Probabilistic social choice

Hans Peters

Moscow November 2019

An example

Ann, Bob, and Chris have to decide on the color to paint their joint apartment. They can choose from five colors: blue, green, red, gray, and black

An example

Ann, Bob, and Chris have to decide on the color to paint their joint apartment. They can choose from five colors: blue, green, red, gray, and black

Their individual rankings are as follows:

Ann: blue > black > red > gray > green

Bob: black > gray > red > green > blue

Chris: red > green > blue > gray > black

Which color should they use?

Borda (1781)

Assign scores $4, \ldots, 0$ to the colors (4 is best), and add up the scores

Borda (1781)

Assign scores $4, \ldots, 0$ to the colors (4 is best), and add up the scores

	blue	green	red	gray	black
Ann	4	0	2	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	4	8	5	7

Borda (1781)

Assign scores $4, \ldots, 0$ to the colors (4 is best), and add up the scores

	blue	green	red	gray	black
Ann	4	0	2	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	4	8	5	7

Winner: red

Is this a good method?

Depends on perspective: we look at **manipulability**

Borda is manipulable (Gibbard, 1973;Satterthwaite, 1975)

	blue	green	red	gray	black
Ann	4	0	2	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	4	8	5	7

Red wins. Now suppose Ann lies about her true ranking:

Borda is manipulable (Gibbard, 1973;Satterthwaite, 1975)

	blue	green	red	gray	black
Ann	4	0	2	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	4	8	5	7

Red wins. Now suppose Ann lies about her true ranking:

	blue	green	red	gray	black
Ann	4	2	0	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	6	6	5	7

Borda is manipulable (Gibbard, 1973;Satterthwaite, 1975)

	blue	green	red	gray	black
Ann	4	0	2	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	4	8	5	7

Red wins. Now suppose Ann lies about her true ranking:

	blue	green	red	gray	black
Ann	4	2	0	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	6	6	6	5	7

Black wins instead of red, which is good for Ann!

Single-peaked preferences

	blue	green	red	gray	black
Ann	<u>4</u>	3	2	1	0
Bob	0	1	2	3	<u>4</u>
Chris	2	3	<u>4</u>	1	0

Red is the median of the peaks

Single-peaked preferences

	blue	green	red	gray	black
Ann	<u>4</u>	3	2	1	0
Bob	0	1	2	3	<u>4</u>
Chris	2	3	<u>4</u>	1	0

Red is the median of the peaks

Preferences are now single-peaked with respect to the ordering: blue - green - red - gray - black

The median cannot be manipulated if preferences are single-peaked and only single-peaked preferences can be reported. (More convincing in case of room temperature)

Goes back to Black (1948). Later: Moulin (1980) and others.

Probabilistic Borda

Assign probabilities to the alternatives based on the Borda scores

	blue	green	red	gray	black
Ann	4	0	2	1	3
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Probabilities	$\frac{6}{30}$	$\frac{4}{30}$	$\frac{8}{30}$	$\frac{5}{30}$	$\frac{7}{30}$

No one can unilaterally increase the probability on his/her best, two best, three best, or four best alternatives! The sincere lottery stochastically dominates any lottery achievable by manipulation

Price paid: every one prefers red above green but still green gets positive probability

Gibbard (1977): only random dictatorship possible, e.g., $\frac{1}{3}$ probability on blue, red, and black

Probabilistic Borda and single-peaked preferences

(Lack of) Pareto optimality is still a problem!

Probabilistic Borda and single-peaked preferences

(Lack of) Pareto optimality is still a problem!

	blue	green	red	gray	black
Ann	3	4	2	1	0
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	$\frac{5}{30}$	$\frac{8}{30}$	$\frac{8}{30}$	$\frac{5}{30}$	$\frac{4}{30}$

Blue is Pareto dominated by green...but still gets positive probability

Probabilistic Borda and single-peaked preferences

(Lack of) Pareto optimality is still a problem!

	blue	green	red	gray	black
Ann	3	4	2	1	0
Bob	0	1	2	3	4
Chris	2	3	4	1	0
Total	$\frac{5}{30}$	$\frac{8}{30}$	$\frac{8}{30}$	$\frac{5}{30}$	$\frac{4}{30}$

Blue is Pareto dominated by green...but still gets positive probability

The central question of this presentation will be: suppose preferences are single-peaked, then which probabilistic rules are non-manipulable (or strategy-proof) and unanimous?

Color choice revisited

We assume anonymity and unanimity

Color choice revisited

We assume anonymity and unanimity

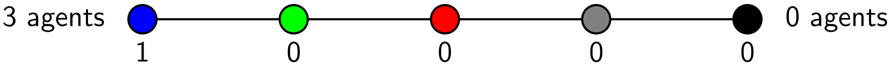

Color choice revisited

We assume anonymity and unanimity

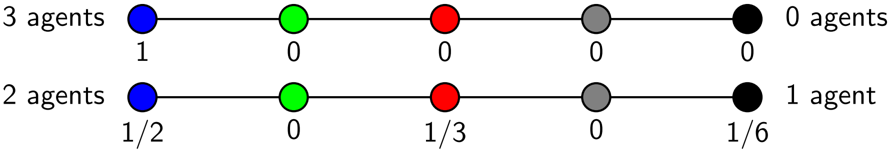

Color choice revisited

We assume anonymity and unanimity

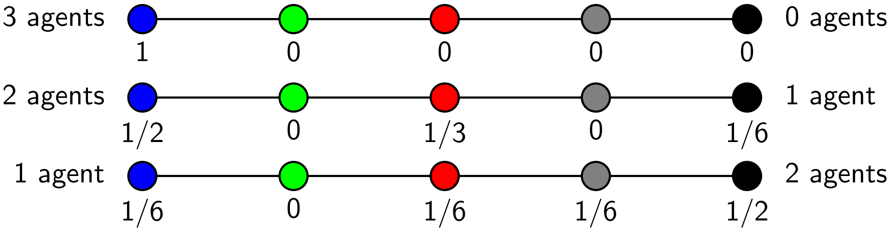

Color choice revisited

We assume anonymity and unanimity

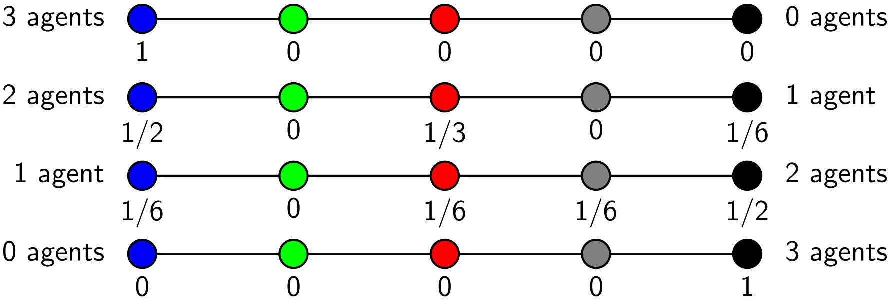

Color choice revisited

We assume anonymity and unanimity

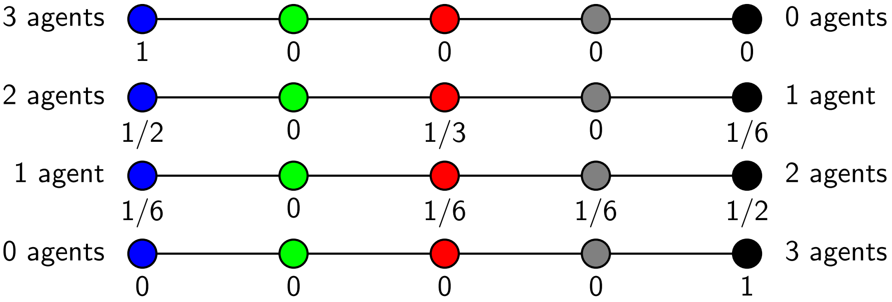

These probability distributions completely determine a probabilistic rule that is unanimous and non-manipulable (and anonymous)

