ICML | 2022

Exact Learning Preference Structure: Single-Peaked Preferences and Beyond

Sonja Kraiczy, Edith Elkind
University of Oxford

In a country far, far away…

In a country far, far away…

- **Estonian Parliament 2021:**
  - Sotsiaaldemokraatlik Erakond
  - Eesti Reformierakond
  - Isamaa
  - Eesti Konservatiivne Rahvaerakond
  - Eesti Keskerakond

In a country far, far away…

- **Estonian Parliament 2021:**
  - Sotsiaaldemokraatlik Erakond
  - Eesti Reformierakond
  - Isamaa
  - Eesti Konservatiivne Rahvaerakond
  - Eesti Keskerakond

Can you order these parties from left to right on the political spectrum?

In a country far, far away…

- **Estonian Parliament 2021:**
  - Sotsiaaldemokraatlik Erakond
  - Eesti Reformierakond
  - Isamaa
  - Eesti Konservatiivne Rahvaerakond
  - Eesti Keskerakond

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

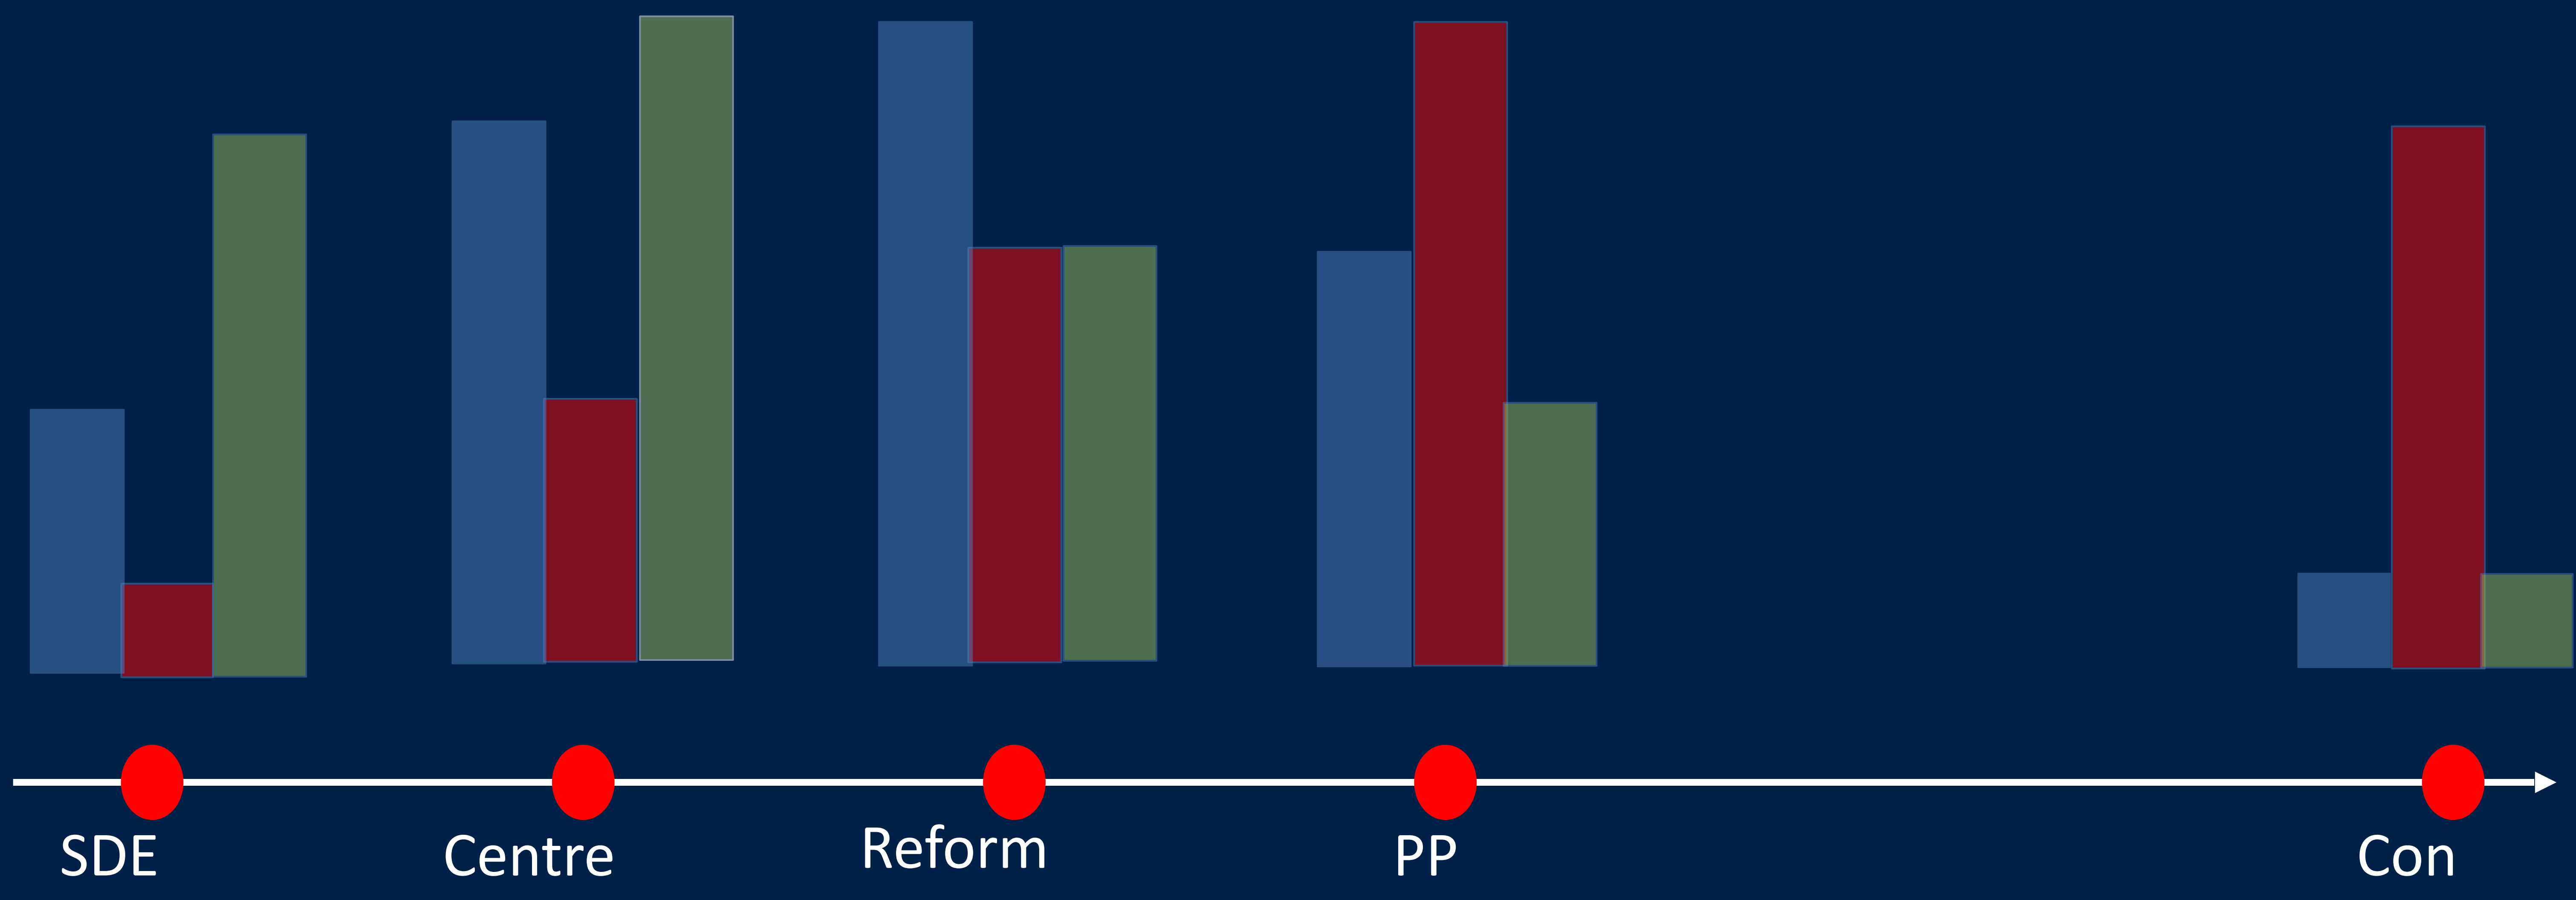

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

