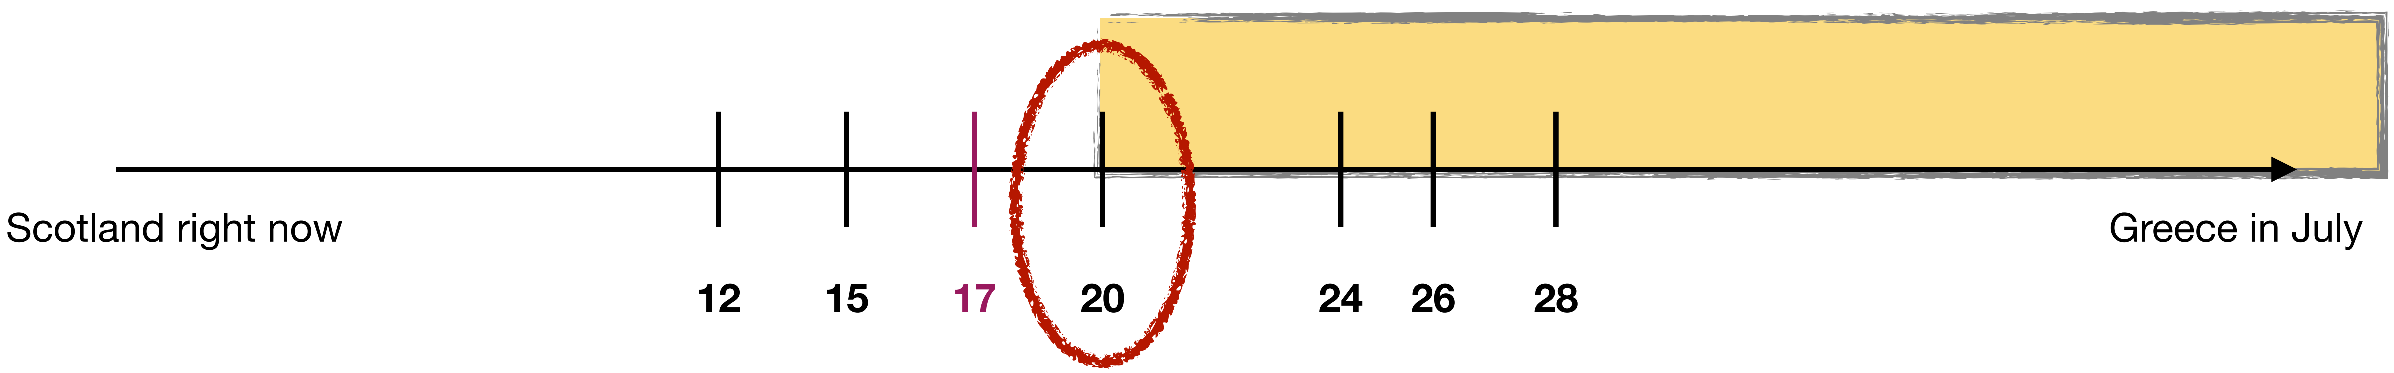

To affect the median, the agent needs to report something here

Example 2: Setting the temperature

<u>Scenario:</u>

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

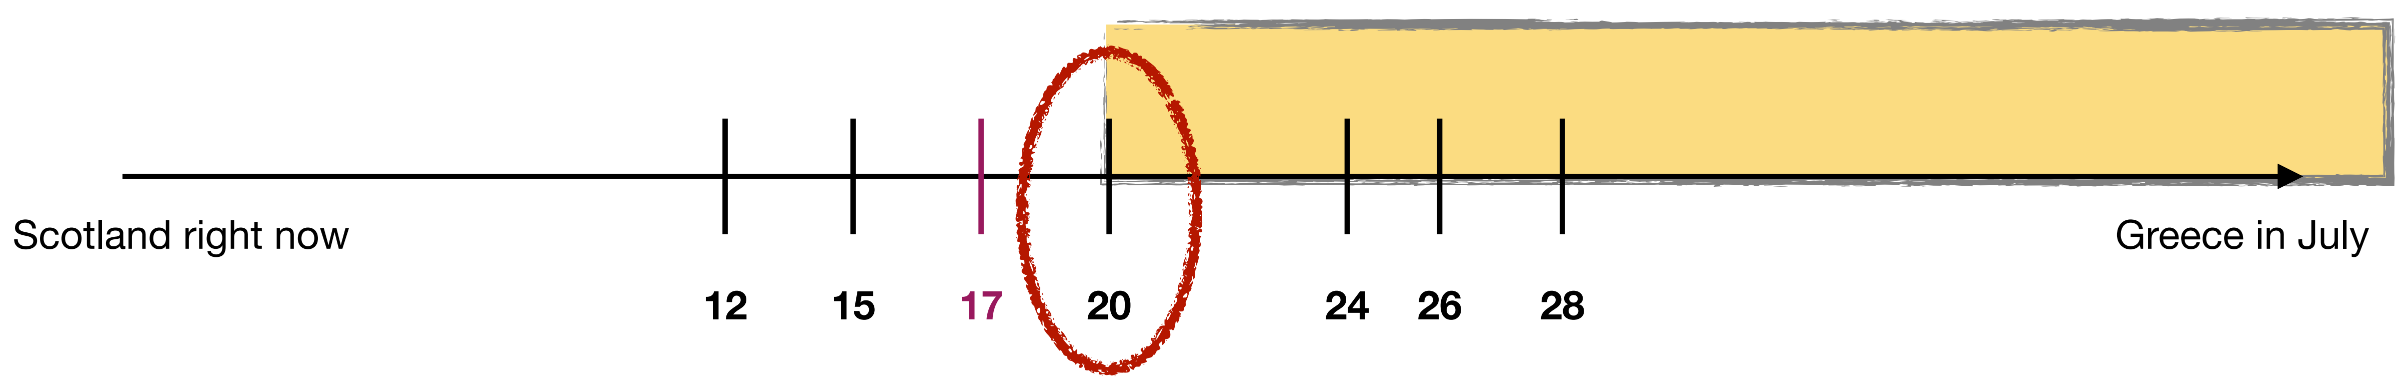

To affect the median, the agent needs to report something here

That would move the outcome farther away!

Example 2:
Setting the temperature

<u>Scenario:</u>

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

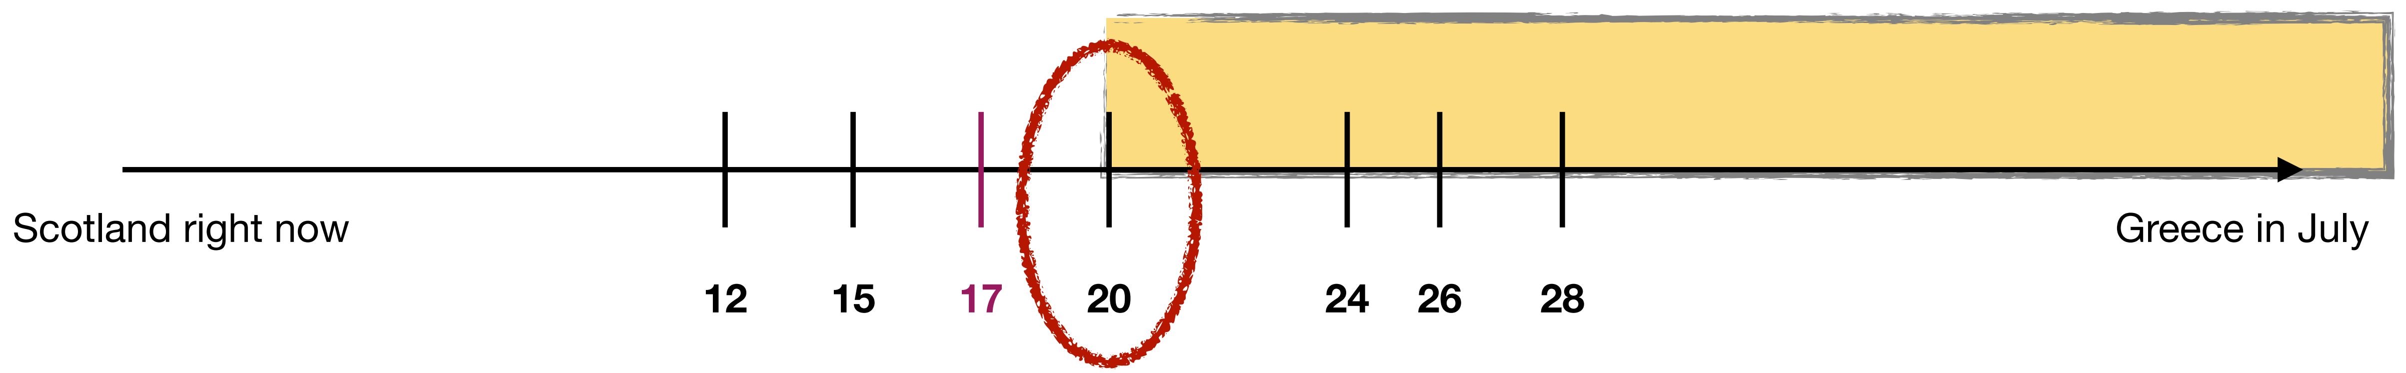

To affect the median, the agent needs to report something here

That would move the outcome farther away!