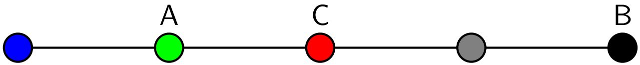
A
C
B

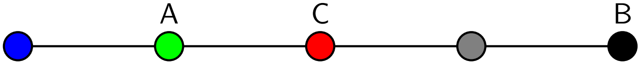

Probability at blue? Shift all agents at the left of blue to the left end point and all others to the right end point. Results in distribution (0,3)

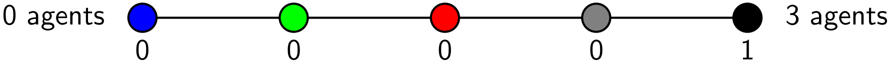

Assign to blue the probability assigned to blue by this distribution. Hence blue gets 0

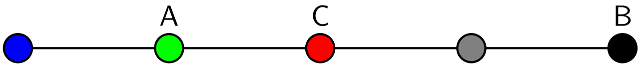
A
C
B

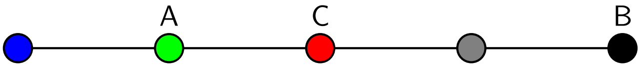

Probability at green?
Shift all agents at the left of green or on green to the left end point and all others to the right end point. Results in distribution (1,2)
Shift all agents at the left of green to the left end point and all others to the right end point. Results in distribution (0,3)

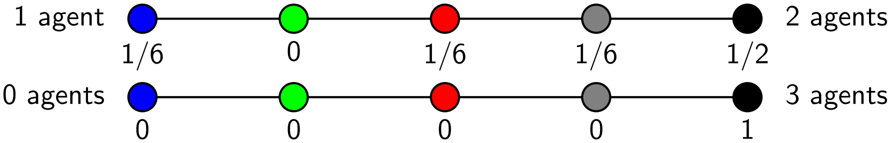

Now green gets $1/6 + 0$ (namely blue and green at distribution (1,2)) minus 0 (namely blue at distribution (0,3)), hence $1/6$

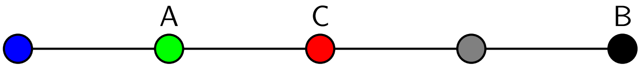
A
C
B

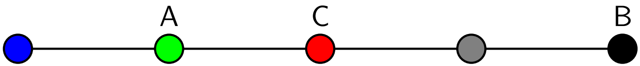

Probability at red?

Shift all agents at the left of red or on red to the left end point and all others to the right end point. Results in distribution (2,1)

Shift all agents at the left of red to the left end point and all others to the right end point. Results in distribution (1,2)

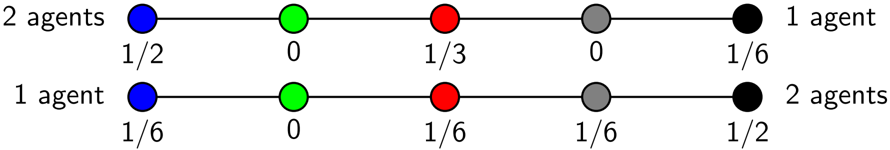

Now red gets $1/2 + 0 + 1/3$ (namely blue, green and red at distribution (2,1)) minus $1/6 + 0$ (namely blue and green at distribution (1,2)), hence $2/3$

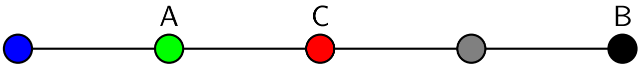
A
C
B

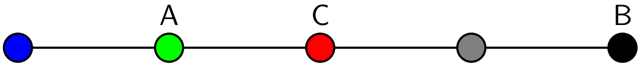

Probability at gray?

There is not peak at gray, so (similar to the case of blue) we take the probability assigned by the (2,1) distribution to gray

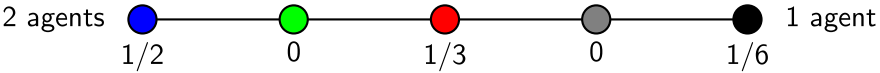

Hence gray gets 0

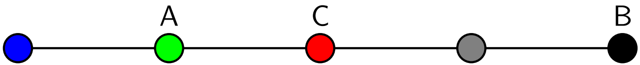
A
C
B

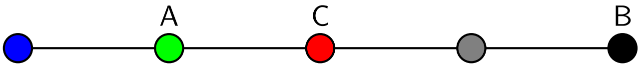

Probability at black?
Shift all agents at the left of black or on black to the left end point and all others to the right end point. Results in distribution (3,0)
Shift all agents at the left of black to the left end point and all others to the right end point. Results in distribution (2,1)

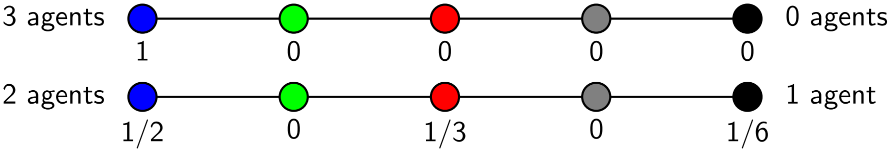

Now black gets 1 (from distribution (3,0)) minus $1/2 + 1/3$ (assigned by distribution (2,1) to all points left from black), hence $1/6$

The last example was based on Ehlers, Peters, Storcken (2002)
We will extend this to more general domains of preferences, where single-peakedness can be defined with respect to a connected graph

The last example was based on Ehlers, Peters, Storcken (2002)
We will extend this to more general domains of preferences, where single-peakedness can be defined with respect to a connected graph

Outline of the rest of the talk

1. Introduction: general model
2. Probabilistic rules
3. Domain restrictions: single-peaked preferences
4. Probabilistic rules and single-peakedness
5. Probabilistic rules for single-peakedness preferences on graphs
6. Trees
7. Leafless graphs
8. General connected graphs
9. Concluding remarks

1.Introduction: general model

Starting point is the classical social choice model:

- $N = \{1, \ldots, n\}$ is the set of (at least two) *agents*
- A is the (usually) finite set of (at least two) *alternatives*
- *Preferences* of agents over alternatives are linear orders
- A *preference profile* is an n-tuple of preferences
- A *social choice function* or *rule* assigns to each preference profile an alternative

Examples: (political) elections, decisions within committees, European songfestival

1.Introduction: general model

Starting point is the classical social choice model:

- $N = \{1, \ldots, n\}$ is the set of (at least two) *agents*
- A is the (usually) finite set of (at least two) *alternatives*
- *Preferences* of agents over alternatives are linear orders
- A *preference profile* is an n-tuple of preferences
- A *social choice function* or *rule* assigns to each preference profile an alternative

Examples: (political) elections, decisions within committees, European songfestival

Throughout we concentrate on strategy-proofness: each agent reports a preference and should not be able to benefit from not reporting the true preference

Formally:

A rule F is *strategy-proof* if for each preference profile R^N, each agent i, and each preference Q^i we have:

$$F(R^N) R^i F(R^{N\setminus i}, Q^i)$$

i.e., truth-telling is a 'weakly dominant strategy'

Formally:

A rule F is *strategy-proof* if for each preference profile R^N, each agent i, and each preference Q^i we have:

$$F(R^N) R^i F(R^{N \setminus i}, Q^i)$$

i.e., truth-telling is a 'weakly dominant strategy'

Strategy-proofness is desirable:

- Makes voting easy for the agents: only knowledge of your own preference is needed to vote optimally

Formally:

A rule F is *strategy-proof* if for each preference profile R^N, each agent i, and each preference Q^i we have:

$$F(R^N) R^i F(R^{N \setminus i}, Q^i)$$

i.e., truth-telling is a 'weakly dominant strategy'

Strategy-proofness is desirable:

- Makes voting easy for the agents: only knowledge of your own preference is needed to vote optimally
- Preserves (ex post) the (other) desirable properties of a rule
- Decisions are made on the basis of the right information
- ...