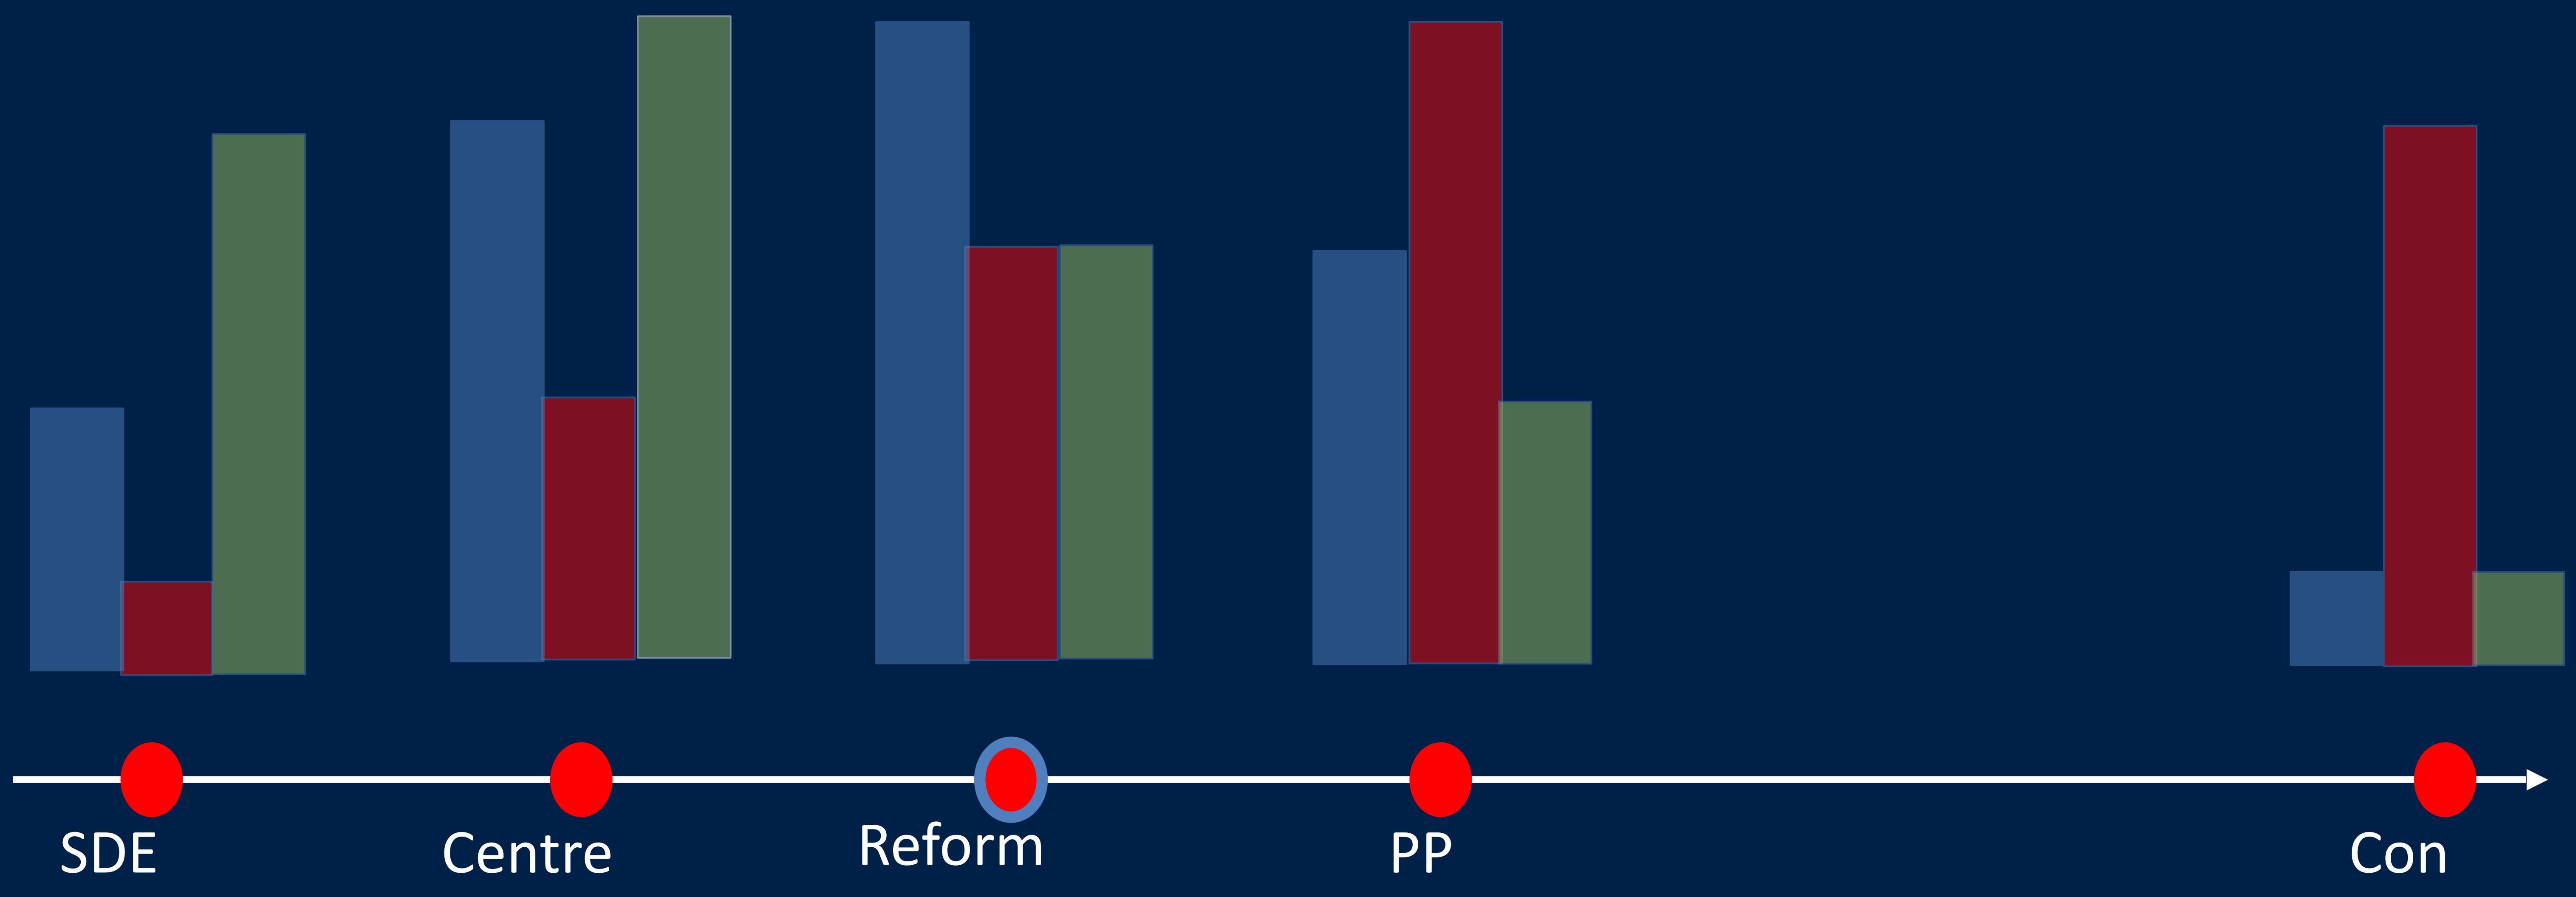

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

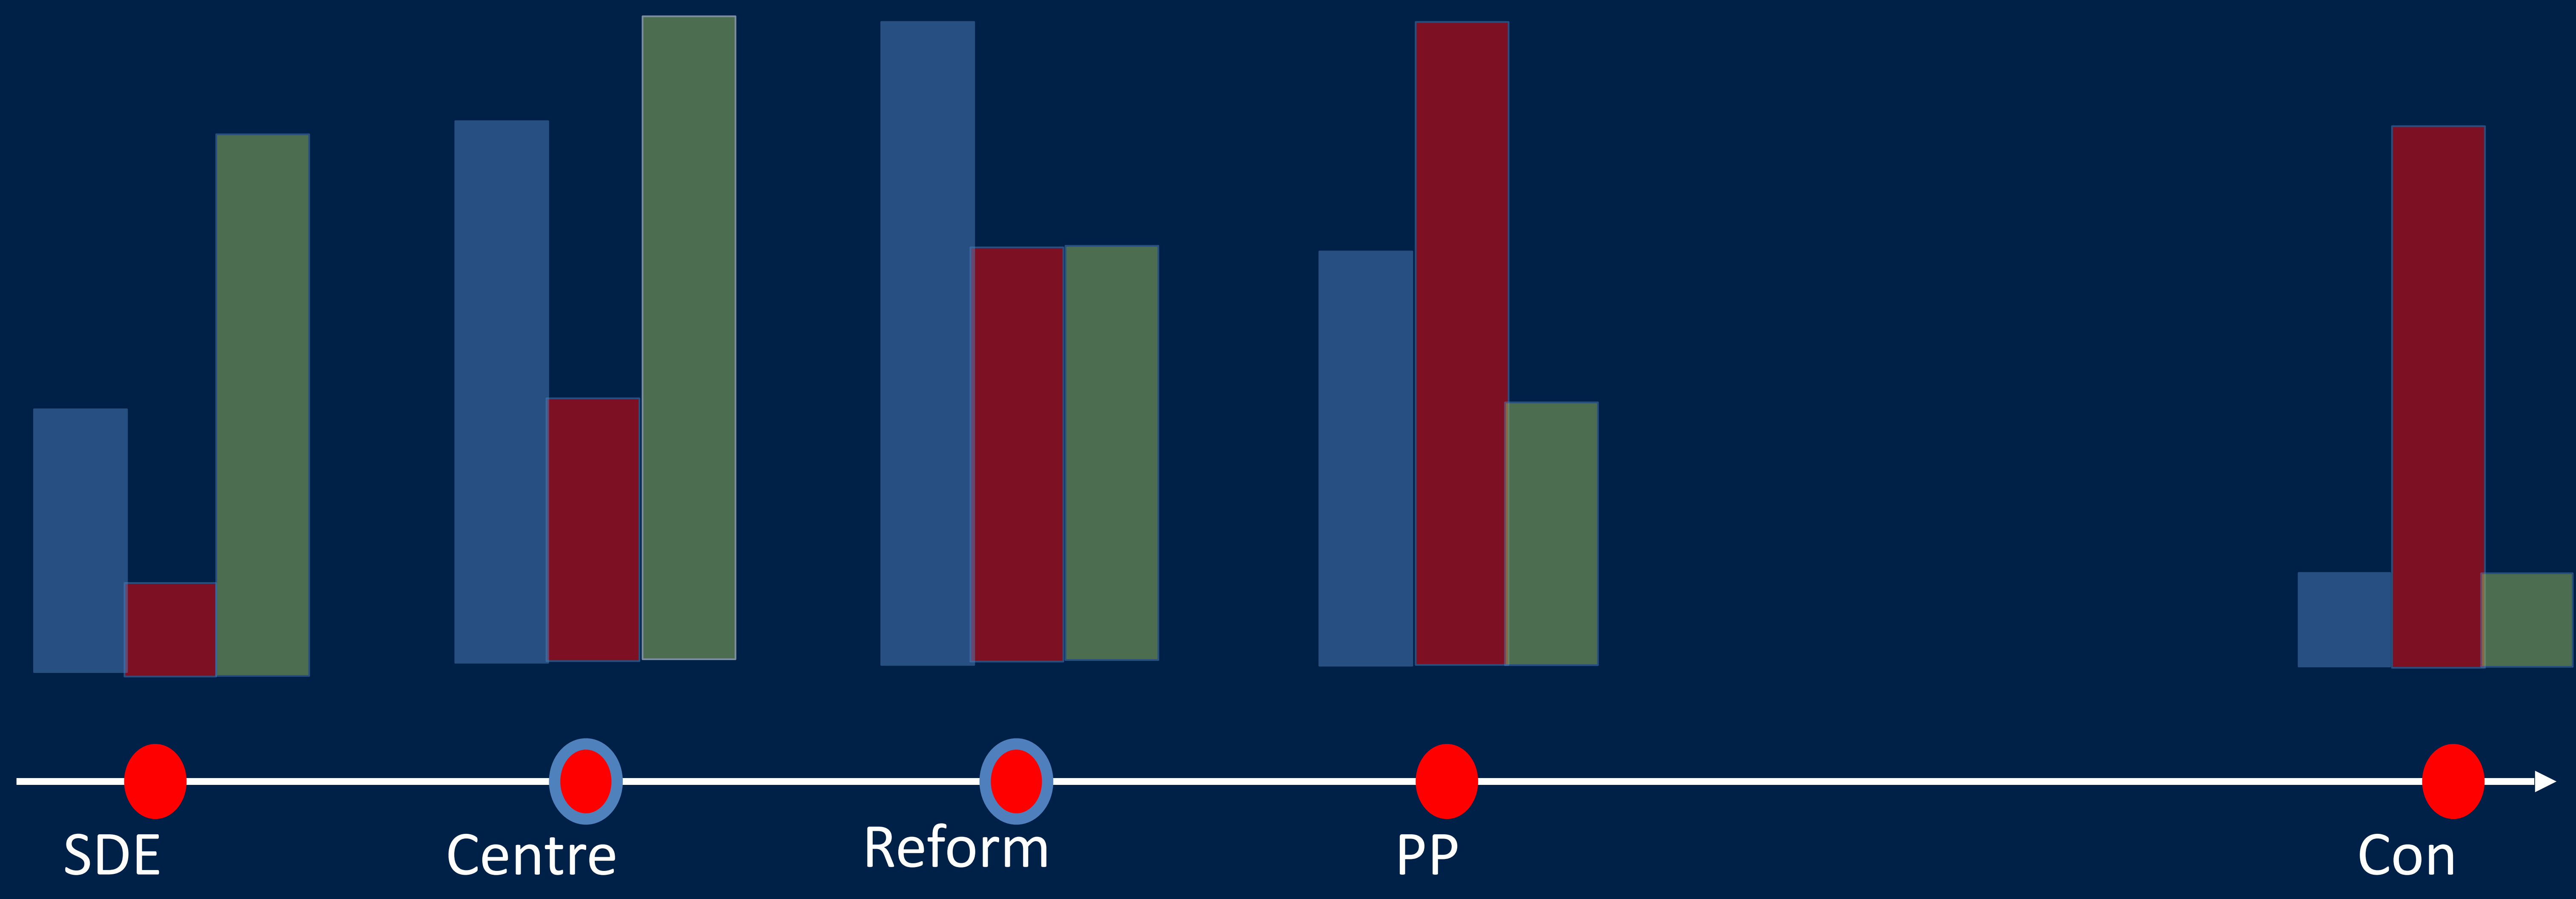

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

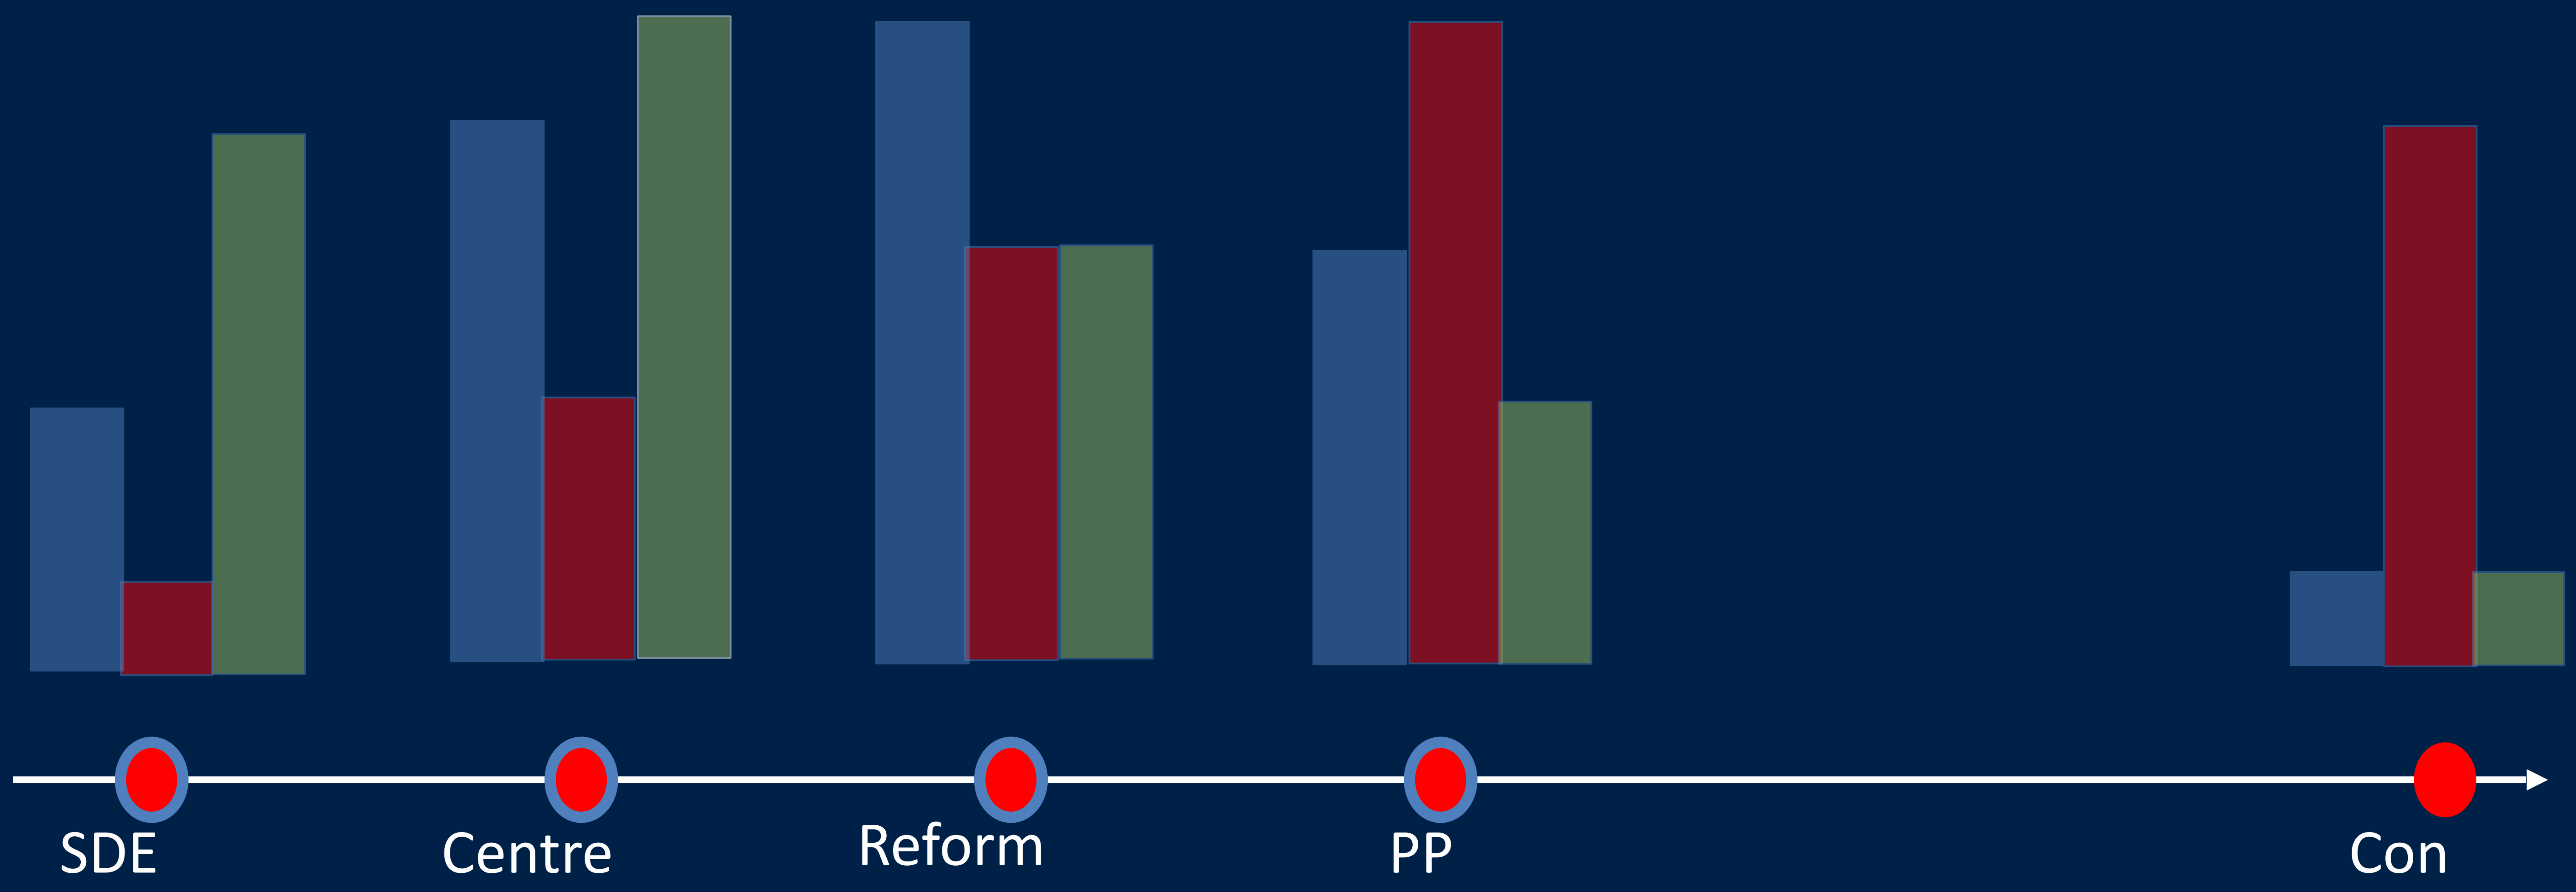

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

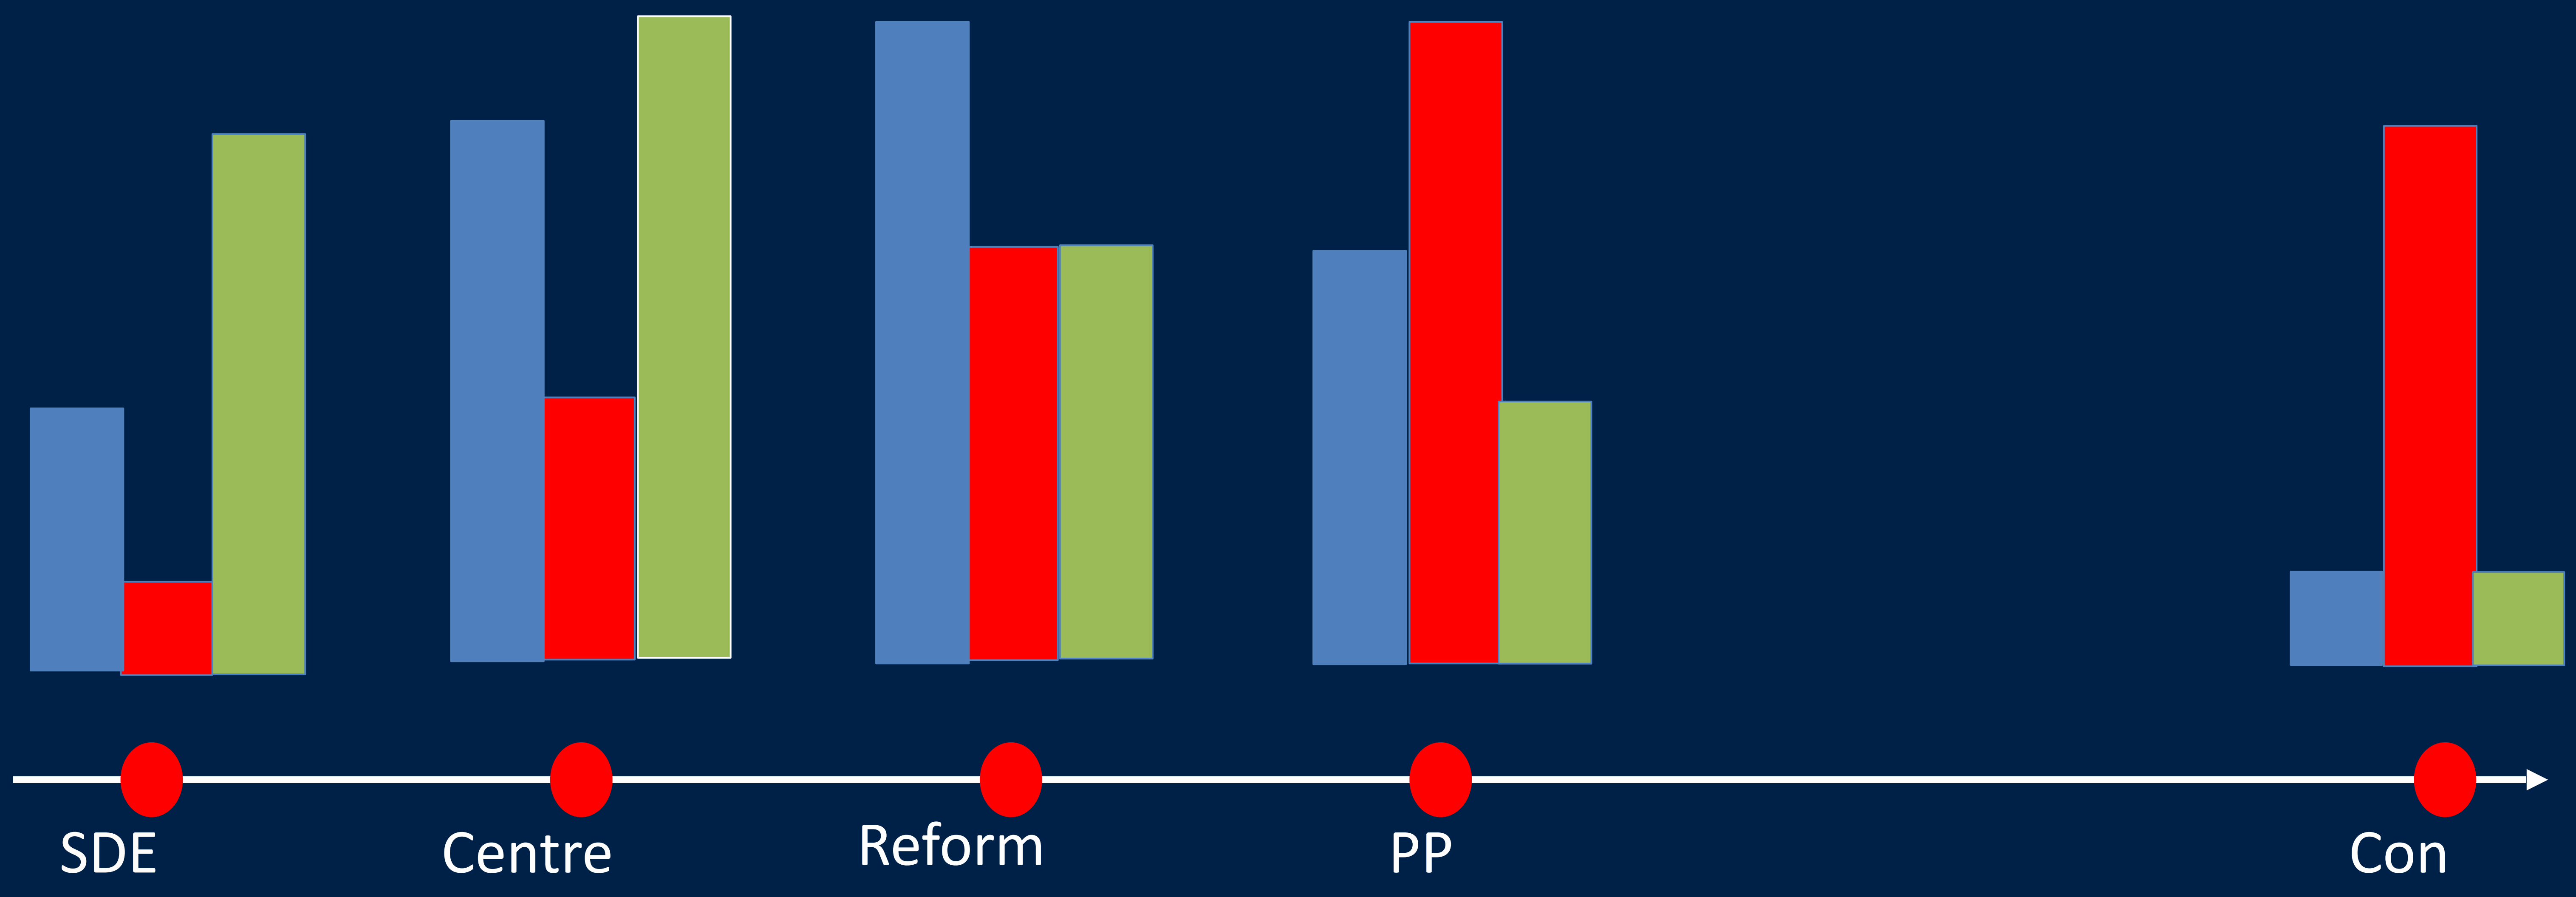

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

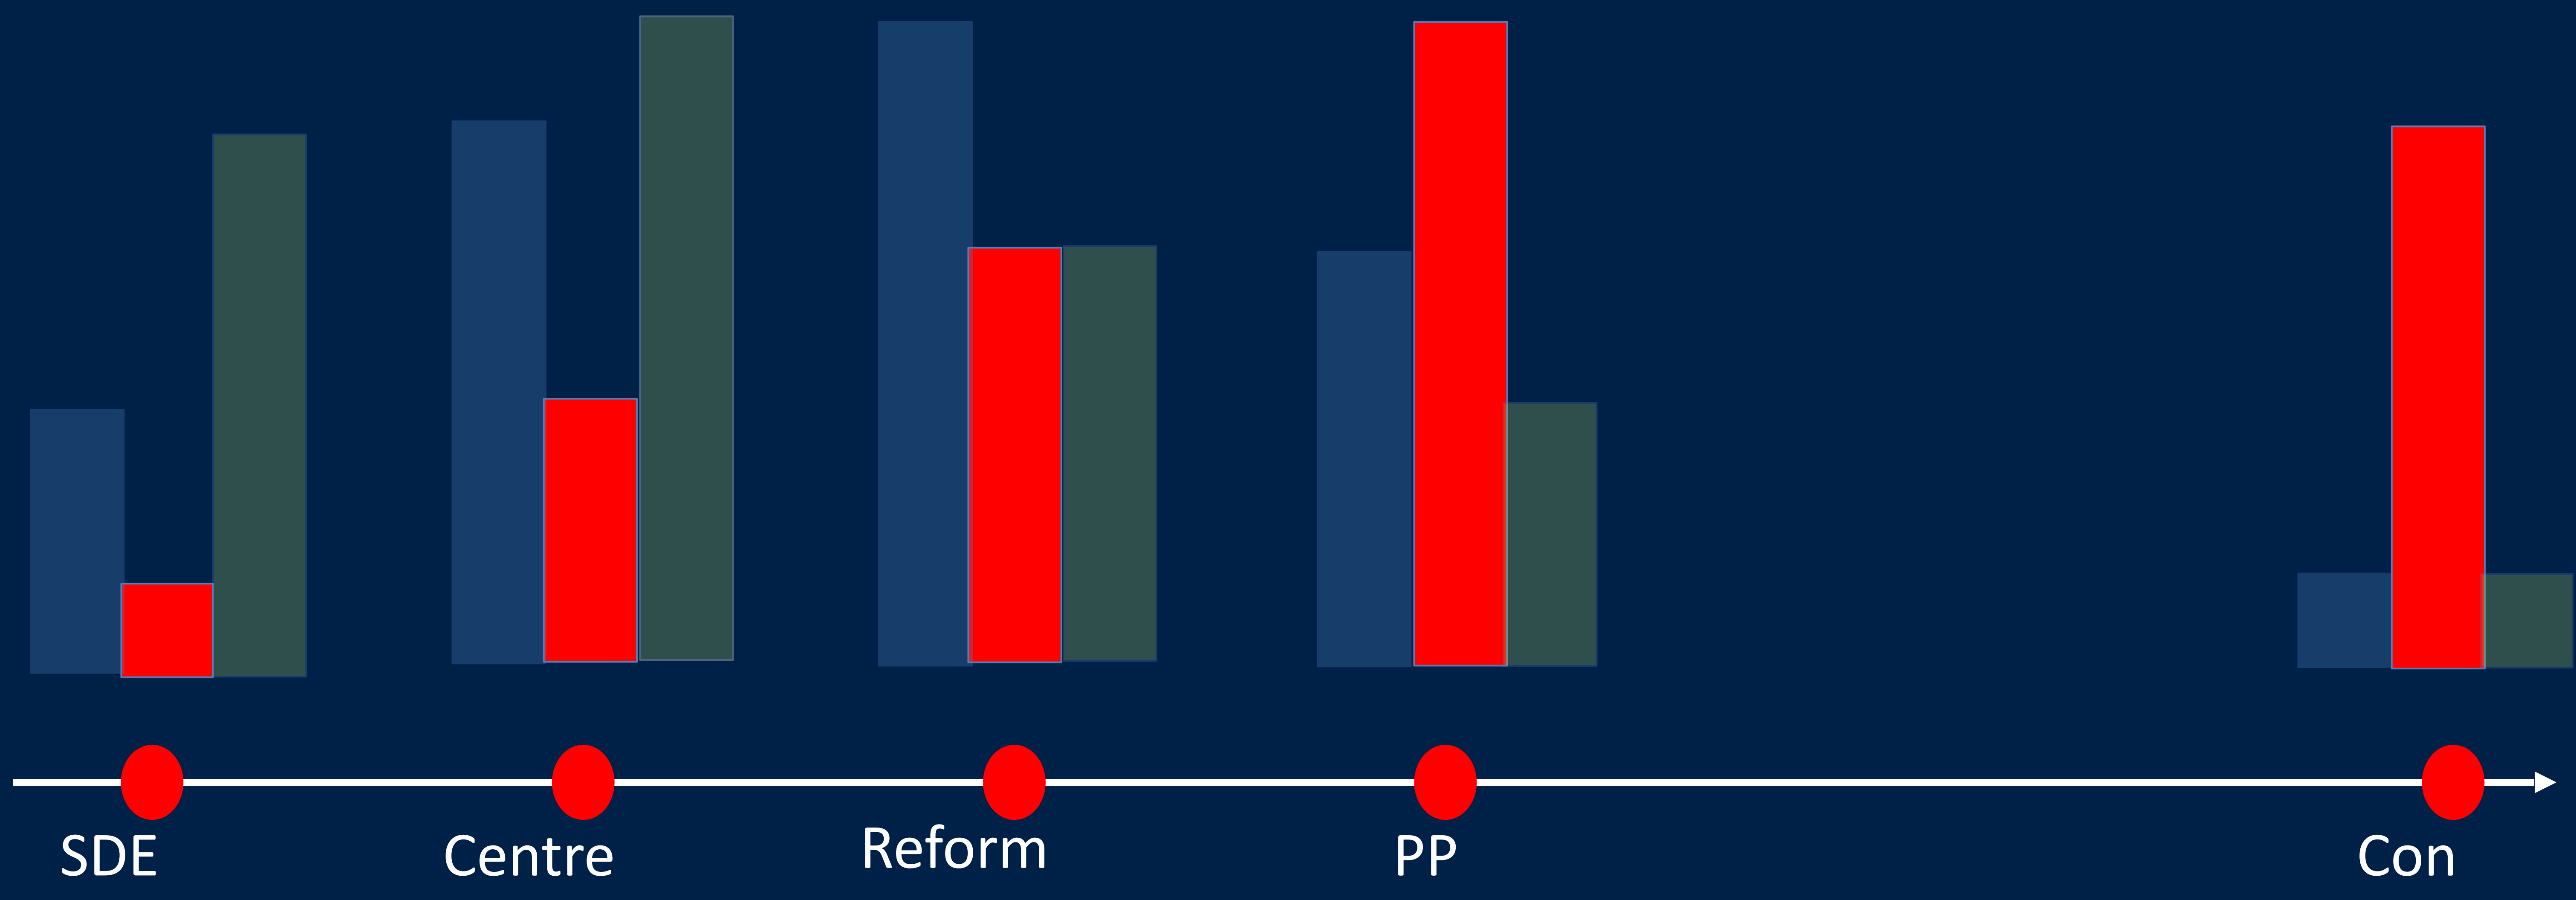

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