The mechanism is *truthful*.

The general social choice setting

The general social choice setting

There is a set $N = \{1, \ldots, n\}$ of agents and a set $A = \{1, \ldots, m\}$ of outcomes.

The general social choice setting

There is a set $N = \{1, \ldots, n\}$ of agents and a set $A = \{1, \ldots, m\}$ of outcomes.

Every agent $i \in N$ has a (cardinal) utility u_{ij} for each outcome $j \in M$.

The general social choice setting

There is a set $N = \{1, \ldots, n\}$ of agents and a set $A = \{1, \ldots, m\}$ of outcomes.

Every agent $i \in N$ has a (cardinal) utility u_{ij} for each outcome $j \in M$.

Let's assume for now that these are all distinct.

The general social choice setting

There is a set $N = \{1, \ldots, n\}$ of agents and a set $A = \{1, \ldots, m\}$ of outcomes.

Every agent $i \in N$ has a (cardinal) utility u_{ij} for each outcome $j \in M$.

Let's assume for now that these are all distinct.

Sometimes it will make sense to only consider the *preference rankings* that these values induce.

The general social choice setting

There is a set $N = \{1, \ldots, n\}$ of agents and a set $A = \{1, \ldots, m\}$ of outcomes.

Every agent $i \in N$ has a (cardinal) utility u_{ij} for each outcome $j \in M$.

Let's assume for now that these are all distinct.

Sometimes it will make sense to only consider the *preference rankings* that these values induce.

We will denote $\succ_i$ the (ordinal) preference ranking of agent $i \in N$, where $\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$

The general social choice setting

There is a set $N = \{1, \ldots, n\}$ of agents and a set $A = \{1, \ldots, m\}$ of outcomes.

Every agent $i \in N$ has a (cardinal) utility u_{ij} for each outcome $j \in M$.

Let's assume for now that these are all distinct.

Sometimes it will make sense to only consider the *preference rankings* that these values induce.

We will denote $\succ_i$ the (ordinal) preference ranking of agent $i \in N$, where $\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$

Let $\succ = (\succ_1, \ldots, \succ_n)$ denote a preference ranking profile.

Convention: Use voting terminology

There is a set $N = \{1, \ldots, n\}$ of voters and a set $A = \{1, \ldots, m\}$ of candidates or alternatives.

Every voter $i \in N$ has a (cardinal) utility u_{ij} for each candidate $j \in M$.

Let's assume for now that these are all distinct.

Sometimes it will make sense to only consider the *preference rankings* that these values induce.

We will denote $\succ_i$ the (ordinal) preference ranking of voter $i \in N$, where $\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$

Let $\succ = (\succ_1, \ldots, \succ_n)$ denote a preference ranking profile.

Mechanisms

Mechanisms

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

Mechanisms

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

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

Mechanisms

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.

Mechanisms

A *social choice function*, or *voting rule*, or *mechanism* is a function $f : (\succ)^n \to A$ 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 $\{1, \ldots, m\}$.

Truthfulness

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ_i', \succ_{-i})$$

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ_i', \succ_{-i})$$

In a perhaps more familiar form:

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ_i', \succ_{-i})$$

In a perhaps more familiar form:

For any preference rankings $\succ_{-i}$ of the other voters

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ'_i$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ'_i, \succ_{-i})$$

In a perhaps more familiar form:

For any preference rankings $\succ_{-i}$ of the other voters

Let $\alpha = f(\succ_i, \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is telling the truth.

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ_i', \succ_{-i})$$

In a perhaps more familiar form:

For any preference rankings $\succ_{-i}$ of the other voters

Let $\alpha = f(\succ_i, \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is telling the truth.

Let $\beta = f(\succ_i', \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is lying by reporting $\succ_i'$.

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ_i', \succ_{-i})$$

In a perhaps more familiar form:

For any preference rankings $\succ_{-i}$ of the other voters

Let $\alpha = f(\succ_i, \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is telling the truth.

Let $\beta = f(\succ_i', \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is lying by reporting $\succ_i'$.

The utility of $v_{i\alpha}$ of the voter for α is larger than the utility $v_{i\beta}$ of the voter for β.

The social choice setting as a game

The social choice setting as a game

Who are the players?

The social choice setting as a game

Who are the players?

The voters.

The social choice setting as a game

Who are the players?

The voters.

What are the pure strategies?

The social choice setting as a game

Who are the players?

The voters.

What are the pure strategies?

The preference rankings $\succ_i$.

The social choice setting as a game

Who are the players?

The voters.

What are the pure strategies?

The preference rankings $\succ_i$.

What are the utilities?

The social choice setting as a game

Who are the players?

The voters.

What are the pure strategies?

The preference rankings $\succ_i$*.*

What are the utilities?

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

The social choice setting as a game

Who are the players?

The voters.

What are the pure strategies?

The preference rankings $\succ_i$.

What are the utilities?

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

What does truthfulness mean?

Truthfulness

A voting rule is *truthful*, or *strategyproof*, or *incentive-compatible*, if for any voter $i \in N$, any *true* preference ranking $\succ_i$ of the voter, any *reported* preference ranking $\succ_i'$ of the voter, and any reported preference rankings $\succ_{-i}$ of the remaining voters, it holds that:

$$f(\succ_i, \succ_{-i}) \succ_i f(\succ_i', \succ_{-i})$$

In a perhaps more familiar form:

For any preference rankings $\succ_{-i}$ of the other voters

Let $\alpha = f(\succ_i, \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is telling the truth.

Let $\beta = f(\succ_i', \succ_{-i})$ be the candidate that gets selected ("elected") when voter i is lying by reporting $\succ_i'$.

The utility of $v_{i\alpha}$ of the voter for α is larger than the utility $v_{i\beta}$ of the voter for β.

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

This assumes that the utilities are distinct. What is we have ties?

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

This assumes that the utilities are distinct. What is we have ties?

This is also contingent on our voting rules being deterministic. What if we have randomisation?

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

This assumes that the utilities are distinct. What is we have ties?

This is also contingent on our voting rules being deterministic. What if we have randomisation?

e.g. f elects α with probability $1/2$ and β with probability $1/2$.

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

This assumes that the utilities are distinct. What is we have ties?

This is also contingent on our voting rules being deterministic. What if we have randomisation?

e.g. f elects α with probability $1/2$ and β with probability $1/2$.

We need to be able to compare γ with e.g. $(1/2) \cdot a + (1/2) \cdot \beta$.

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

This assumes that the utilities are distinct. What is we have ties?

This is also contingent on our voting rules being deterministic. What if we have randomisation?

e.g. f elects α with probability $1/2$ and β with probability $1/2$.

We need to be able to compare γ with e.g. $(1/2) \cdot a + (1/2) \cdot \beta$.

This also assumes that our voting rule was defined to only depend on the *ordinal* preference rankings.

An subtle but important note

These are the utilities v_{ij}, although these are equivalently implicitly defined by the preference rankings $\succ_i$.

$$\alpha \succ_i \beta \Leftrightarrow v_{i\alpha} > v_{i\beta}$$

This assumes that the utilities are distinct. What is we have ties?

This is also contingent on our voting rules being deterministic. What if we have randomisation?

e.g. f elects α with probability $1/2$ and β with probability $1/2$.

We need to be able to compare γ with e.g. $(1/2) \cdot a + (1/2) \cdot \beta$.

This also assumes that our voting rule was defined to only depend on the *ordinal* preference rankings.

Cardinal vs Ordinal (Randomised) Rules

An **ordinal** *social choice function*, or *voting rule*, or *mechanism* is a function $f : (\succ)^n \rightarrow \Delta(A)$ mapping (reported) preference profiles to (probability distributions over) candidates,

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

Cardinal vs Ordinal (Randomised) Rules

An **ordinal** *social choice function*, or *voting rule*, or *mechanism* is a function $f : (\succ)^n \rightarrow \Delta(A)$ mapping (reported) preference profiles to (probability distributions over) candidates,

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

A **cardinal** *social choice function*, or *voting rule*, or *mechanism* is a function $f : U^n \rightarrow \Delta(A)$ mapping (reported) utility profiles to (probability distributions over) candidates,

Cardinal vs Ordinal (Randomised) Rules

An **ordinal** *social choice function*, or *voting rule*, or *mechanism* is a function $f : (\succ)^n \rightarrow \Delta(A)$ mapping (reported) preference profiles to (probability distributions over) candidates,

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

A **cardinal** *social choice function*, or *voting rule*, or *mechanism* is a function $f : U^n \rightarrow \Delta(A)$ mapping (reported) utility profiles to (probability distributions over) candidates,

where U^n is the space of all possible preference profiles.