But strategy-proofness is hard to get:

Theorem (Gibbard, 1973; Satterthwaite, 1975)

Let F have range of at least three. Then F is strategy-proof if and only if it is *dictatorial* on its range, i.e., there is an agent d such that $F(R^N)$ is the top alternative of d's preference R^d in the range of F, for each preference profile R^N

But strategy-proofness is hard to get:

Theorem (Gibbard, 1973; Satterthwaite, 1975)

Let F have range of at least three. Then F is strategy-proof if and only if it is *dictatorial* on its range, i.e., there is an agent d such that $F(R^N)$ is the top alternative of d's preference R^d in the range of F, for each preference profile R^N

If the range is two or if $|A| = 2$, then we can simply use majority (plurality) voting (May, 1952)

How to ‘escape’ from Gibbard-Satterthwaite?

How to ‘escape’ from Gibbard-Satterthwaite?

- ‘Minimize’ manipulability (Kelly, 1988; Fristrup and Keiding, 1998; Aleskerov and Kurbanov, 1999; Maus et al, 2007; Arribillaga and Masso, 2016; etc.)

How to 'escape' from Gibbard-Satterthwaite?

- 'Minimize' manipulability (Kelly, 1988; Fristrup and Keiding, 1998; Aleskerov and Kurbanov, 1999; Maus et al, 2007; Arribillaga and Masso, 2016; etc.)
- Make manipulability hard to compute (Bartholdi et al, 1989; Friedgut et al, 2011; etc.)

How to 'escape' from Gibbard-Satterthwaite?

- 'Minimize' manipulability (Kelly, 1988; Fristrup and Keiding, 1998; Aleskerov and Kurbanov, 1999; Maus et al, 2007; Arribillaga and Masso, 2016; etc.)
- Make manipulability hard to compute (Bartholdi et al, 1989; Friedgut et al, 2011; etc.)
- Make manipulation 'unsafe' (e.g., Reijngoud and Endriss, 2012; Veselova, ongoing work)

How to 'escape' from Gibbard-Satterthwaite?

- 'Minimize' manipulability (Kelly, 1988; Fristrup and Keiding, 1998; Aleskerov and Kurbanov, 1999; Maus et al, 2007; Arribillaga and Masso, 2016; etc.)
- Make manipulability hard to compute (Bartholdi et al, 1989; Friedgut et al, 2011; etc.)
- Make manipulation 'unsafe' (e.g., Reijngoud and Endriss, 2012; Veselova, ongoing work)

But we are going to concentrate on:

- Restrict the set of preference profiles (Black, 1948; Moulin, 1980; etc.)

How to ‘escape’ from Gibbard-Satterthwaite?

- ‘Minimize’ manipulability (Kelly, 1988; Fristrup and Keiding, 1998; Aleskerov and Kurbanov, 1999; Maus et al, 2007; Arribillaga and Masso, 2016; etc.)
- Make manipulability hard to compute (Bartholdi et al, 1989; Friedgut et al, 2011; etc.)
- Make manipulation ‘unsafe’ (e.g., Reijngoud and Endriss, 2012; Veselova, ongoing work)

But we are going to concentrate on:

- Restrict the set of preference profiles (Black, 1948; Moulin, 1980; etc.)
- Probabilistic rules (Gibbard, 1977, 1978; Barberà, 1979; Dutta et al, 2002; Ehlers et al, 2002; Chatterji et al, 2014; ....etc.)

2. Probabilistic rules

A *probabilistic rule* F assigns to every preference profile R^N a probability distribution $F(R^N)$ over A

2. Probabilistic rules

A *probabilistic rule* F assigns to every preference profile R^N a probability distribution $F(R^N)$ over A

A probabilistic rule is *strategy-proof* if for every preference profile R^N, every agent i and every preference Q^i, the distribution $F(R^N)$ stochastically dominates the distribution $F(R^{N\setminus\{i\}}, Q^i)$ according to R^i, that is,

2. Probabilistic rules

A *probabilistic rule* F assigns to every preference profile R^N a probability distribution $F(R^N)$ over A

A probabilistic rule is *strategy-proof* if for every preference profile R^N, every agent i and every preference Q^i, the distribution $F(R^N)$ stochastically dominates the distribution $F(R^{N\setminus\{i\}}, Q^i)$ according to R^i, that is,

- for every alternative x, the probability assigned to the set $\{y \in A \mid yR^i x\}$ by $F(R^N)$ is at least as large as the probability assigned by $F(R^{N\setminus\{i\}}, Q^i)$

2. Probabilistic rules

A *probabilistic rule* F assigns to every preference profile R^N a probability distribution $F(R^N)$ over A

A probabilistic rule is *strategy-proof* if for every preference profile R^N, every agent i and every preference Q^i, the distribution $F(R^N)$ stochastically dominates the distribution $F(R^{N\setminus\{i\}}, Q^i)$ according to R^i, that is,

- for every alternative x, the probability assigned to the set $\{y \in A \mid yR^i x\}$ by $F(R^N)$ is at least as large as the probability assigned by $F(R^{N\setminus\{i\}}, Q^i)$

or equivalently,

- $\mathbb{E}_{F(R^N)}(u^i) \geq \mathbb{E}_{F(R^{N\setminus\{i\}}, Q^i)}(u^i)$ for every u^i representing R^i

2. Probabilistic rules

A *probabilistic rule* F assigns to every preference profile R^N a probability distribution $F(R^N)$ over A

A probabilistic rule is *strategy-proof* if for every preference profile R^N, every agent i and every preference Q^i, the distribution $F(R^N)$ stochastically dominates the distribution $F(R^{N\setminus\{i\}}, Q^i)$ according to R^i, that is,

- for every alternative x, the probability assigned to the set $\{y \in A \mid yR^ix\}$ by $F(R^N)$ is at least as large as the probability assigned by $F(R^{N\setminus\{i\}}, Q^i)$

or equivalently,

- $\mathbb{E}_{F(R^N)}(u^i) \geq \mathbb{E}_{F(R^{N\setminus\{i\}}, Q^i)}(u^i)$ for every u^i representing R^i

Observe: this is a strong condition as it implies comparability of these two distributions (stochastic dominance is not a complete relation)

A probabilistic rule is a *random dictatorship* if it picks the top alternative of each agent i with a fixed probability α_i

A probabilistic rule is a *random dictatorship* if it picks the top alternative of each agent i with a fixed probability α_i

A probabilistic rule is *unanimous* if it assigns probability 1 to an alternative x if every agent has x at top

A probabilistic rule is a *random dictatorship* if it picks the top alternative of each agent i with a fixed probability α_i

A probabilistic rule is *unanimous* if it assigns probability 1 to an alternative x if every agent has x at top

Theorem (Gibbard 1977)

Let F be a strategy-proof and unanimous probabilistic rule (defined on all possible preference profiles), and let there be at least three alternatives. Then F is a random dictatorship

A probabilistic rule is a *random dictatorship* if it picks the top alternative of each agent i with a fixed probability α_i

A probabilistic rule is *unanimous* if it assigns probability 1 to an alternative x if every agent has x at top

Theorem (Gibbard 1977)

Let F be a strategy-proof and unanimous probabilistic rule (defined on all possible preference profiles), and let there be at least three alternatives. Then F is a random dictatorship

Is this an 'escape' from the Gibbard-Satterthwaite Theorem?