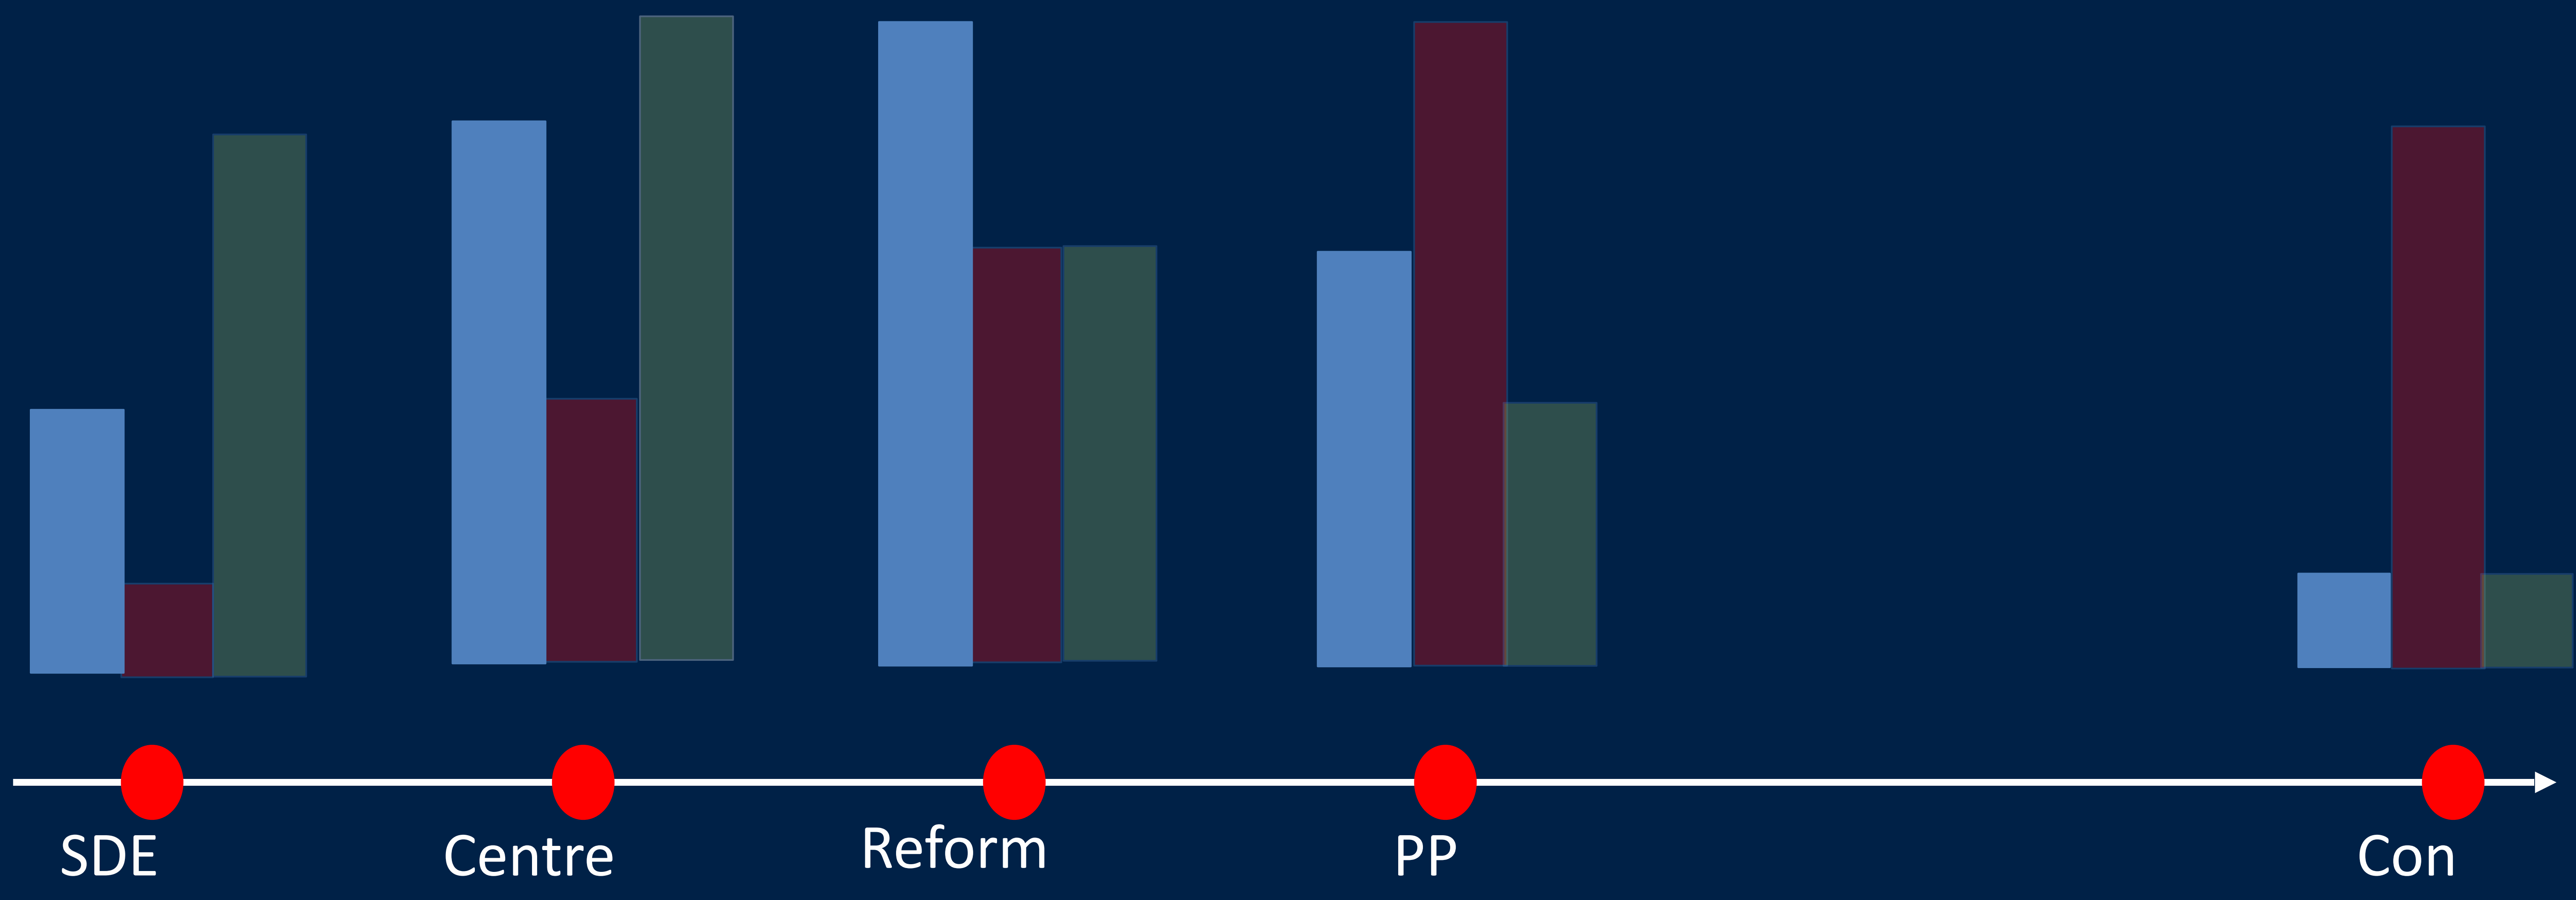

Example: Estonian Parliament

v_1: PP > Con > Reform > Centre > SDE

v_2: Reform > Centre > PP > SDE > Con

v_3: Centre > SDE > Reform > PP > Con

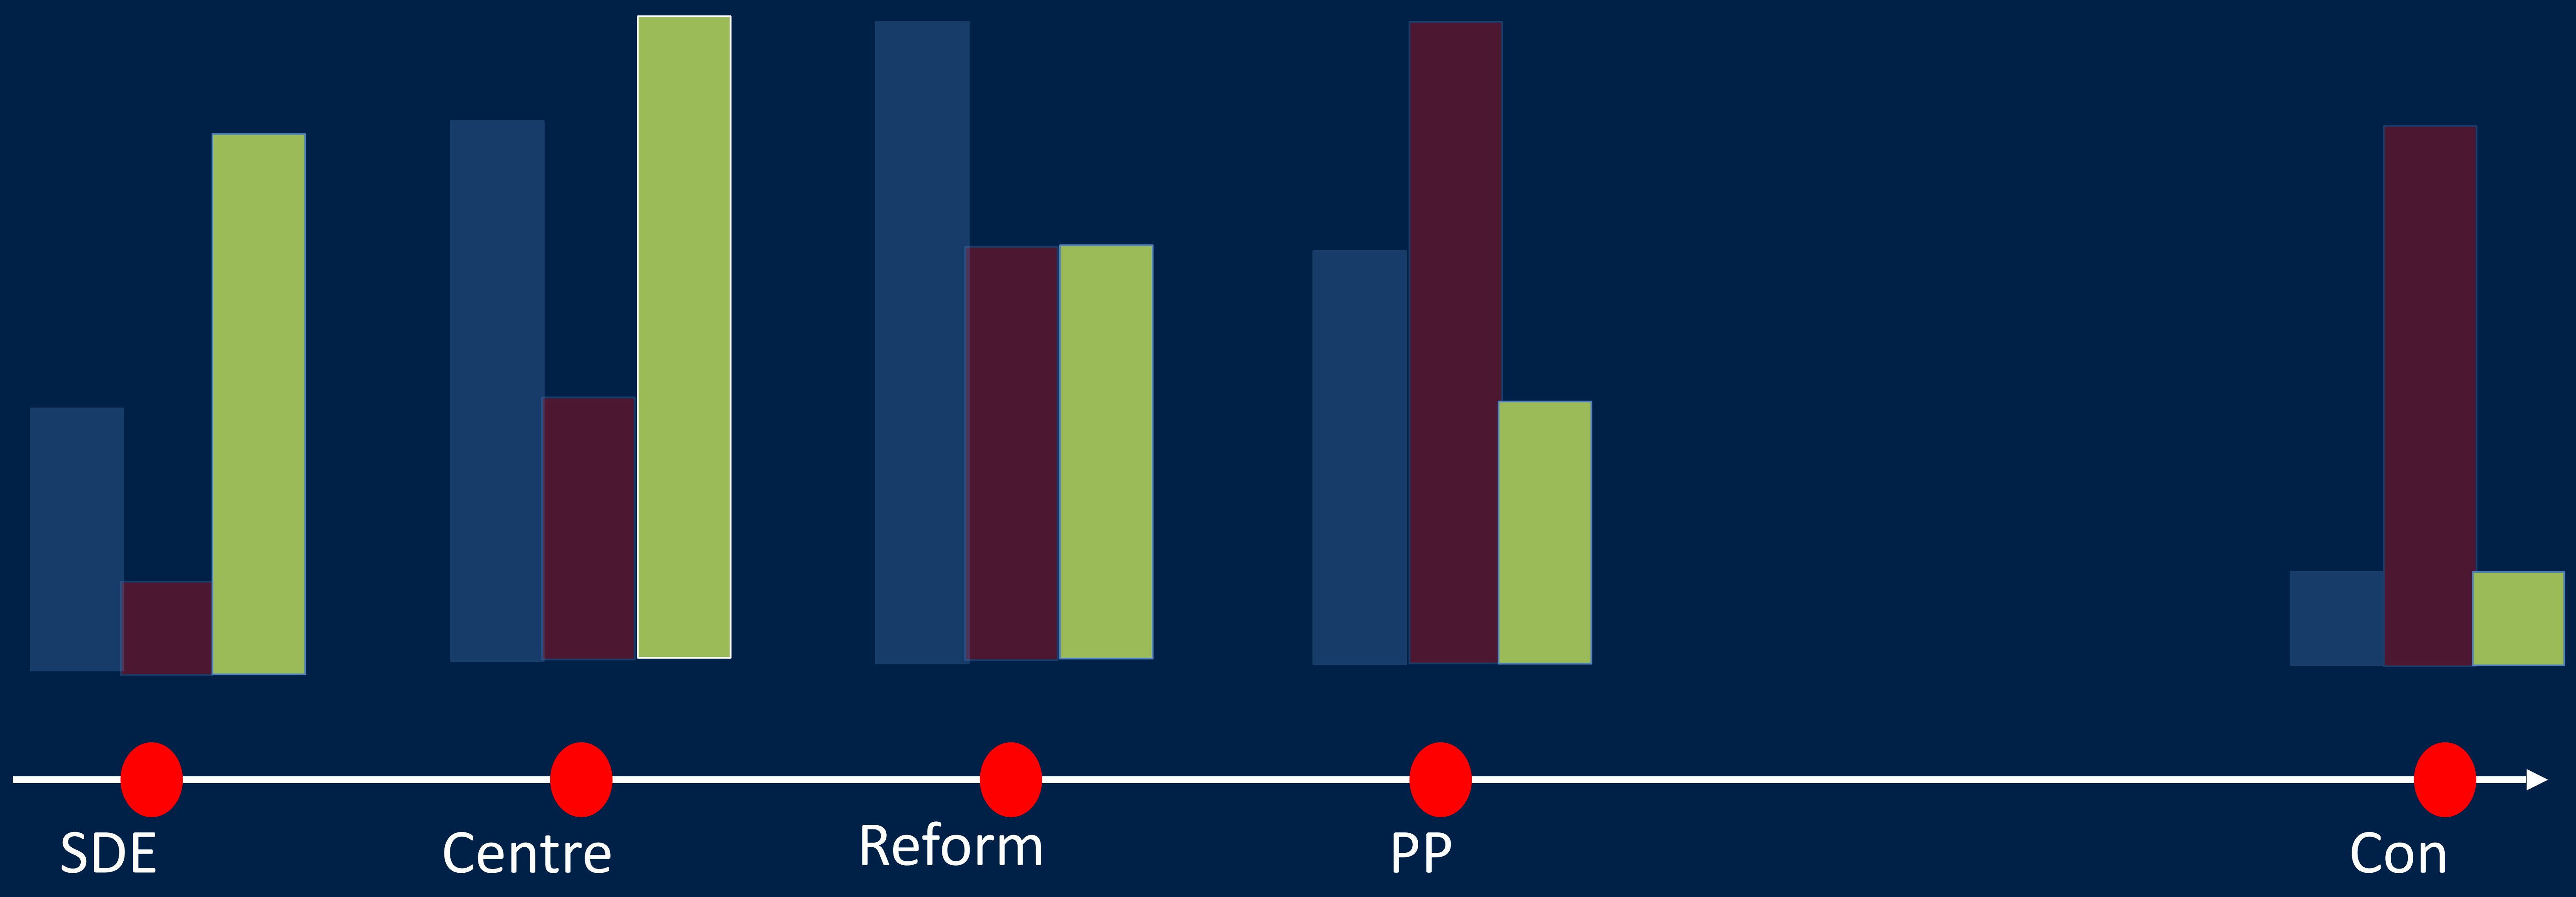

Single-peaked preferences

Setup:

- a set of m candidates C
- each voter ranks candidates from best to worst
- top(v): most preferred candidate of voter v

Definition:

A vote v is single-peaked (SP) wrt an ordering < of candidates (axis) if it holds that:

- if top(v) < d < e, v prefers d to e
- if a < b < top(v), v prefers b to a

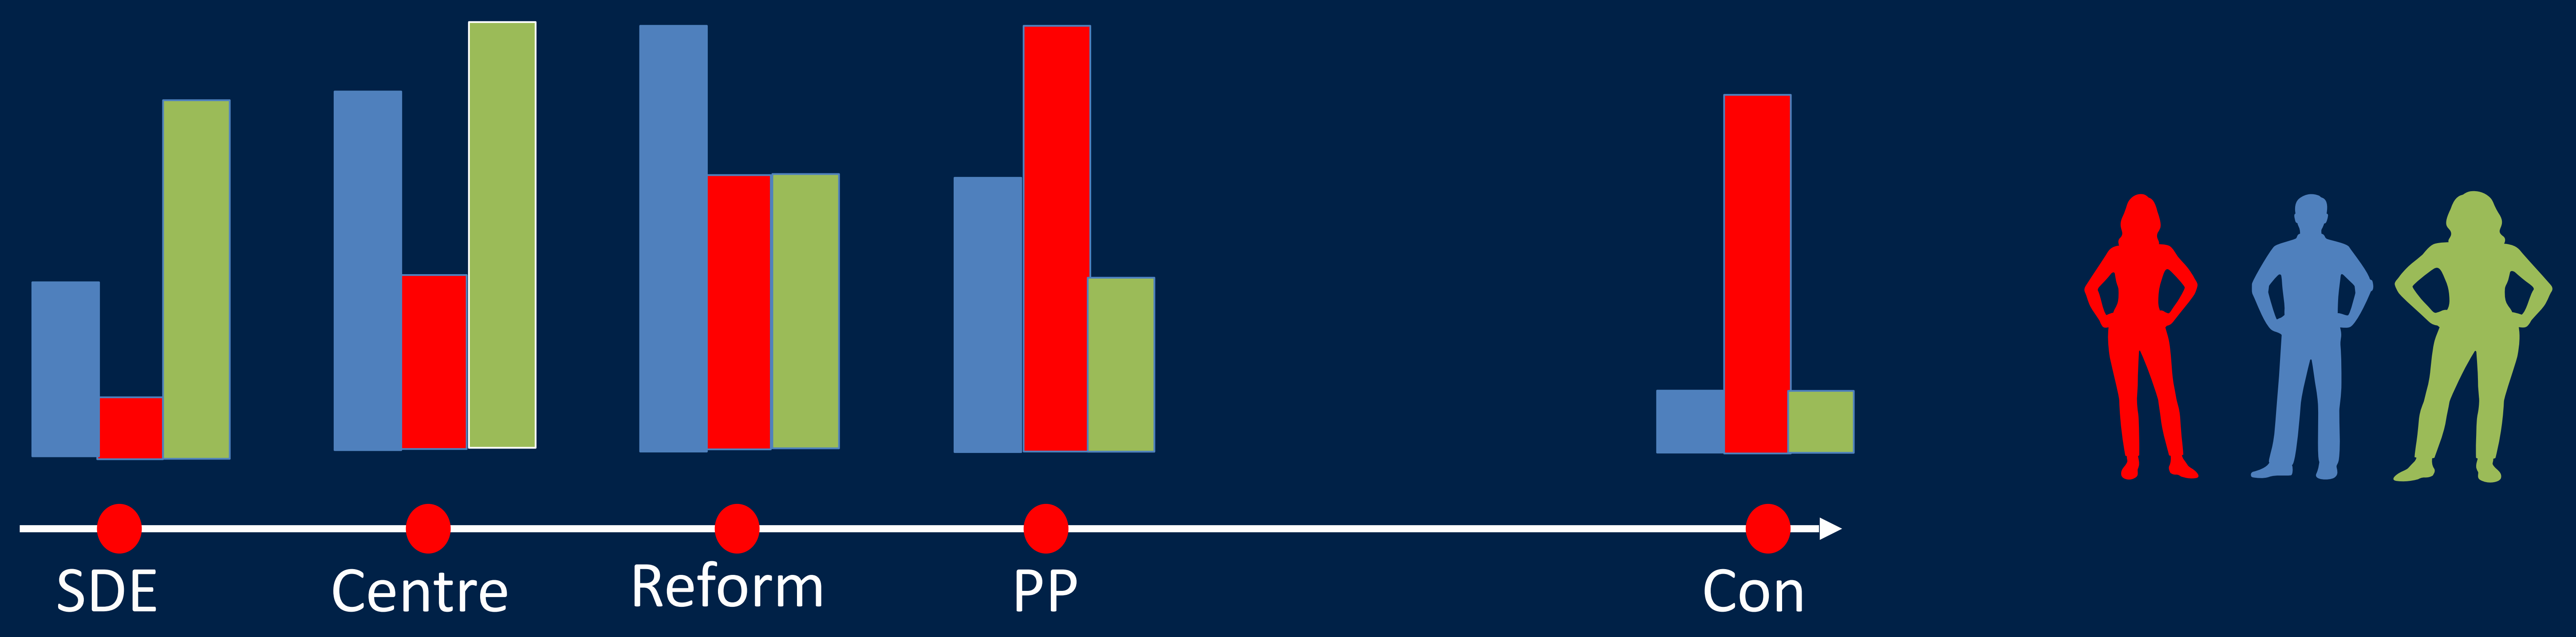

Single-peaked preferences: Challenge

For a set of votes, we can decide in linear time whether it is single-peaked (Doignon et Falmagne, 1994)...

...but if there is a unique axis <, what can we say about how much information we have to see before we learn it exactly?