Cardinal vs Ordinal (Randomised) Rules

In simple words: A cardinal voting rule is ordinal if it disregards the numbers and only keeps the information about the relative ranking between the candidates.

Expected Utility

Expected Utility

Given a reported utility profile s as input, the expected utility of a voter $i \in N$ from a randomised voting rule f is $\mathbb{E}[u_i(f(s))]$, where the expectation is taken over the randomness of the algorithm.

Expected Utility

Given a reported utility profile s as input, the expected utility of a voter $i \in N$ from a randomised voting rule f is $\mathbb{E}[u_i(f(s))]$, where the expectation is taken over the randomness of the algorithm.

A mechanism is truthful (*in expectation*) if for any reported utilities s_{-i} of the other voters, reporting the true utility u_i maximises the voter's expected utility, for all voters.

For the time being

For the time being

Unless stated otherwise, we will assume:

For the time being

Unless stated otherwise, we will assume:

- Our voting rules are deterministic.

For the time being

Unless stated otherwise, we will assume:

- Our voting rules are deterministic.
- Our voting rules are ordinal.

For the time being

Unless stated otherwise, we will assume:

- Our voting rules are deterministic.
- Our voting rules are ordinal.

Therefore we will work with the ordinal preference rankings $\succ_i$ without worrying about the cardinal utilities.

In the quest for truthful voting rules

Our first goal will be to design a voting rule that is truthful for the unrestricted domain.

Mechanisms

A *social choice function*, or *voting rule*, or *mechanism* is a function $f : (\succ)^n \to A$ 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 $\{1, \ldots, m\}$.

In the quest for truthful voting rules

Our first goal will be to design a voting rule that is truthful for the unrestricted domain.

Any ideas?

Fixed candidate

Fixed candidate

Completely disregard the preferences of the voters and always elect a fixed candidate, e.g., candidate α.

Fixed candidate

Completely disregard the preferences of the voters and always elect a fixed candidate, e.g., candidate α.

Formally, $f(\succ) = \alpha$ for all $\succ \in \succ^n$.

Fixed candidate

Completely disregard the preferences of the voters and always elect a fixed candidate, e.g., candidate α.

Formally, $f(\succ) = \alpha$ for all $\succ \in \succ^n$.

Intuitively, not a good solution. How to formalise this?

Fixed candidate

Completely disregard the preferences of the voters and always elect a fixed candidate, e.g., candidate a.

Formally, $f(\succ) = a$ for all $\succ \in \succ^n$.

Intuitively, not a good solution. How to formalise this?

It fails to satisfy the following *property* or *axiom*:

Fixed candidate

Completely disregard the preferences of the voters and always elect a fixed candidate, e.g., candidate α.

Formally, $f(\succ) = \alpha$ for all $\succ \in \succ^n$.

Intuitively, not a good solution. How to formalise this?

It fails to satisfy the following *property* or *axiom*:

Property (onto): A voting rule is *onto* or *surjective* if for any candidate β, there exists some preference profile $\succ$ such that $f(\succ) = \beta$.

Fixed candidate

Completely disregard the preferences of the voters and always elect a fixed candidate, e.g., candidate α.

Formally, $f(\succ) = \alpha$ for all $\succ \in \succ^n$.

Intuitively, not a good solution. How to formalise this?

It fails to satisfy the following *property* or *axiom*:

Property (onto): A voting rule is *onto* or *surjective* if for any candidate β, there exists some preference profile $\succ$ such that $f(\succ) = \beta$.

No candidate is *a-priori* excluded from consideration.

In the quest for truthful voting rules

Our refined goal will be to design a voting rule that is truthful and onto for the unrestricted domain.

Any ideas?

Dictatorship

Always elect the favourite candidate of a fixed voter, e.g., voter 1.

Dictatorship

Always elect the favourite candidate of a fixed voter, e.g., voter 1.

Formally, $f(\succ_i, \succ_{-i}) = \text{top}_{\succ_i}$ where $\text{top}_{\succ_i}$ is such that $\text{top}_{\succ_i} \succ_i \beta$ for any other candidate β.

Dictatorship

Always elect the favourite candidate of a fixed voter, e.g., voter 1.

Formally, $f(\succ_i, \succ_{-i}) = \text{top}_{\succ_i}$ where $\text{top}_{\succ_i}$ is such that $\text{top}_{\succ_i} \succ_i \beta$ for any other candidate β.

Is this mechanism onto?

In the quest for truthful voting rules

Our refined goal will be to design a voting rule that is truthful, onto and non-dictatorial for the unrestricted domain.

Any ideas?

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.

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.

This type of result is called a "characterisation". It identifies exactly the class of rules that are truthful and onto, as that of dictatorships.

Proving the GS Theorem

Proving the GS Theorem

There are several proofs of this theorem out there.

Proving the GS Theorem

There are several proofs of this theorem out there.

See the course webpage with links to some of them.

Proving the GS Theorem

There are several proofs of this theorem out there.

See the course webpage with links to some of them.

The proofs in the AGT book and the MAS book go via another impossibility theorem, called "Arrow's Theorem".

Proving the GS Theorem

There are several proofs of this theorem out there.

See the course webpage with links to some of them.

The proofs in the AGT book and the MAS book go via another impossibility theorem, called "Arrow's Theorem".

Here we will present a more direct proof that does not use that.

Proving the GS Theorem

There are several proofs of this theorem out there.

See the course webpage with links to some of them.

The proofs in the AGT book and the MAS book go via another impossibility theorem, called "Arrow's Theorem".

Here we will present a more direct proof that does not use that.

The proof has many steps, so we only present a sketch here.

A useful property

A useful property

Property (Pareto Optimality): A voting rule f is *Pareto optimal* or *Pareto efficient* if, for any candidate $\alpha \in A$, if there is another candidate $\beta \in A$ that is at least as good as α for all voters, and strictly better for one voter, then $f(\succ) \neq \alpha$, i.e., α cannot be the winner.

A useful property

Property (Pareto Optimality): A voting rule f is *Pareto optimal* or *Pareto efficient* if, for any candidate $\alpha \in A$, if there is another candidate $\beta \in A$ that is at least as good as α for all voters, and strictly better for one voter, then $f(\succ) \neq \alpha$, i.e., α cannot be the winner.

In our setup, we have strict preferences, so this translates to:

A useful property

Property (Pareto Optimality): A voting rule f is *Pareto optimal* or *Pareto efficient* if, for any candidate $\alpha \in A$, if there is another candidate $\beta \in A$ that is at least as good as α for all voters, and strictly better for one voter, then $f(\succ) \neq \alpha$, i.e., α cannot be the winner.

In our setup, we have strict preferences, so this translates to:

If there exists $\beta \in A$ such that $\beta \succ_i \alpha$ for all voters $i \in N$, then $f(\succ) \neq \alpha$.

A useful property

Property (Pareto Optimality): A voting rule f is *Pareto optimal* or *Pareto efficient* if, for any candidate $\alpha \in A$, if there is another candidate $\beta \in A$ that is at least as good as α for all voters, and strictly better for one voter, then $f(\succ) \neq \alpha$, i.e., α cannot be the winner.

In our setup, we have strict preferences, so this translates to:

If there exists $\beta \in A$ such that $\beta \succ_i \alpha$ for all voters $i \in N$, then $f(\succ) \neq \alpha$.

Then we prove the following useful lemma:

A useful property

Property (Pareto Optimality): A voting rule f is *Pareto optimal* or *Pareto efficient* if, for any candidate $\alpha \in A$, if there is another candidate $\beta \in A$ that is at least as good as α for all voters, and strictly better for one voter, then $f(\succ) \neq \alpha$, i.e., α cannot be the winner.

In our setup, we have strict preferences, so this translates to:

If there exists $\beta \in A$ such that $\beta \succ_i \alpha$ for all voters $i \in N$, then $f(\succ) \neq \alpha$.

Then we prove the following useful lemma:

Lemma: A voting rule f is *Pareto optimal* if it is *truthful* and *onto*.

A useful property

Property (Pareto Optimality): A voting rule f is *Pareto optimal* or *Pareto efficient* if, for any candidate $\alpha \in A$, if there is another candidate $\beta \in A$ that is at least as good as α for all voters, and strictly better for one voter, then $f(\succ) \neq \alpha$, i.e., α cannot be the winner.

In our setup, we have strict preferences, so this translates to:

If there exists $\beta \in A$ such that $\beta \succ_i \alpha$ for all voters $i \in N$, then $f(\succ) \neq \alpha$.