Not always convincing, e.g. suppose there are 10 agents, each $\alpha_i = 0.1$, 11 alternatives and profile:

1	2	$\cdots$	9	10
x_1	x_2	$\cdots$	x_9	x_{10}
x_{11}	x_{11}	$\cdots$	x_{11}	x_{11}
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

A probabilistic rule is a *random dictatorship* if it picks the top alternative of each agent i with a fixed probability α_i

A probabilistic rule is *unanimous* if it assigns probability 1 to an alternative x if every agent has x at top

Theorem (Gibbard 1977)

Let F be a strategy-proof and unanimous probabilistic rule (defined on all possible preference profiles), and let there be at least three alternatives. Then F is a random dictatorship

Is this an 'escape' from the Gibbard-Satterthwaite Theorem?

Not always convincing, e.g. suppose there are 10 agents, each $\alpha_i = 0.1$, 11 alternatives and profile:

1	2	$\cdots$	9	10
x_1	x_2	$\cdots$	x_9	x_{10}
x_{11}	x_{11}	$\cdots$	x_{11}	x_{11}
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Result still holds under cardinal utilities: Hylland (1980), Dutta et al (2007)

3. Domain restrictions: single-peaked preferences

One of the most studied domain restrictions: single-peaked preferences (Black, 1948)

3. Domain restrictions: single-peaked preferences

One of the most studied domain restrictions: single-peaked preferences (Black, 1948)

Theorem (deterministic case, Moulin 1980)

Let $A = \mathbb{R}$ and let $\mathcal{S}$ be the set of single-peaked preferences on A. Then $F : \mathcal{S}^N \to A$ is strategy-proof, anonymous, and Pareto optimal iff there are $a_1 \leq \ldots \leq a_{n-1} \in \mathbb{R} \cup \{-\infty, \infty\}$ such that

$$F(R^N) = \text{median}\{a_1, \ldots, a_{n-1}, t(R^1), \ldots, t(R^n)\}$$

where $t(R^i)$ is the top alternative (peak) of R^i, for every $R^N \in S^N$

3. Domain restrictions: single-peaked preferences

One of the most studied domain restrictions: single-peaked preferences (Black, 1948)

Theorem (deterministic case, Moulin 1980)

Let $A = \mathbb{R}$ and let $\mathcal{S}$ be the set of single-peaked preferences on A. Then $F : \mathcal{S}^N \to A$ is strategy-proof, anonymous, and Pareto optimal iff there are $a_1 \leq \ldots \leq a_{n-1} \in \mathbb{R} \cup \{-\infty, \infty\}$ such that

$$F(R^N) = \text{median}\{a_1, \ldots, a_{n-1}, t(R^1), \ldots, t(R^n)\}$$

where $t(R^i)$ is the top alternative (peak) of R^i, for every $R^N \in S^N$

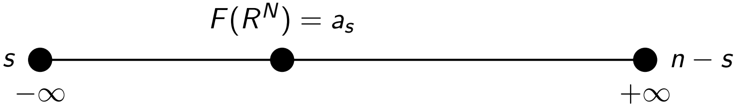

4. Probabilistic rules and single-peakedness

Assume $A \subseteq \mathbb{R}$ is finite, $A = \{x_1, \ldots, x_m\}$ with $x_1 < \ldots < x_m$. Let $\mathcal{S}$ be the set of single-peaked preferences on A

Let the probabilistic rule F on $\mathcal{S}^N$ be strategy-proof, anonymous, and unanimous

4. Probabilistic rules and single-peakedness

Assume $A \subseteq \mathbb{R}$ is finite, $A = \{x_1, \ldots, x_m\}$ with $x_1 < \ldots < x_m$. Let $\mathcal{S}$ be the set of single-peaked preferences on A

Let the probabilistic rule F on $\mathcal{S}^N$ be strategy-proof, anonymous, and unanimous

Any coalition S with s voters determines a probability distribution on A, namely $F(R^N)$ with S having peaks on x_1 and $N \setminus S$ having peaks on x_m

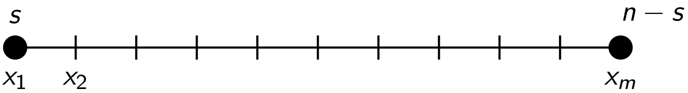

Generates probability distribution D_s. If s decreases then probability shifts to the right

How, in turn, is the rule determined by these distributions D_s, $1 \leq s \leq n-1$?

How, in turn, is the rule determined by these distributions D_s, $1 \leq s \leq n-1$?
Assume $p^1 \leq \ldots \leq p^n$ are the reported peaks

How, in turn, is the rule determined by these distributions D_s, $1 \leq s \leq n-1$?

Assume $p^1 \leq \ldots \leq p^n$ are the reported peaks

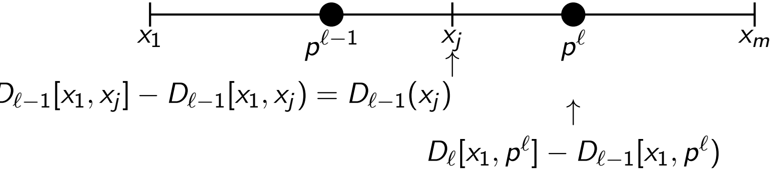

How, in turn, is the rule determined by these distributions D_s, $1 \leq s \leq n-1$?

Assume $p^1 \leq \ldots \leq p^n$ are the reported peaks

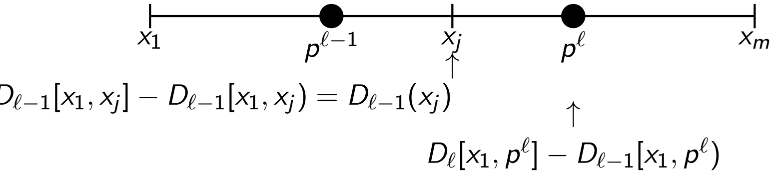

Theorem (EPS 2002)

All strategy-proof, anonymous, and unanimous probabilistic rules are determined by such fixed distributions (on $\mathbb{R}$ or a subset of $\mathbb{R}$)

How, in turn, is the rule determined by these distributions D_s, $1 \leq s \leq n-1$?