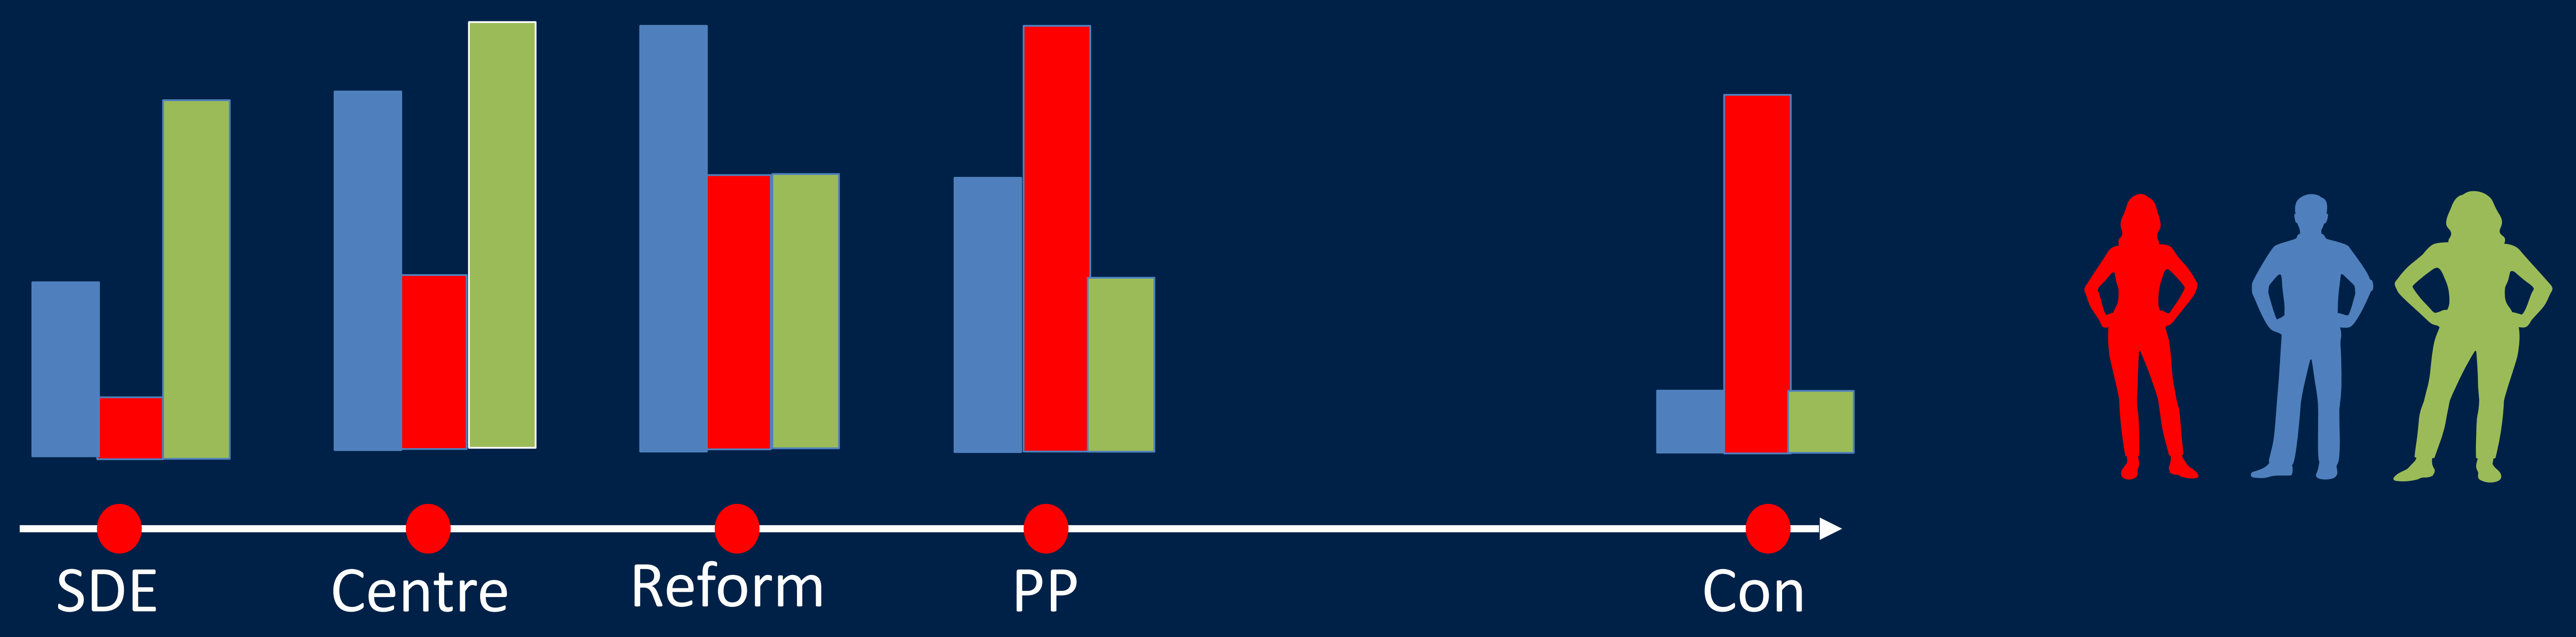

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

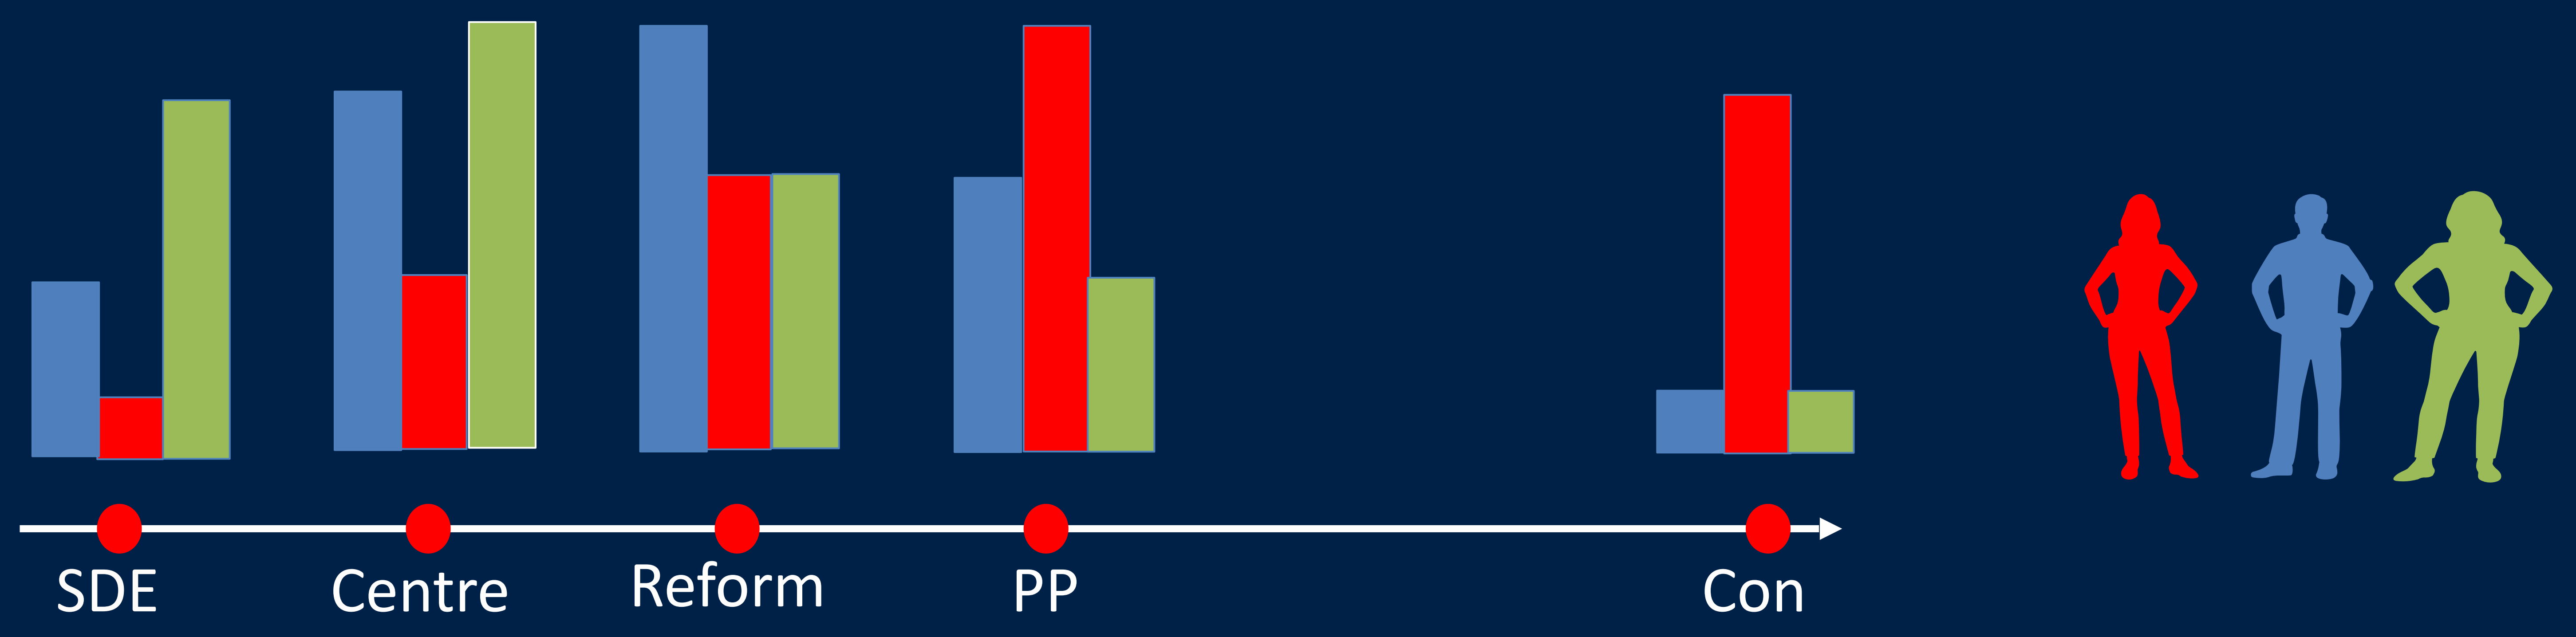

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

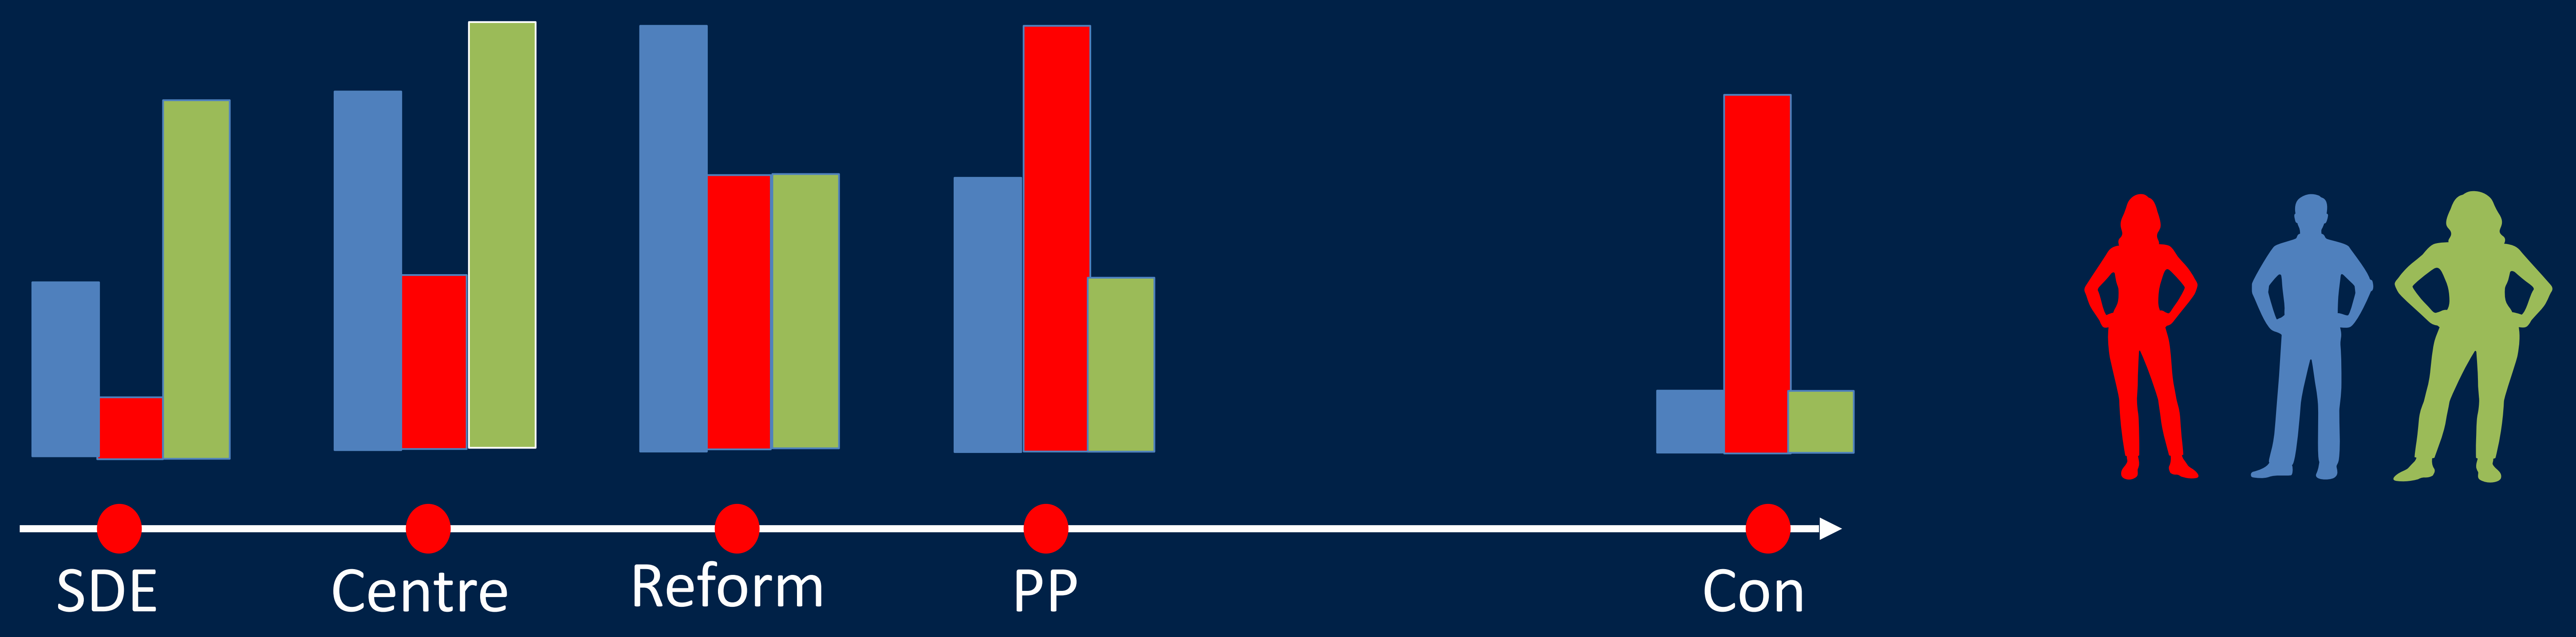

ICML | 2022

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