Then we prove the following useful lemma:

Lemma: A voting rule f is *Pareto optimal* if it is *truthful* and *onto*.

Proof: Tutorial.

GS Rephrased

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

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Voter i	Other Voters fixed
f	?
$\vdots$	$\vdots$
a	?
$\vdots$	$\vdots$
a'	?
$\vdots$	$\vdots$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i , \succ_{-i}) = \alpha$, and $f(\succ_i' , \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Voter i	Other Voters fixed
f	$?$
$\vdots$	$\vdots$
a	$?$
a'	
$\vdots$	$\vdots$
	$?$
$\vdots$	$\vdots$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Voter i	Other Voters fixed
f	$?$
$\vdots$	$\vdots$
a	$?$
$\vdots$	$\vdots$
	$?$
$\vdots$ a'	$\vdots$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Voter i	Other Voters fixed
f	$?$
$\vdots$	$\vdots$
a	$?$
$\vdots$	$\vdots$
a'	$?$
$\vdots$	$\vdots$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Voter i	Other Voters fixed
f	?
$\vdots$ a'	$\vdots$
a	?
$\vdots$	$\vdots$
	?
$\vdots$	$\vdots$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ'_i, \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ'_i \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ'_i$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Voter i	Other Voters fixed
f	?
$\vdots\ a'$	$\vdots$
a	?
$\vdots$	$\vdots$
	?
$\vdots$	$\vdots$

α' could be the winner

Another useful property

Property (Monotonicity): Let $f(\succ_i , \succ_{-i}) = \alpha$, and $f(\succ_i' , \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ'_i, \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ'_i \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ'_i$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$
A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.
In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

Voter i	Other Voters fixed
f	?
$\vdots$	$\vdots$
a	?
$\vdots$	$\vdots$
a'	?
$\vdots$	$\vdots$

Another useful property

Property (Monotonicity): Let $f(\succ_i , \succ_{-i}) = \alpha$, and $f(\succ_i' , \succ_{-i}) = \alpha'$
A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.
In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i	Other Voters fixed
	f	$?$
$\succ_i$	$\vdots$	$\vdots$
	a	$?$
	$\vdots$	$\vdots$
	a'	$?$
	$\vdots$	$\vdots$

Another useful property

Property (Monotonicity): Let $f(\succ_i , \succ_{-i}) = \alpha$, and $f(\succ_i' , \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i	Other Voters fixed
	f	?
$\succ_i$	$\vdots$	$\vdots$
	a	?
	$\vdots$	$\vdots$
	a'	?
	$\vdots$	$\vdots$

This is preference profile $(\succ_i , \succ_{-i})$

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ'_i, \succ_{-i}) = \alpha'$
A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ'_i \alpha$.
In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ'_i$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i		Other Voters fixed
$\succ_i$	f	g	?
	$\vdots$	$\vdots$	$\vdots$
	a	a	?
	$\vdots$	$\vdots$	$\vdots$
	a'	a'	?
	$\vdots$	$\vdots$	$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ'_i, \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ'_i \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ'_i$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ'_i$	?
	$\vdots$	$\vdots$		$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α

Another useful property

<u>Property (Monotonicity):</u> Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$
A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.
In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

<u>Lemma:</u> A voting rule f is *truthful* if and only if it is *monotone.*

<u>Proof:</u>

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ_i'$	?
	$\vdots$	$\vdots$		$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α

This is preference profile $(\succ_i', \succ_{-i})$

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ_i'$	?
	⋮	⋮		⋮
	a	a		?
	⋮	⋮		⋮
	a'	a'		?
	⋮	⋮		⋮

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α

This is preference profile $(\succ_i', \succ_{-i})$

Winner: α'

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

Assume that $\succ_i'$ is the true preference ranking.

	Voter i ($\succ_i$)	Voter i ($\succ_i'$)	Other Voters fixed
	f	g	?
	$\vdots$	$\vdots$	$\vdots$
	a	a	?
	$\vdots$	$\vdots$	$\vdots$
	a'	a'	?
	$\vdots$	$\vdots$	$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α

This is preference profile $(\succ_i', \succ_{-i})$

Winner: α'

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

Assume that $\succ_i'$ is the true preference ranking.

Voter i could report $\succ_i$ and make α win, which the voter prefers!

	Voter i ($\succ_i$)	Voter i ($\succ_i'$)	Other Voters fixed
	f	g	?
	⋮	⋮	⋮
	a	a	?
	⋮	⋮	⋮
	a'	a'	?
	⋮	⋮	⋮

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α

This is preference profile $(\succ_i', \succ_{-i})$

Winner: α'

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i			Other Voters fixed
	f	g		?
$\succ_i$	$\vdots$	$\vdots$	$\succ_i'$	$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$
A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.
In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ_i'$	?
	$\vdots$	$\vdots$		$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$
A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.
In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i		Other Voters fixed
	$\succ_i$	$\succ_i'$	
	f	g	?
	$\vdots$	$\vdots$	$\vdots$
	a	a	?
	$\vdots$	$\vdots$	$\vdots$
	a'	a'	?
	$\vdots$	$\vdots$	$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α'

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ'_i, \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ'_i \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ'_i$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i			Other Voters fixed
$\succ_i$	$\succ'_i$			
	f	g		?
	$\vdots$	$\vdots$		$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α'

This is preference profile $(\succ'_i, \succ_{-i})$

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ_i'$	?
	$\vdots$	$\vdots$		$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α'

This is preference profile $(\succ_i', \succ_{-i})$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

Assume that $\succ_i$ is the true preference ranking.

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ_i'$	?
	$\vdots$	$\vdots$		$\vdots$
	a	a		?
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		?
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α'

This is preference profile $(\succ_i', \succ_{-i})$

Winner: α

Another useful property

Property (Monotonicity): Let $f(\succ_i, \succ_{-i}) = \alpha$, and $f(\succ_i', \succ_{-i}) = \alpha'$

A voting rule f is *monotone* if $a \neq a'$ implies that $\alpha \succ_i \alpha'$ and $\alpha' \succ_i' \alpha$.

In other words, if voter i changes its (reported) ranking from $\succ_i$ to $\succ_i'$, the winner changes from α to α' only if α' moved from below α to above α in the ranking.

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Proof:

Assume that $\succ_i$ is the true preference ranking.

Voter i could report $\succ_i'$ and make α' win, which the voter prefers!

	Voter i			Other Voters fixed
$\succ_i$	f	g	$\succ_i'$	$?$
	$\vdots$	$\vdots$		$\vdots$
	a	a		$?$
	$\vdots$	$\vdots$		$\vdots$
	a'	a'		$?$
	$\vdots$	$\vdots$		$\vdots$

This is preference profile $(\succ_i, \succ_{-i})$

Winner: α'

This is preference profile $(\succ_i', \succ_{-i})$

Winner: α

Standard Truthfulness Argument

This is a standard argument using truthfulness:

Consider two profiles $(\succ_i , \succ_{-i})$ and $(\succ'_i , \succ_{-i})$.

Assume that $f(\succ_i , \succ_{-i}) = \alpha$. Let α' be such that the relative order of α and α' is the same in both $\succ_i$ and $\succ'_i$.

Truthfulness implies that $f(\succ'_i , \succ_{-i}) = \alpha$.

Another useful property

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Corollary:

Voter i		Voter j	
f		f	
$\vdots$		a	
a	$\cdots$	$\vdots$	$\cdots$
$\vdots$		$\vdots$	
a'		a'	
$\vdots$		$\vdots$	

Winner: α

Another useful property

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Corollary:

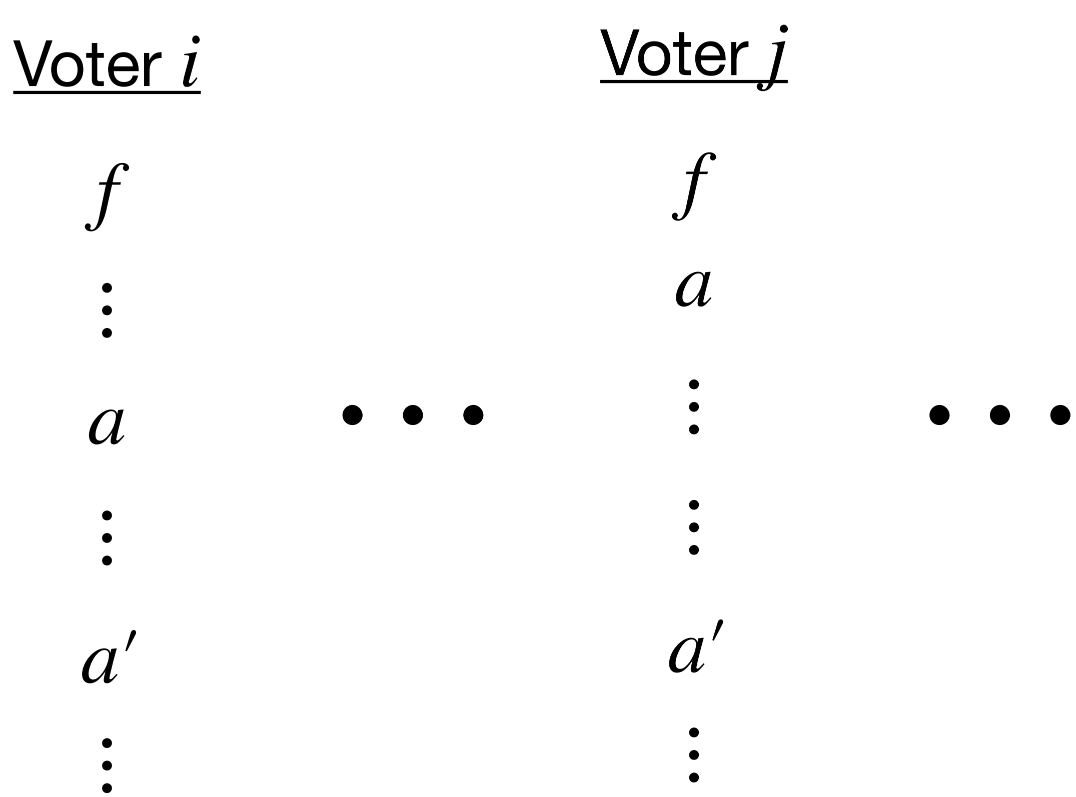

Another useful property

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Corollary:

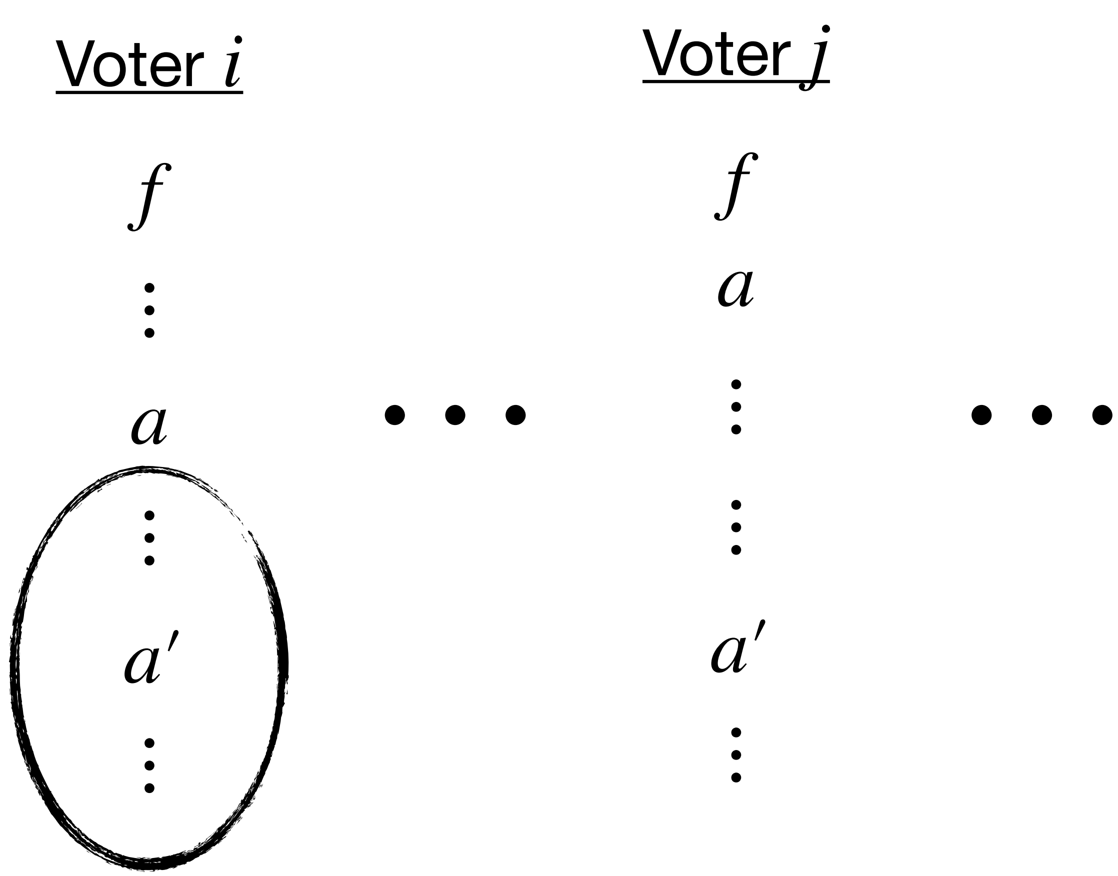

Another useful property

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Corollary:

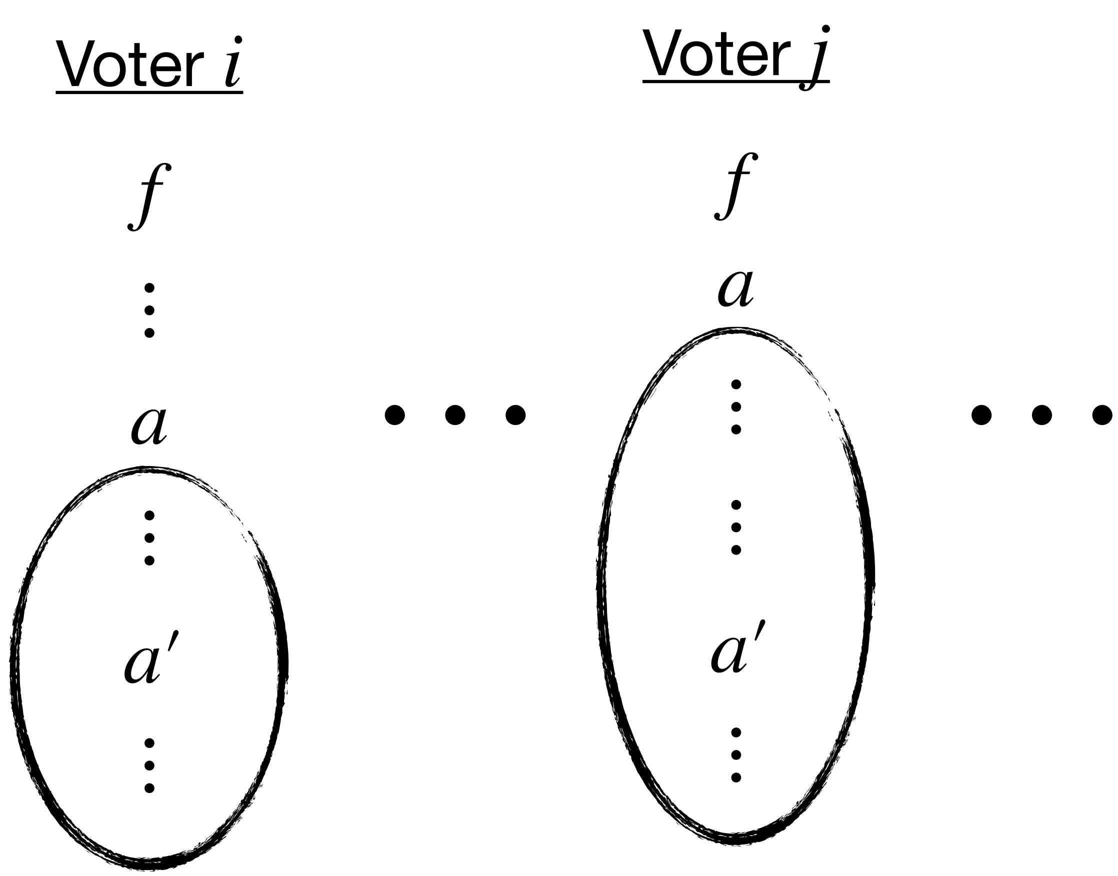

Another useful property

Lemma: A voting rule f is *truthful* if and only if it is *monotone.*

Corollary:

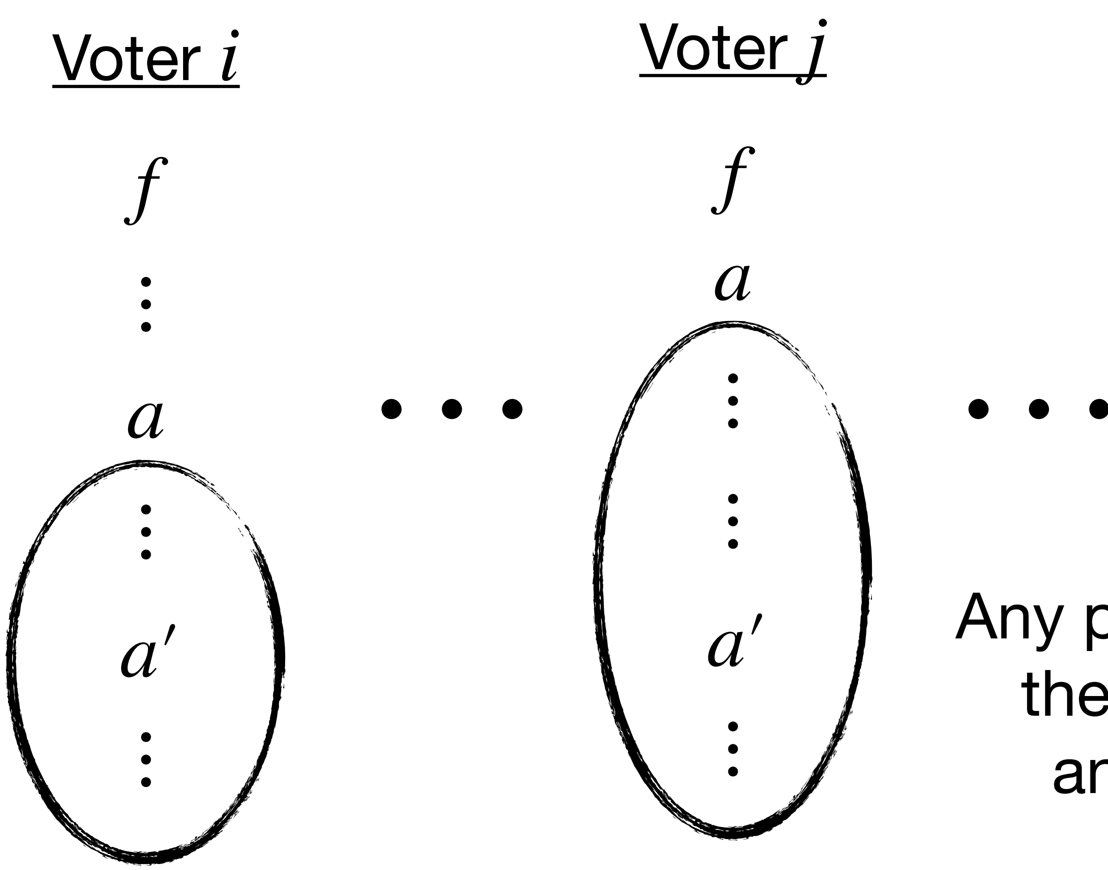

Any permutation of the rankings for these candidates cannot make any of these candidates win.

A final useful lemma

Lemma 1: Consider a truthful voting rule f such that $f(\succ_i, \succ_{-i}) = \beta$. Let $\succ_i'$ be such that some candidate x is ranked higher in $\succ_i'$ than in $\succ$. Then $f(\succ_i', \succ_{-i}) \in \{\beta, x\}$.

Voter i		Voter j	
f		?	
$\vdots$		?	
β	$\cdots$	$\vdots$	$\cdots$
$\vdots$		$\vdots$	
x		?	
$\vdots$		$\vdots$	

Winner: β

A final useful lemma

Lemma 1: Consider a truthful voting rule f such that $f(\succ_i, \succ_{-i}) = \beta$. Let $\succ_i'$ be such that some candidate x is ranked higher in $\succ_i'$ than in $\succ$. Then $f(\succ_i', \succ_{-i}) \in \{\beta, x\}$.

Voter i		Voter j	
f		?	
$\vdots$		?	
x	$\cdots$	$\vdots$	$\cdots$
$\vdots$		$\vdots$	
β		?	
$\vdots$		$\vdots$	

Winner: β

A final useful lemma

Lemma 1: Consider a truthful voting rule f such that $f(\succ_i, \succ_{-i}) = \beta$. Let $\succ_i'$ be such that some candidate x is ranked higher in $\succ_i'$ than in $\succ$. Then $f(\succ_i', \succ_{-i}) \in \{\beta, x\}$.

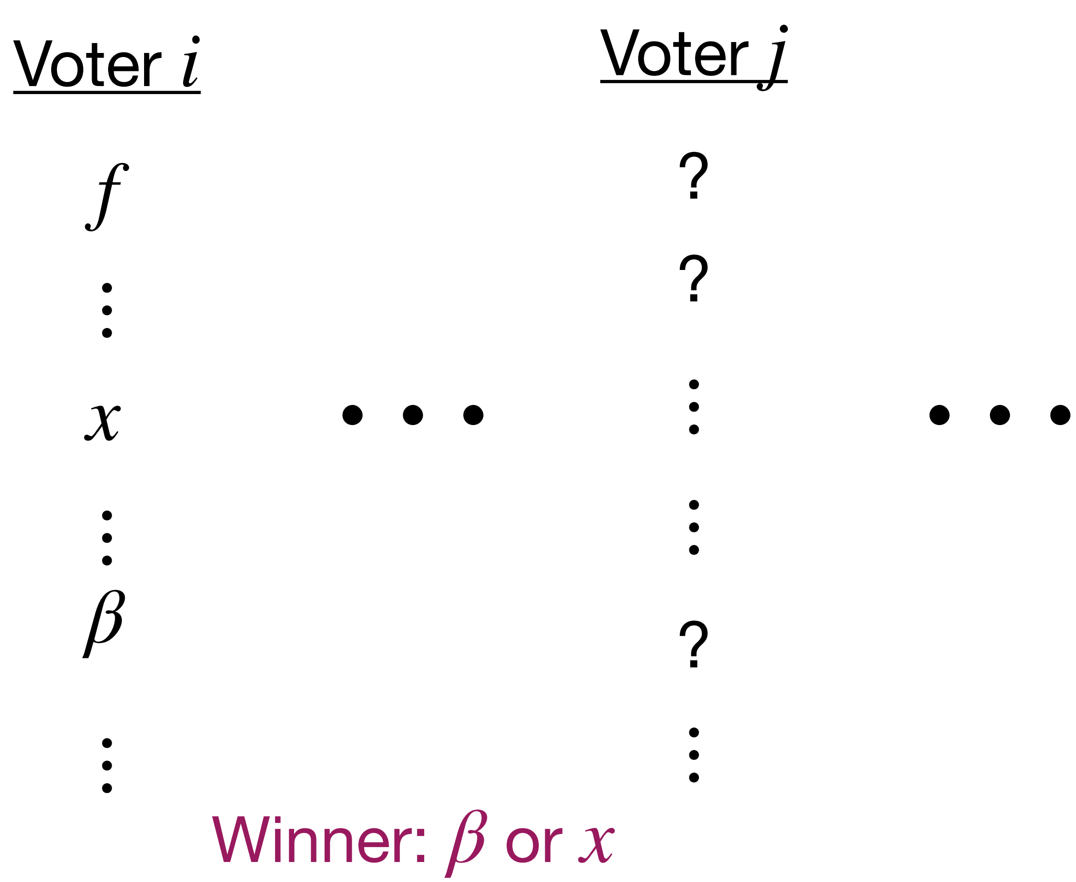

A final useful lemma

Lemma 1: Consider a truthful voting rule f such that $f(\succ_i, \succ_{-i}) = \beta$. Let $\succ_i'$ be such that some candidate x is ranked higher in $\succ_i'$ than in $\succ_i$. Then $f(\succ_i', \succ_{-i}) \in \{\beta, x\}$.

The proof follows by monotonicity.

Voter i		Voter j	
f		?	
$\vdots$		?	
x	$\cdots$	$\vdots$	$\cdots$
$\vdots$		$\vdots$	
β		?	
$\vdots$		$\vdots$	

Winner: β or x

Corollary

Corollary:

Voter i		Voter j	
f		f	
$\vdots$		a	
a	$\cdots$	$\vdots$	$\cdots$
$\vdots$		$\vdots$	
a'		a'	
$\vdots$		$\vdots$	

Winner: α

Corollary

Corollary:

Voter i		Voter j	
f		f	
$\vdots$		a	
a	$\cdots$	$\vdots$	$\cdots$
$\vdots$		$\vdots$	
a'		a'	
$\vdots$		$\vdots$	

Corollary

Corollary:

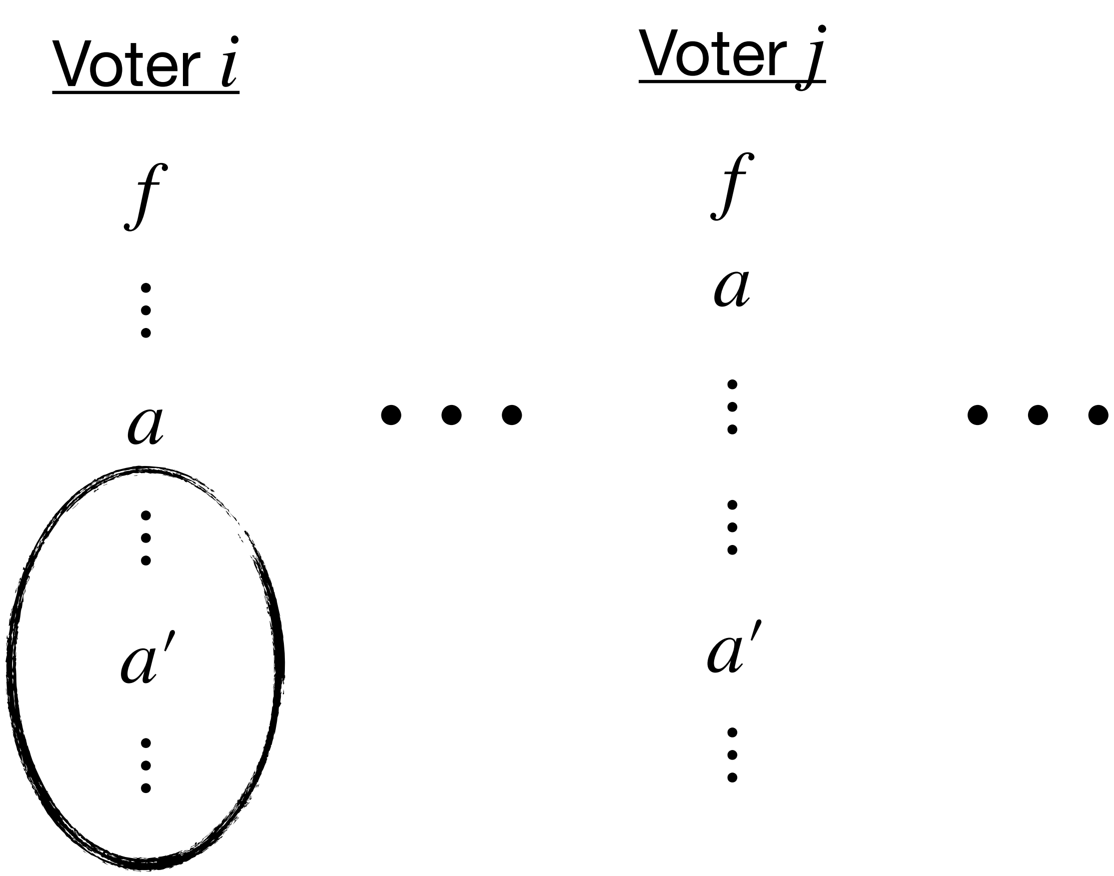

Corollary

Corollary:

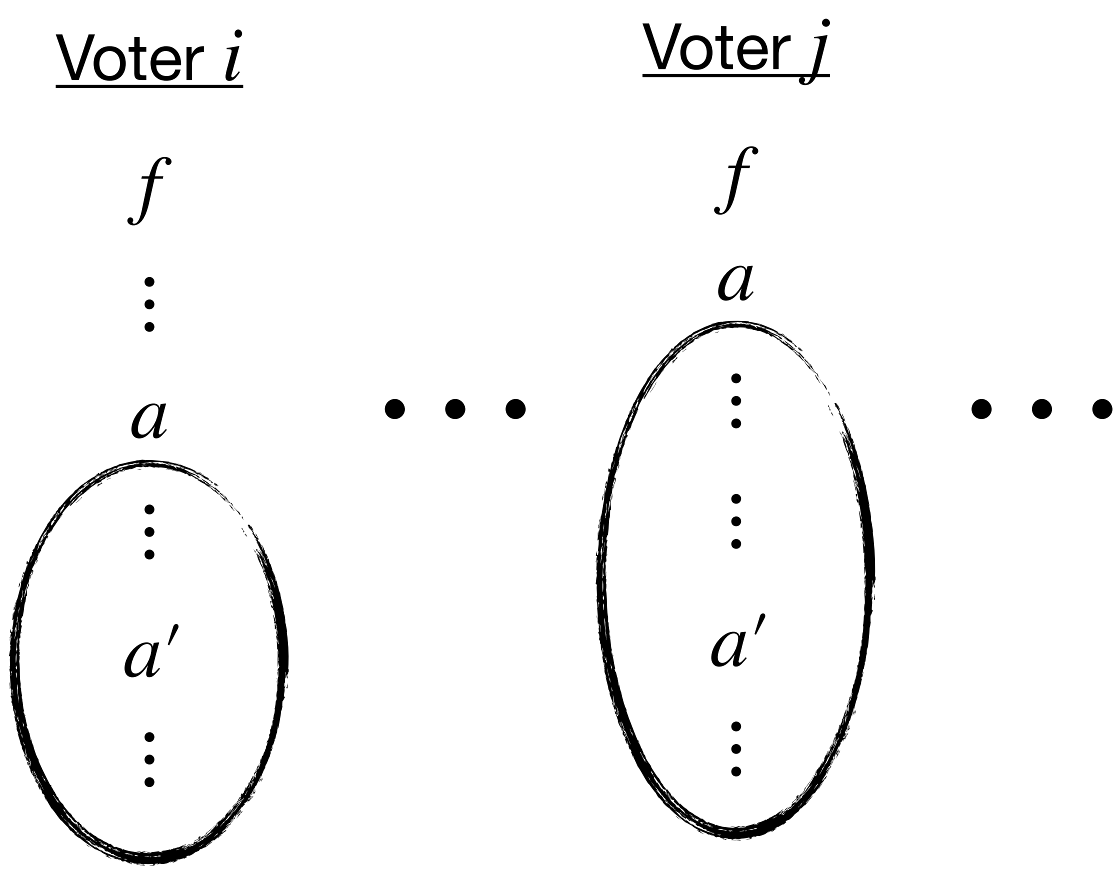

Corollary

Corollary:

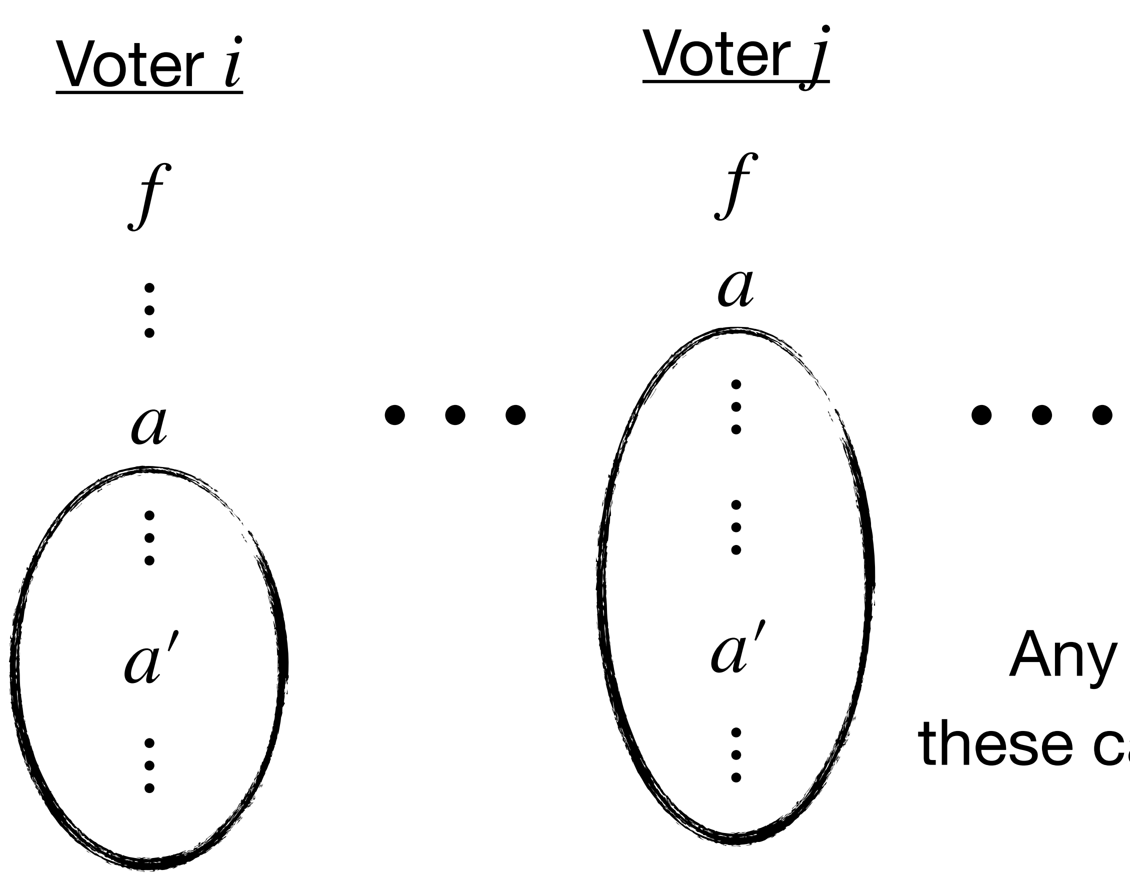

Any permutation of the rankings for these candidates still results in α winning.