Assume $p^1 \leq \ldots \leq p^n$ are the reported peaks

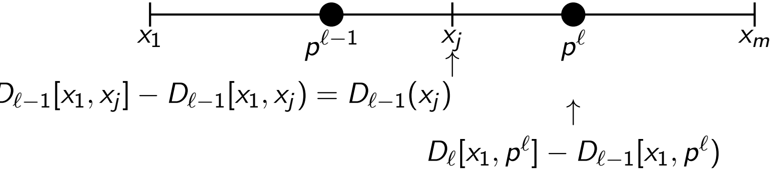

Theorem (EPS 2002)

All strategy-proof, anonymous, and unanimous probabilistic rules are determined by such fixed distributions (on $\mathbb{R}$ or a subset of $\mathbb{R}$)

Further results

- For A finite, every such rule is a convex combination of deterministic rules with the same properties (PRSS 2014)
- Under the analogous conditions, for $A \subseteq \mathbb{R}^n$ with $n > 1$ we have random dictatorship (DPS 2002)

5. Probabilistic rules for single-peaked preferences on graphs

(Joint work with Souvik Roy and Soumyarup Sadhukhan, 2019)

5. Probabilistic rules for single-peaked preferences on graphs

(Joint work with Souvik Roy and Soumyarup Sadhukhan, 2019)

We now assume that the alternatives are the nodes in a graph, and that (single-peaked) preferences are related to this graph

5. Probabilistic rules for single-peaked preferences on graphs

(Joint work with Souvik Roy and Soumyarup Sadhukhan, 2019)

We now assume that the alternatives are the nodes in a graph, and that (single-peaked) preferences are related to this graph
There are several possible interpretations

- Literal interpretation: the graph is a network of roads or railway tracks, and a public facility is to be located at some node in this network (special case: line graph, as considered before)

5. Probabilistic rules for single-peaked preferences on graphs

(Joint work with Souvik Roy and Soumyarup Sadhukhan, 2019)

We now assume that the alternatives are the nodes in a graph, and that (single-peaked) preferences are related to this graph
There are several possible interpretations

- Literal interpretation: the graph is a network of roads or railway tracks, and a public facility is to be located at some node in this network (special case: line graph, as considered before)
- A graph as a general means to express preferences. Think of some of the nodes representing meals of different spiciness, and others representing the amounts of meat. Preferences can be single-peaked with respect to each category separately, but not between categories

An example

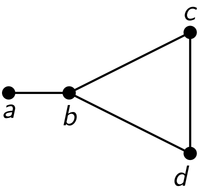

$N = \{1, 2, 3\}$

An example

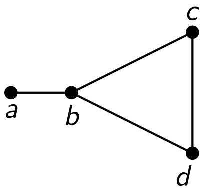

$N = \{1, 2, 3\}$

- Each preference is single-peaked with respect to some spanning tree. E.g. if the peak is at b and edge $\{b, d\}$ is left out, then b is preferred to a, b is preferred to c and c to d

An example

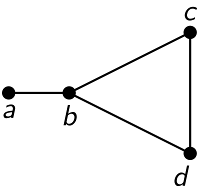

$N = \{1, 2, 3\}$

- Each preference is single-peaked with respect to some spanning tree. E.g. if the peak is at b and edge $\{b, d\}$ is left out, then b is preferred to a, b is preferred to c and c to d
- Each agent has weight $\frac{1}{3}$. If all agents have peaks in $\{b, c, d\}$ then random dictatorship

An example

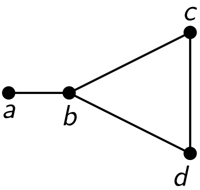

$N = \{1, 2, 3\}$

- Each preference is single-peaked with respect to some spanning tree. E.g. if the peak is at b and edge $\{b, d\}$ is left out, then b is preferred to a, b is preferred to c and c to d
- Each agent has weight $\frac{1}{3}$. If all agents have peaks in $\{b, c, d\}$ then random dictatorship
- If there are one or two agents on a their weight is split evenly between a and b

An example

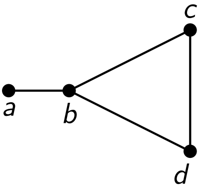

$N = \{1, 2, 3\}$

- Each preference is single-peaked with respect to some spanning tree. E.g. if the peak is at b and edge $\{b, d\}$ is left out, then b is preferred to a, b is preferred to c and c to d
- Each agent has weight $\frac{1}{3}$. If all agents have peaks in $\{b, c, d\}$ then random dictatorship
- If there are one or two agents on a their weight is split evenly between a and b
- For instance, if 1 is at a, 2 at b, and 3 at d, then a gets probability $\frac{1}{2} \cdot \frac{1}{3} = \frac{1}{6}$, b gets probability $(\frac{1}{2} \cdot \frac{1}{3}) + \frac{1}{3} = \frac{1}{2}$ and d gets probability $\frac{1}{3}$.

Formally

- $N = \{1, \ldots, n\}$ is the set of agents

Formally

- $N = \{1, \ldots, n\}$ is the set of agents
- $G = (A, E)$ is a connected graph, where A is the finite set of alternatives (nodes) and $E \subseteq \{\{x, y\} : x, y \in A\}$ is the set of edges

Formally

- $N = \{1, \ldots, n\}$ is the set of agents
- $G = (A, E)$ is a connected graph, where A is the finite set of alternatives (nodes) and $E \subseteq \{\{x, y\} : x, y \in A\}$ is the set of edges
- A *path* in G between $x, y \in A$ is a sequence $x = x^1, \ldots, x^k = y$ of distinct alternatives such that $\{x^j, x^{j+1}\} \in E$ for every $j = 1, \ldots, k - 1$

Formally

- $N = \{1, \ldots, n\}$ is the set of agents
- $G = (A, E)$ is a connected graph, where A is the finite set of alternatives (nodes) and $E \subseteq \{\{x, y\} : x, y \in A\}$ is the set of edges
- A *path* in G between $x, y \in A$ is a sequence $x = x^1, \ldots, x^k = y$ of distinct alternatives such that $\{x^j, x^{j+1}\} \in E$ for every $j = 1, \ldots, k-1$
- A *spanning tree* $T = (A, E_T)$, where $E_T \subseteq E$, is a graph such that between every pair of alternatives $x, y \in A$ there is a unique path, denoted $[x, y]$

- A preference R is a linear ordering on A. Its top alternative is denoted by $t(R)$

- A preference R is a linear ordering on A. Its top alternative is denoted by $t(R)$
- A preference R is *single-peaked* if there is a spanning tree such that for all distinct $x, y \in A$ with $y \neq t(R)$,

$$x \in [t(R), y] \implies xRy$$

- A preference R is a linear ordering on A. Its top alternative is denoted by $t(R)$
- A preference R is *single-peaked* if there is a spanning tree such that for all distinct $x, y \in A$ with $y \neq t(R)$,

$$x \in [t(R), y] \implies xRy$$

- A *probabilistic rule* φ assigns to each profile of single-peaked preferences probability distribution over A

A probabilistic rule is *unanimous* it it assigns probability 1 to an alternative if that alternative is the peak of every agent

A probabilistic rule is *unanimous* it it assigns probability 1 to an alternative if that alternative is the peak of every agent

We will characterize all strategy-proof and unanimous probabilistic rules for connected graphs $G = (A, E)$

A probabilistic rule is *unanimous* it it assigns probability 1 to an alternative if that alternative is the peak of every agent

We will characterize all strategy-proof and unanimous probabilistic rules for connected graphs $G = (A, E)$

A *leaf* of G is an alternative $a \in A$ that has degree 1, i.e., there is a unique $b \in A$ with $\{a, b\} \in E$

A probabilistic rule is *unanimous* it it assigns probability 1 to an alternative if that alternative is the peak of every agent

We will characterize all strategy-proof and unanimous probabilistic rules for connected graphs $G = (A, E)$