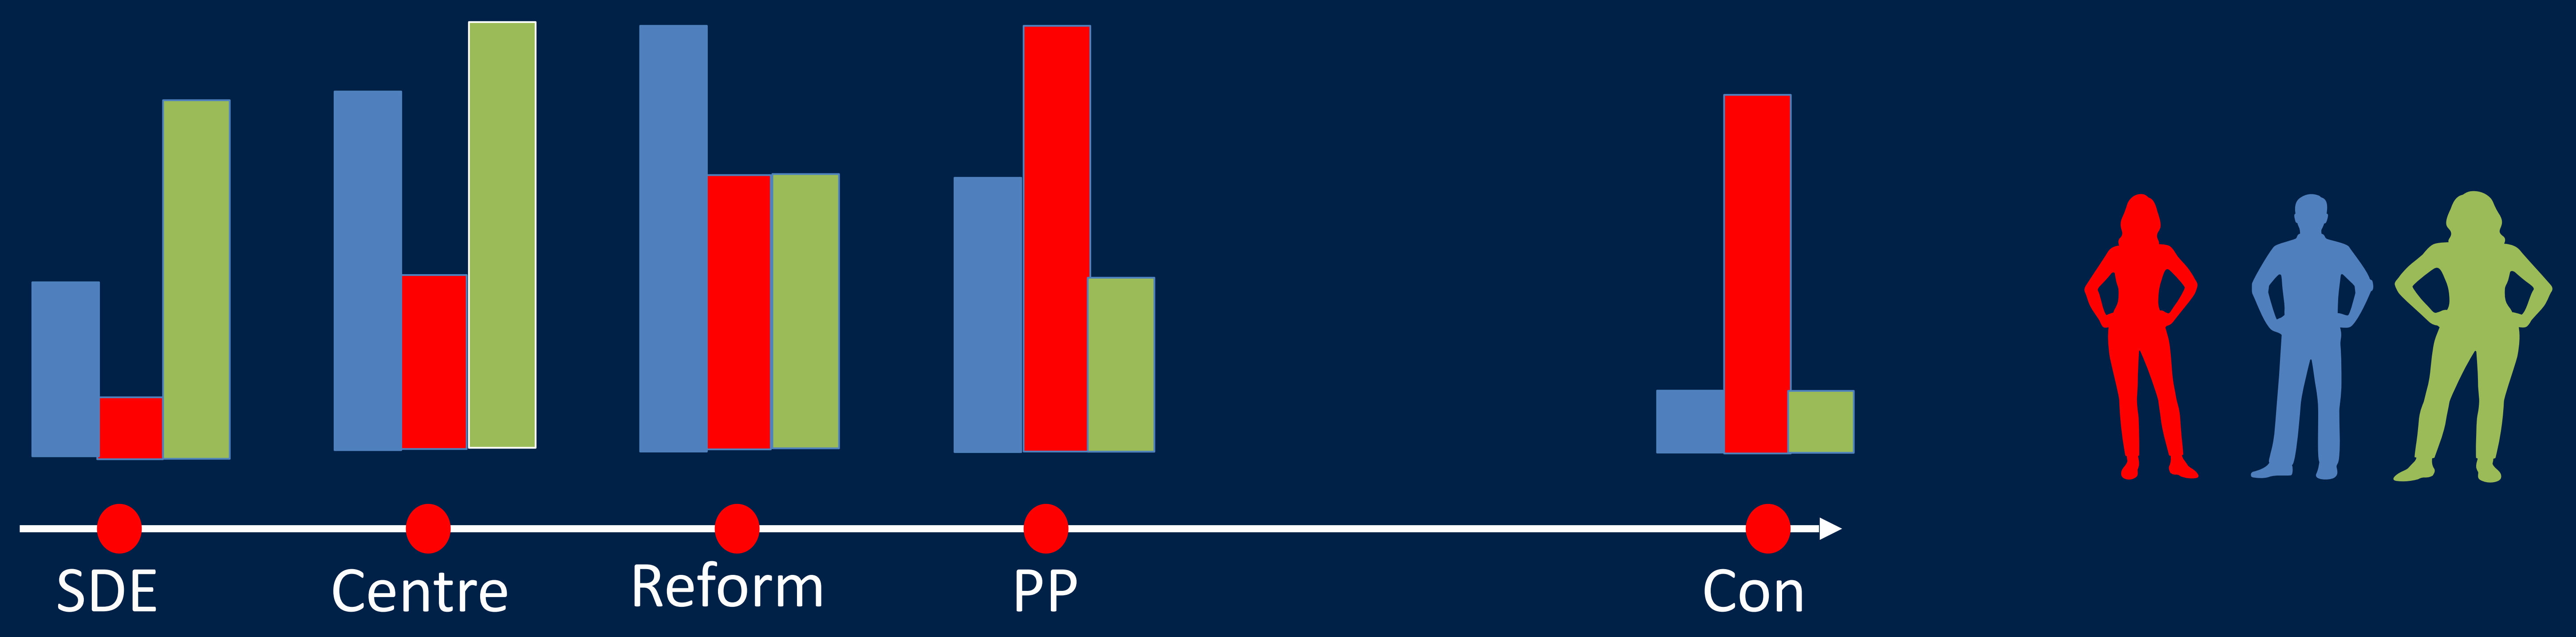

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

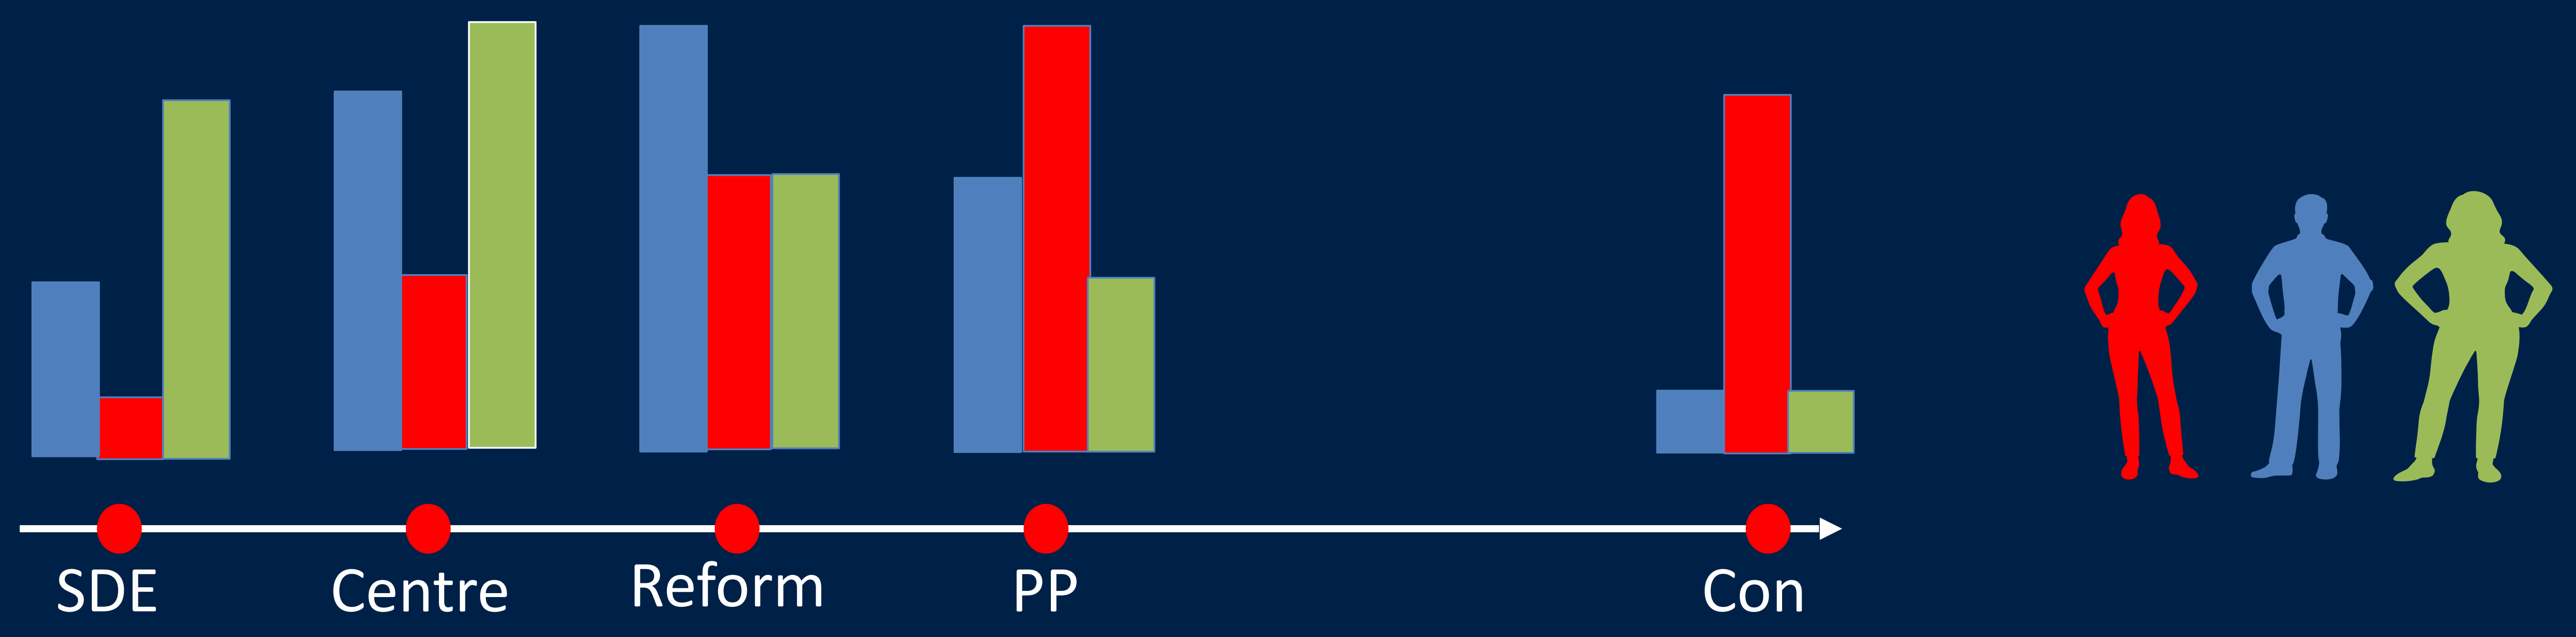

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Independent of # candidates!

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