Proof of GS

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

Claim: $f(\succ) \neq b$. Why?

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

Claim: $f(\succ) \neq b$. By *Pareto optimality*.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
f	a		ℓ	a	a		z
a	g		h	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
t	s		s	t	s		g
b	b		b	b	b		b

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	a		ℓ	a	a		z
f	g		h	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	s		s	t	s		g
t	b		b	b	b		b

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	a		ℓ	a	a		z
f	g		h	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	s		s	t	s		g
t	b		b	b	b		b

By Lemma 1, either we still select y, or we switch to b.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		ℓ	a	a		z
f	a		h	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		s	t	s		g
t	s		b	b	b		b

By Lemma 1, either we still select y, or we switch to b.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	a	a		z
f	a		ℓ	z	g		a
⋮	⋮		⋮	⋮	⋮		⋮
g	ℓ		y	t	s		g
t	s		s	b	b		b

By Lemma 1, either we still select y, or we switch to b.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	a	a		z
f	a		ℓ	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	t	s		g
t	s		s	b	b		b

By Lemma 1, either we still select y, or we switch to b.

At some point we reach the *pivotal voter* r, where we switch to b.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

By Lemma 1, either we still select y, or we switch to b.

At some point we reach the *pivotal voter* r, where we switch to b.

Proof of GS

Consider any preference profile $\succ$ in which all voters rank candidate b last.

We established: $f(\succ) \neq b$. Assume $f(\succ) = y$ for some $y \in A$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

By Lemma 1, either we still select y, or we switch to b.

At some point we reach the *pivotal voter* r, where we switch to b. Why?

Proof of GS

On this profile $f(\succ) = b$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

Proof of GS

On this profile $f(\succ) = b$.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

For any other ranking $\succ_i'$ of any voter in here, what is the outcome?

Proof of GS

On this profile $f(\succ) = b$.

It has to be b by monotonicity.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

For any other ranking $\succ_i'$ of any voter in here, what is the outcome?

Proof of GS

On this profile $f(\succ) = b$.

It has to be b by monotonicity.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

For any other ranking $\succ_i'$ that has b first of any voter in here, what is the outcome?

For any other ranking $\succ_i'$ of any voter in here, what is the outcome?

Proof of GS

On this profile $f(\succ) = b$.

It has to be b by monotonicity.

It has to be b by monotonicity.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	a		z
f	a		ℓ	a	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	h	s		g
t	s		s	t	b		b

For any other ranking $\succ'_i$ that has b first of any voter in here, what is the outcome?

For any other ranking $\succ'_i$ of any voter in here, what is the outcome?

Proof of GS

We have established: Whenever the first r voters rank b first, then b is elected.

Consider the profile before the change of the pivotal voter (where the outcome is still y).

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	a	a		z
f	a		ℓ	z	g		a
⋮	⋮		⋮	⋮	⋮		⋮
g	ℓ		y	t	s		g
t	s		s	b	b		b

Proof of GS

We have established: Whenever the first r voters rank b first, then b is elected.

Consider the profile before the change of the pivotal voter (where the outcome is still y).

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	a	a		z
f	a		ℓ	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	t	s		g
t	s		s	b	b		b

For any other ranking $\succ'_i$ of any voter in here, what is not the outcome?

Proof of GS

We have established: Whenever the first r voters rank b first, then b is elected.

Consider the profile before the change of the pivotal voter (where the outcome is still y).

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	a	a		z
f	a		ℓ	z	g		a
$\vdots$	$\vdots$		$\vdots$	$\vdots$	$\vdots$		$\vdots$
g	ℓ		y	t	s		g
t	s		s	b	b		b

For any other ranking $\succ_i'$ of any voter in here, what is not the outcome?

For any other ranking $\succ_i'$ of any voter in here where b is ranked last, what is not the outcome?

Established Facts

Established Facts

We have established:

Established Facts

We have established:

- Whenever the first r voters rank b first, then b is elected.

Established Facts

We have established:

- Whenever the first r voters rank b first, then b is elected.
- Whenever the last $n - r + 1$ voters rank b last, then b is not elected.

Established Facts

We have established:

- Whenever the first r voters rank b first, then b is elected.
- Whenever the last $n - r + 1$ voters rank b last, then b is not elected.

We will prove that voter r is a dictator.

Proof of GS

Consider the following profile, called Profile 1:

What is the elected candidate?

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	b	b		b

Proof of GS

Consider the following profile, called Profile 1:

What is the elected candidate?

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	b	b		b

By Pareto optimality, the only options are k and b.

Proof of GS

Consider the following profile, called Profile 1:

What is the elected candidate?

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	b	b		b

By Pareto optimality, the only options are k and b.
By our established facts, it cannot be b.

Proof of GS

From Profile 1, obtain Profile 2 by raising b to the second position in the ranking of voter r.

What is the elected candidate?

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	b	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

From Profile 1, obtain Profile 2 by raising b to the second position in the ranking of voter r.

What is the elected candidate? By monotonicity, it is still k.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	b	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

Consider any profile of the following form (call this Profile 3).

Assume that some candidate $g \neq k$ is chosen.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
?	?		?	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
b	b		b	b	b		b

Proof of GS

From Profile 3, obtain Profile 4 by moving b to the top of the first r-1 voters.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	b	b		b

Proof of GS

From Profile 3, obtain Profile 4 by moving b to the top of the first r-1 voters.

By Lemma 1, only g and b can be chosen. By our established facts, b is not chosen.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	b	b		b

Proof of GS

From Profile 4, obtain Profile 5 by moving b to the second position in the ranking of voter r.

By Lemma 1, either g or b will be chosen.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	?		?
?	?		?	b	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

If we switch the position of b and k in the ranking of voter r, the same outcome as before should be chosen by monotonicity (since the relative order of g and b as not changed).

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	?		?
?	?		?	k	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

If we switch the position of b and k in the ranking of voter r, the same outcome as before should be chosen by monotonicity (since the relative order of g and b as not changed).

By our established facts 1, b will be chosen, so b was also chosen before.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	b	?		?
?	?		?	k	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

From Profile 4, obtain Profile 5 by moving b to the second position in the ranking of voter r.

We established that b will be chosen.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	?		?
?	?		?	b	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

Then we can move k to the second position of the first r-1 voters and the first position of the remaining voters, and argue that b is still chosen. This follows from monotonicity, since the relative order of b and k has not changed for any voter.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	b	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

But this is Profile 2 that we saw before, and argued that the elected candidate is k. This is a contradiction, which means that…

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
b	b		b	k	k		k
k	k		k	b	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	b		b

Proof of GS

Consider any profile of the following form (call this Profile 3).

Assume that some candidate $g \neq k$ is chosen.

<u>This is not possible</u>, so k is chosen.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
?	?		?	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
b	b		b	b	b		b

Proof of GS

Consider any profile of the following form (call this Profile 3).

Assume that some candidate $g \neq k$ is chosen.

<u>This is not possible</u>, so k is chosen.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
?	?		?	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
b	b		b	b	b		b

We are almost there. We have proved that for almost any preference profile, the elected candidate is the top choice of voter r.

Proof of GS

To conclude the proof, we need to construct a few more arguments of the same flavour, to argue that voter r's top choice is always selected on every profile.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
?	?		?	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	?		?

Proof of GS

To conclude the proof, we need to construct a few more arguments of the same flavour, to argue that voter r's top choice is always selected on every profile.

Voter 1	Voter 2	…	Voter r-1	Voter r	Voter r+1	…	Voter n
?	?		?	k	?		?
?	?		?	?	?		?
⋮	⋮		⋮	⋮	⋮		⋮
?	?		?	?	?		?
?	?		?	?	?		?

The proof uses again and again the same arguments: it "moves things around" and argues using monotonicity , Pareto optimality, Lemma 1, and our established facts.

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.

This type of result is called a "characterisation". It identifies exactly the class of rules that are truthful and onto, as that of dictatorships.

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.

This type of result is called a "characterisation". It identifies exactly the class of rules that are truthful and onto, as that of dictatorships.

Should we give up all hope?

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.

This type of result is called a "characterisation". It identifies exactly the class of rules that are truthful and onto, as that of dictatorships.

Should we give up all hope?

No! Many domains of interest are restricted.

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.

This type of result is called a "characterisation". It identifies exactly the class of rules that are truthful and onto, as that of dictatorships.

Should we give up all hope?

No! Many domains of interest are restricted.

In other words, the preferences have some structure and the GS theorem does not apply! More on that next time.