A *leaf* of G is an alternative $a \in A$ that has degree 1, i.e., there is a unique $b \in A$ with $\{a, b\} \in E$

We consider to 'extreme' subcases namely

- $G = (A, E)$ is a tree
- $G = (A, E)$ has no leafs

The general case will follow from a 'combination' of these two cases

6. Trees

Until further notice $G = (A, E)$ is a tree

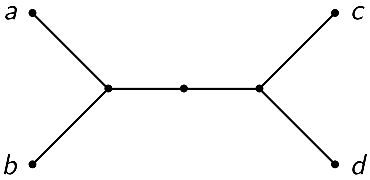

6. Trees

Until further notice $G = (A, E)$ is a tree

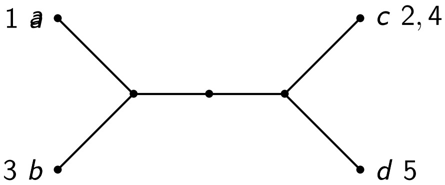

An *leaf assignment* μ assigns every agent to a leaf of the tree, and for every μ there is a probability distribution β_μ over A, such that

6. Trees

Until further notice $G = (A, E)$ is a tree

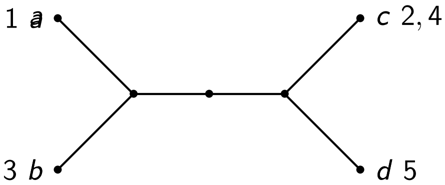

An *leaf assignment* μ assigns every agent to a leaf of the tree, and for every μ there is a probability distribution β_μ over A, such that

- If all agents are at the same leaf then this leaf gets probability 1

6. Trees

Until further notice $G = (A, E)$ is a tree

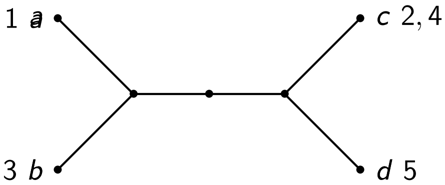

An *leaf assignment* μ assigns every agent to a leaf of the tree, and for every μ there is a probability distribution β_μ over A, such that

- If all agents are at the same leaf then this leaf gets probability 1
- If an agent moves to a different leaf then probability shifts along the path from the former to the new leaf; the probabilities off this path stay the same

Such a *monotonic collection* $B = (\beta_\mu)_\mu$ determines a probabilistic rule φ^B for single-peaked preference profiles over $G = (A, E)$

Such a *monotonic collection* $B = (\beta_\mu)_\mu$ determines a probabilistic rule φ^B for single-peaked preference profiles over $G = (A, E)$

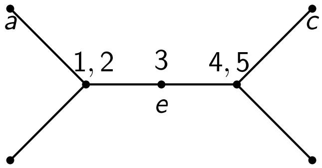

What is $\varphi_e^B(R_N)$?

Such a *monotonic collection* $B = (\beta_\mu)_\mu$ determines a probabilistic rule φ^B for single-peaked preference profiles over $G = (A, E)$

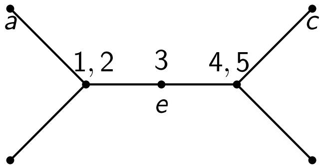

What is $\varphi^B_e(R_N)$?

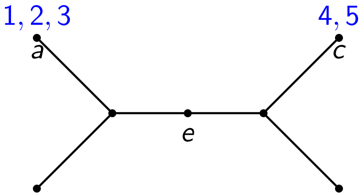

Such a *monotonic collection* $B = (\beta_\mu)_\mu$ determines a probabilistic rule φ^B for single-peaked preference profiles over $G = (A, E)$

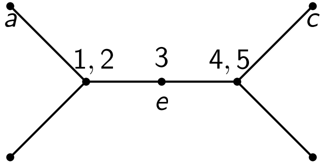

What is $\varphi_e^B(R_N)$?

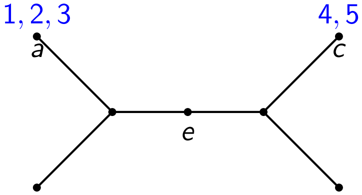

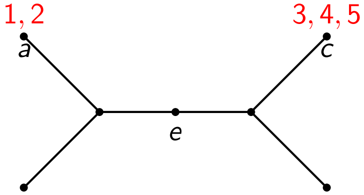

Such a *monotonic collection* $B = (\beta_\mu)_\mu$ determines a probabilistic rule φ^B for single-peaked preference profiles over $G = (A, E)$

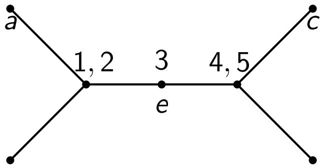

What is $\varphi_e^B(R_N)$?

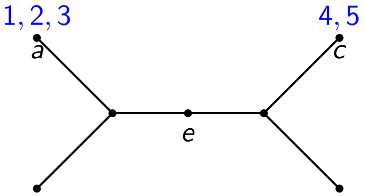

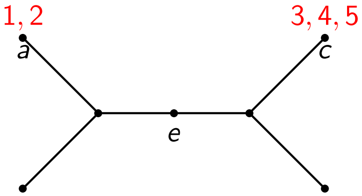

Now $\varphi_e^B(R_N) = \beta_{\mu}([e, a]) - \beta_{\hat{\mu}}((e, a])$

Theorem for trees

Let $G = (A, E)$ be a tree and let φ be a probabilistic rule defined for all single-peaked preference profiles on G. Then φ is unanimous and strategy-proof if and only if there is a monotonic collection $B = (\beta_\mu)_\mu$ such that $\varphi = \varphi^B$

Theorem for trees

Let $G = (A, E)$ be a tree and let φ be a probabilistic rule defined for all single-peaked preference profiles on G. Then φ is unanimous and strategy-proof if and only if there is a monotonic collection $B = (\beta_\mu)_\mu$ such that $\varphi = \varphi^B$

[In the proof we show that unanimity and strategy-proofness imply *peaks-onliness*: φ depends only on the peaks of the preferences. We use a result of Chatterji and Zeng (2018). In turn, this implies that φ is *uncompromising*: if an agent shifts its peak to another alternative then all probabilities off the path between the the old and the new peak stay the same (Border and Jordan, 1983)]

7. Leafless graphs

We first consider 2-connected graphs: G is 2-connected if for each pair of distinct alternatives a and b there is a cycle containing a and b

Clearly, a 2-connected graph is leafless

7. Leafless graphs

We first consider 2-connected graphs: G is 2-connected if for each pair of distinct alternatives a and b there is a cycle containing a and b

Clearly, a 2-connected graph is leafless

A probabilistic rule φ is a *random dictatorship* if there are $\alpha_1, \dots, \alpha_n \geq 0$ with $\sum_{i \in N} \alpha_i = 1$, such that for every preference profile R_N and every $a \in A$ we have:

$$\varphi_a(R_N) = \sum_{i \in N : t(R_i) = a} \alpha_i$$

7. Leafless graphs

We first consider 2-connected graphs: G is 2-connected if for each pair of distinct alternatives a and b there is a cycle containing a and b

Clearly, a 2-connected graph is leafless

A probabilistic rule φ is a *random dictatorship* if there are $\alpha_1, \ldots, \alpha_n \geq 0$ with $\sum_{i \in N} \alpha_i = 1$, such that for every preference profile R_N and every $a \in A$ we have:

$$\varphi_a(R_N) = \sum_{i \in N : t(R_i) = a} \alpha_i$$

Theorem for 2-connected graphs

Let $G = (A, E)$ be a 2-connected graph and let φ be a probabilistic rule defined for all single-peaked preference profiles on G. Then φ is unanimous and strategy-proof if and only if it is a random dictatorship

Any (hence also a leafless) graph can be decomposed into maximal 2-connected subgraphs, giving rise to a so-called *block-tree* (Menger, 1927)

Any (hence also a leafless) graph can be decomposed into maximal 2-connected subgraphs, giving rise to a so-called *block-tree* (Menger, 1927)

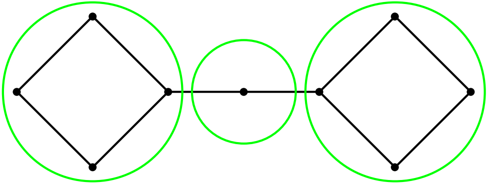

Any (hence also a leafless) graph can be decomposed into maximal 2-connected subgraphs, giving rise to a so-called *block-tree* (Menger, 1927)

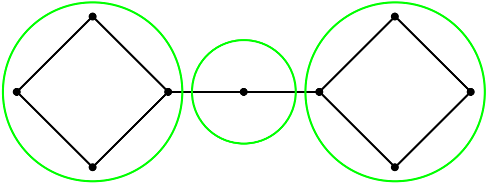

A unanimous and strategy-proof probabilistic rule generates random dictatorships on these maximal 2-connected subgraphs

Any (hence also a leafless) graph can be decomposed into maximal 2-connected subgraphs, giving rise to a so-called *block-tree* (Menger, 1927)

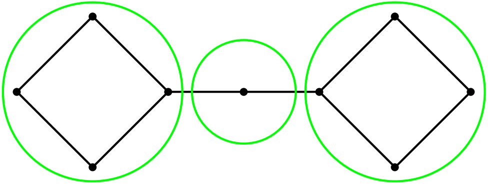

A unanimous and strategy-proof probabilistic rule generates random dictatorships on these maximal 2-connected subgraphs

We show that these are the same random dictatorships and, moreover, they spread out over the rest of the graph

Any (hence also a leafless) graph can be decomposed into maximal 2-connected subgraphs, giving rise to a so-called *block-tree* (Menger, 1927)

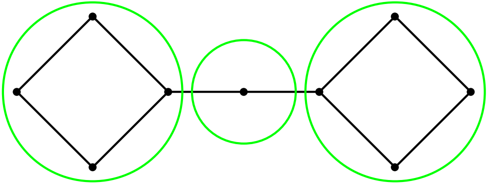

A unanimous and strategy-proof probabilistic rule generates random dictatorships on these maximal 2-connected subgraphs

We show that these are the same random dictatorships and, moreover, they spread out over the rest of the graph

Theorem for leafless graphs

Let $G = (A, E)$ be a leafless graph and let φ be a probabilistic rule defined for all single-peaked preference profiles on G. Then φ is unanimous and strategy-proof if and only if it is a random dictatorship

A random dictatorship φ with weights $\alpha_1, \ldots, \alpha_n$ is a special case of rule φ^B, for a specific collection $B = (\beta_\mu)_\mu$, as follows

A random dictatorship φ with weights $\alpha_1, \ldots, \alpha_n$ is a special case of rule φ^B, for a specific collection $B = (\beta_\mu)_\mu$, as follows

Take an arbitrary spanning tree T of G and for any leaf assignment μ and leaf a let $\beta_\mu(a) = \sum_{i \in N : \mu(i) = a} \alpha_i$. Then the random dictatorship φ coincides with φ^B for this collection $B = (\beta_\mu)_\mu$

A random dictatorship φ with weights $\alpha_1, \ldots, \alpha_n$ is a special case of rule φ^B, for a specific collection $B = (\beta_\mu)_\mu$, as follows

Take an arbitrary spanning tree T of G and for any leaf assignment μ and leaf a let $\beta_\mu(a) = \sum_{i \in N : \mu(i) = a} \alpha_i$. Then the random dictatorship φ coincides with φ^B for this collection $B = (\beta_\mu)_\mu$

Say $N = \{1, 2, 3\}$, $(\alpha_1, \alpha_2, \alpha_3) = (\frac{1}{6}, \frac{1}{3}, \frac{1}{2})$

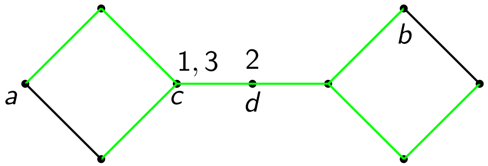

A random dictatorship φ with weights $\alpha_1, \ldots, \alpha_n$ is a special case of rule φ^B, for a specific collection $B = (\beta_\mu)_\mu$, as follows

Take an arbitrary spanning tree T of G and for any leaf assignment μ and leaf a let $\beta_\mu(a) = \sum_{i \in N: \mu(i)=a} \alpha_i$. Then the random dictatorship φ coincides with φ^B for this collection $B = (\beta_\mu)_\mu$

Say $N = \{1, 2, 3\}$, $(\alpha_1, \alpha_2, \alpha_3) = (\frac{1}{6}, \frac{1}{3}, \frac{1}{2})$

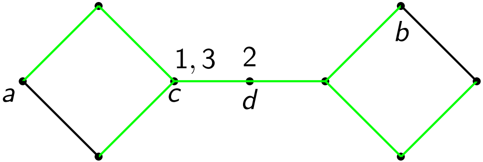

Hence $\varphi_c(R_N) = \frac{1}{6} + \frac{1}{2} = \frac{2}{3}$, $\varphi_d(R_N) = \frac{1}{3}$

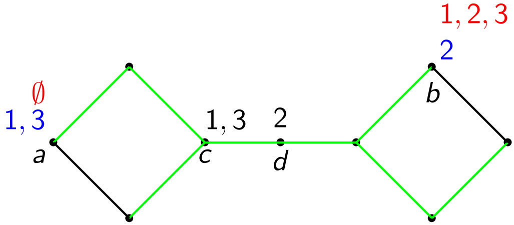

Now $\varphi_c^B(P_N) = \beta_{\mu}([c,a]) - \beta_{\hat{\mu}}((c,a]) = \frac{1}{6} + \frac{1}{2} - 0 = \frac{2}{3} = \varphi_c(P_N)$

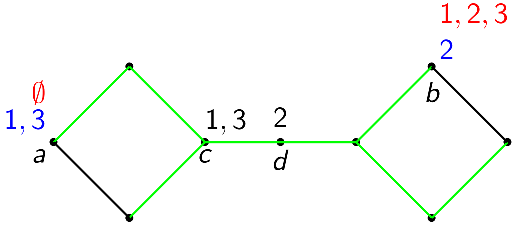

Now $\varphi^B_c(P_N) = \beta_{\mu}([c,a]) - \beta_{\hat{\mu}}((c,a]) = \frac{1}{6} + \frac{1}{2} - 0 = \frac{2}{3} = \varphi_c(P_N)$

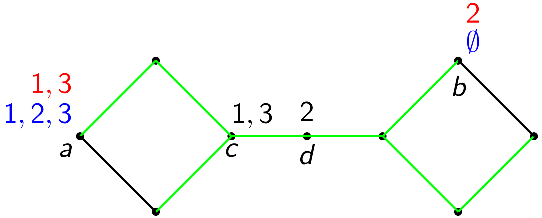

And $\varphi^B_d(P_N) = \beta_{\mu}([d,a]) - \beta_{\hat{\mu}}((d,a]) = 1 - \frac{1}{6} - \frac{1}{2} = \frac{1}{3} = \varphi_d(P_N)$

8. General connected graphs

Finally, $G = (A, E)$ is an arbitrary connected graph

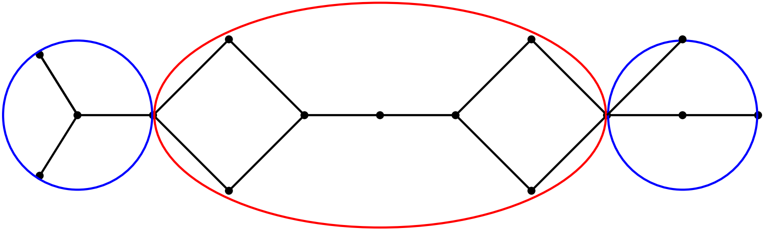

The subgraph in the red ellipse is the maximal leafless subgraph

The parts in the blue circles are *branches* (trees)

Every connected graph can be split up this way

We take an arbitrary spanning tree T of G

We take an arbitrary spanning tree T of G

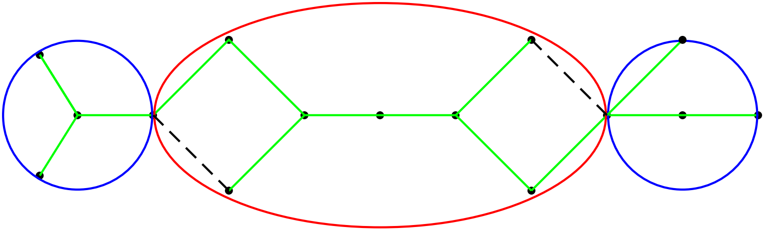

We take an arbitrary spanning tree T of G

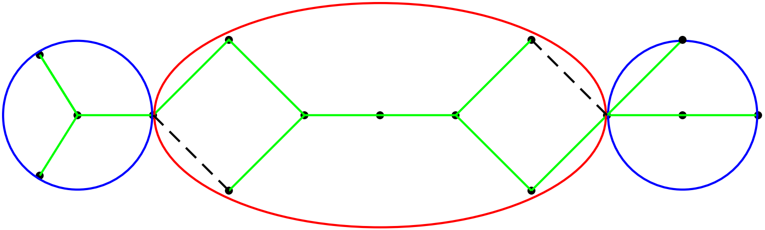

We consider again monotonic collections $B = (\beta_\mu)_\mu$, plus random dictatorship weights $\alpha_1, \ldots, \alpha_n$, with **additional conditions**

We take an arbitrary spanning tree T of G

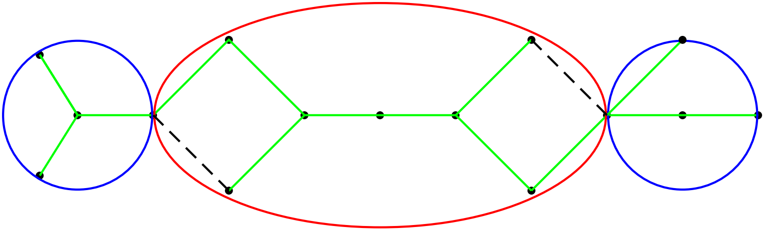

We consider again monotonic collections $B = (\beta_\mu)_\mu$, plus random dictatorship weights $\alpha_1, \ldots, \alpha_n$, with **additional conditions**

(a) For any β_μ, the sum of the probabilities on a branch equals the sum of the α_i of agents i assigned to leafs of that branch

We take an arbitrary spanning tree T of G

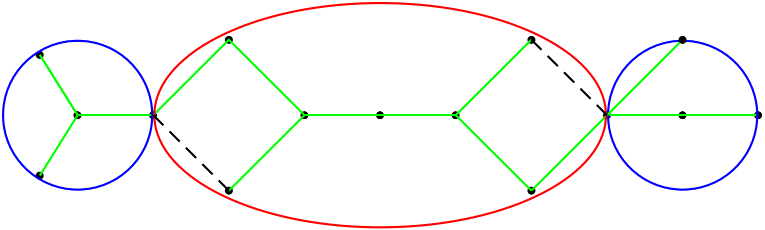

We consider again monotonic collections $B = (\beta_\mu)_\mu$, plus random dictatorship weights $\alpha_1, \ldots, \alpha_n$, with **additional conditions**

(a) For any β_μ, the sum of the probabilities on a branch equals the sum of the α_i of agents i assigned to leafs of that branch

(b) Any leaf of T in the leafless red subgraph gets sum of the α_i of agents assigned to that leaf

Theorem for general connected graphs

Let $G = (A, E)$ be a connected graph and let φ be a probabilistic rule defined for all single-peaked preference profiles on G. Fix an arbitrary spanning tree T of G. Then φ is unanimous and strategy-proof if and only if there are weights $\alpha_1, \ldots, \alpha_n$ and a monotonic collection $B = (\beta_\mu)_\mu$ satisfying (a) and (b) above such that $\varphi = \varphi^B$

An example

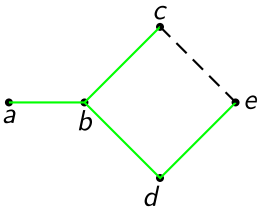

Let $N = \{1, 2, 3\}$, $\alpha_1 = \alpha_2 = \alpha_3 = \frac{1}{3}$, and let each β_μ assign equal probabilities to a and b if the number of agents assigned to a is below 3

An example

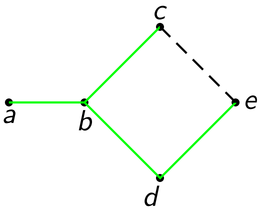

Let $N = \{1, 2, 3\}$, $\alpha_1 = \alpha_2 = \alpha_3 = \frac{1}{3}$, and let each β_μ assign equal probabilities to a and b if the number of agents assigned to a is below 3

Then, for instance, if R_N is such that $t(R_1) = a$, $t(R_2) = c$, and $t(R_3) = d$, then φ assigns $(\frac{1}{6}, \frac{1}{6}, \frac{1}{3}, \frac{1}{3}, 0)$ to (a, b, c, d, e)

9. Concluding remarks

We have characterized the set of all unanimous and strategy-proof rules when preferences are single-peaked with respect to a connected arbitrary graph

9. Concluding remarks

We have characterized the set of all unanimous and strategy-proof rules when preferences are single-peaked with respect to a connected arbitrary graph

About deterministic versus probabilistic rules:

9. Concluding remarks

We have characterized the set of all unanimous and strategy-proof rules when preferences are single-peaked with respect to a connected arbitrary graph

About deterministic versus probabilistic rules:

- Deterministic rules are a special case of probabilistic rules

9. Concluding remarks

We have characterized the set of all unanimous and strategy-proof rules when preferences are single-peaked with respect to a connected arbitrary graph

About deterministic versus probabilistic rules:

- Deterministic rules are a special case of probabilistic rules
- Clearly, random dictatorship rules are convex combinations of deterministic rules

9. Concluding remarks

We have characterized the set of all unanimous and strategy-proof rules when preferences are single-peaked with respect to a connected arbitrary graph

About deterministic versus probabilistic rules:

- Deterministic rules are a special case of probabilistic rules
- Clearly, random dictatorship rules are convex combinations of deterministic rules
- This result extends to probabilistic rules for line graphs (PRSS 2014)
- It no longer holds for other trees or for general connected graphs

The end