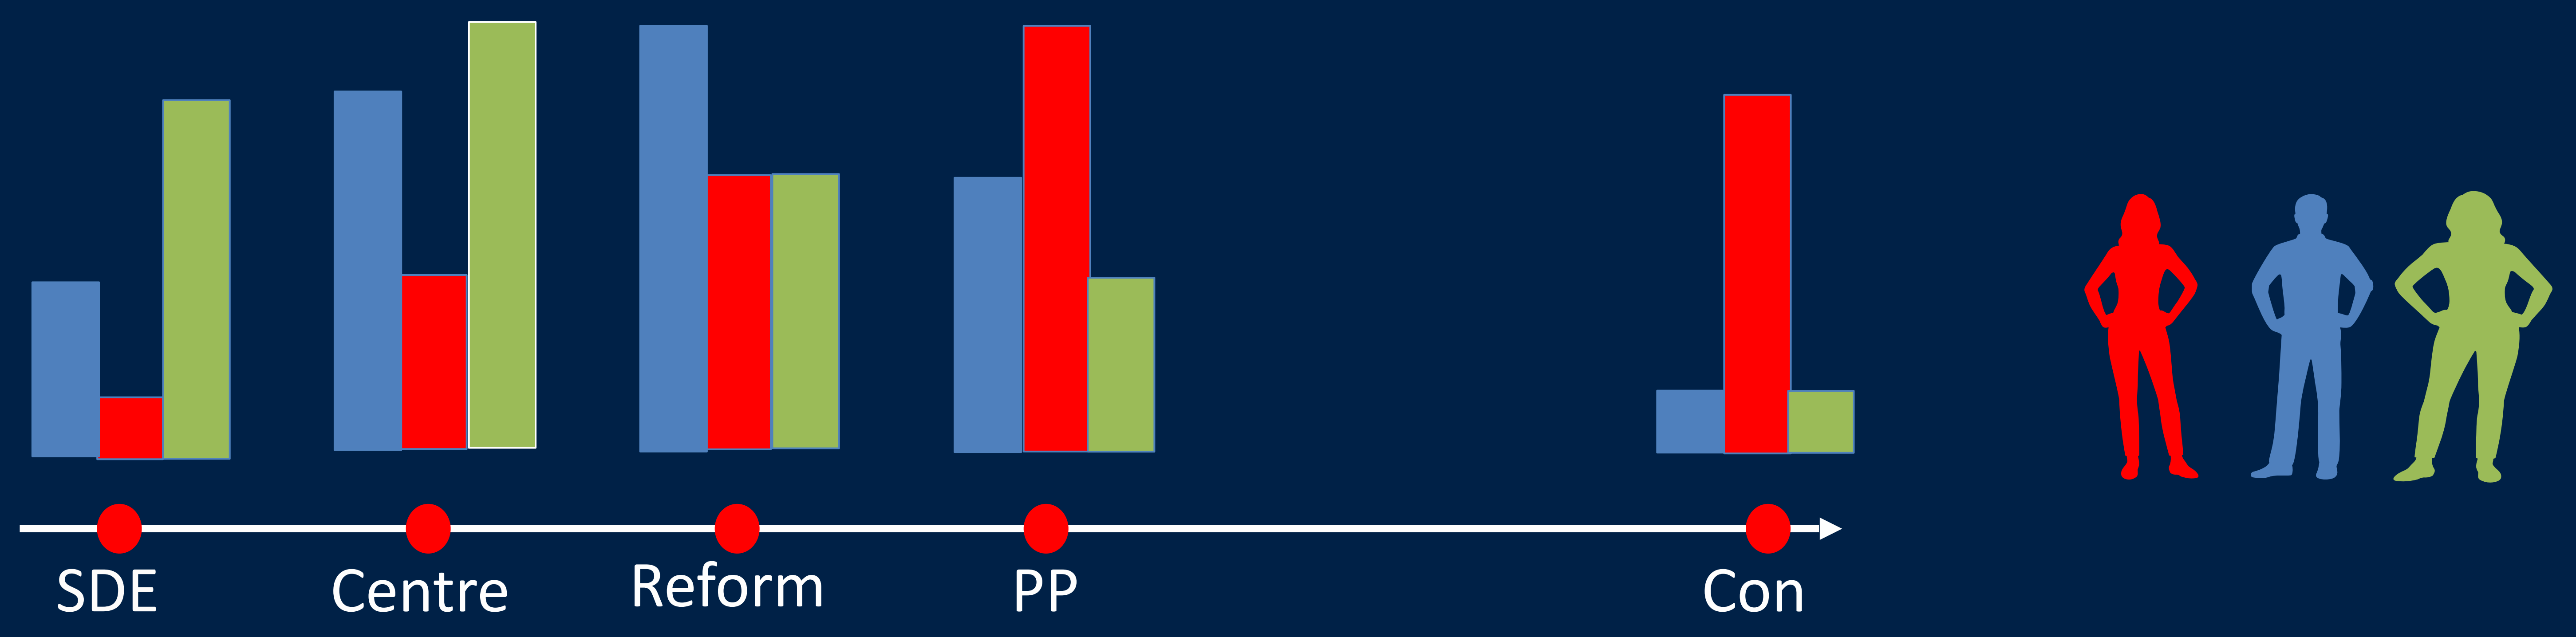

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

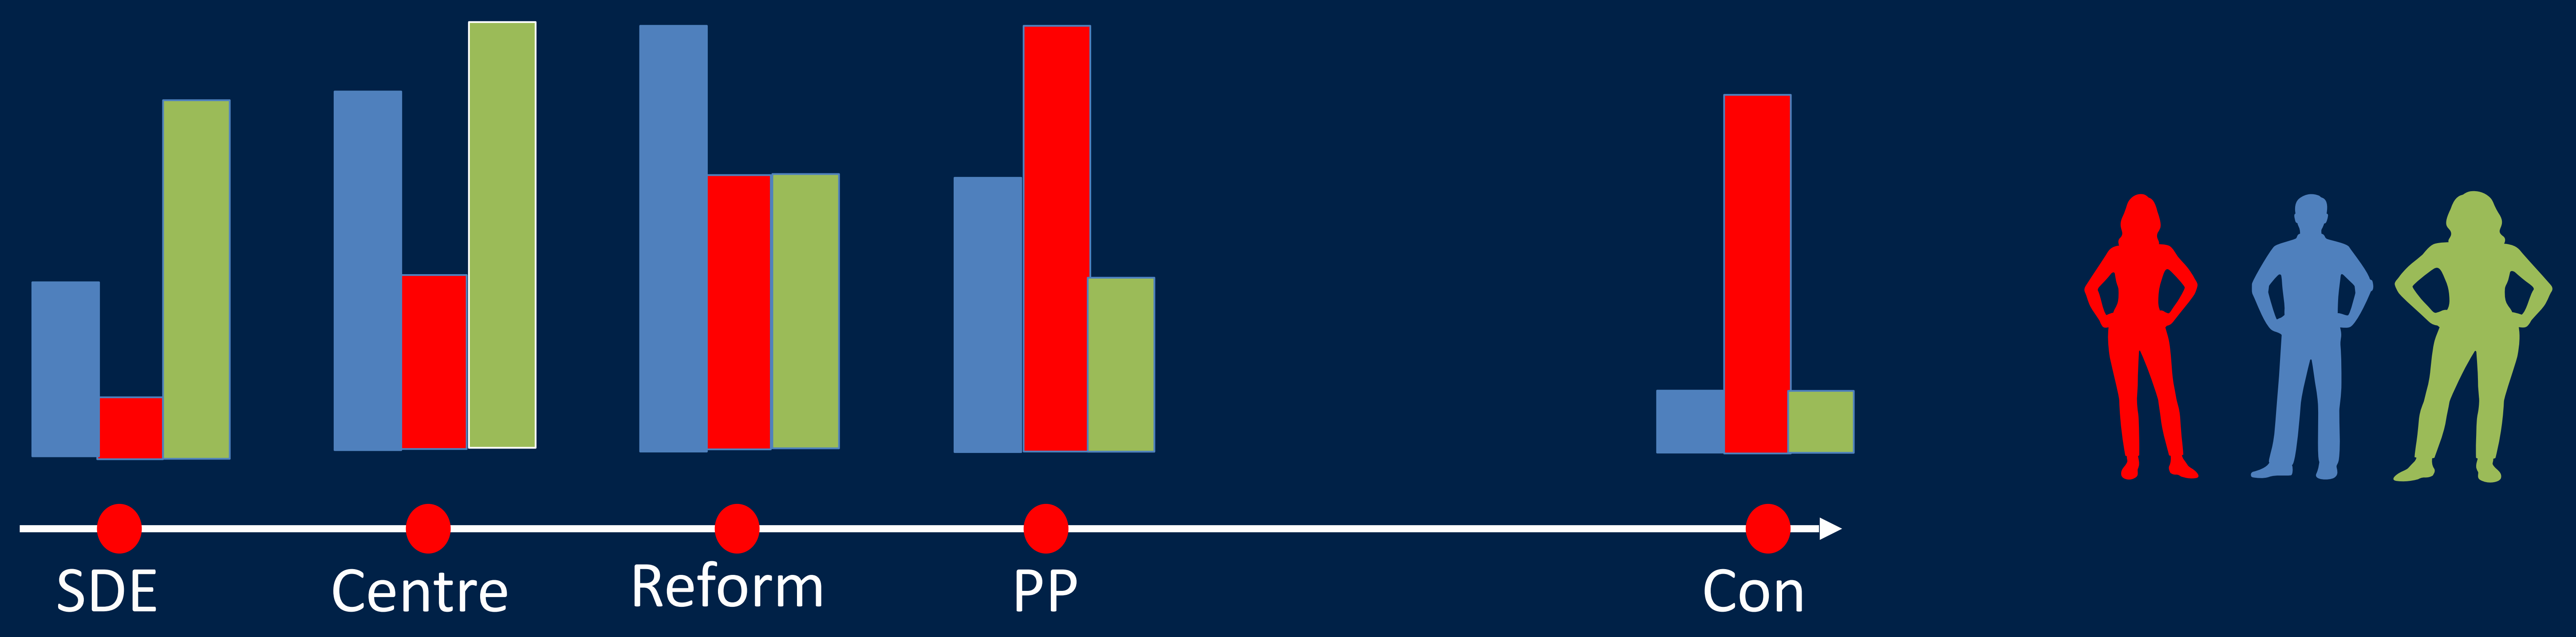

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model

#candidates-1

2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

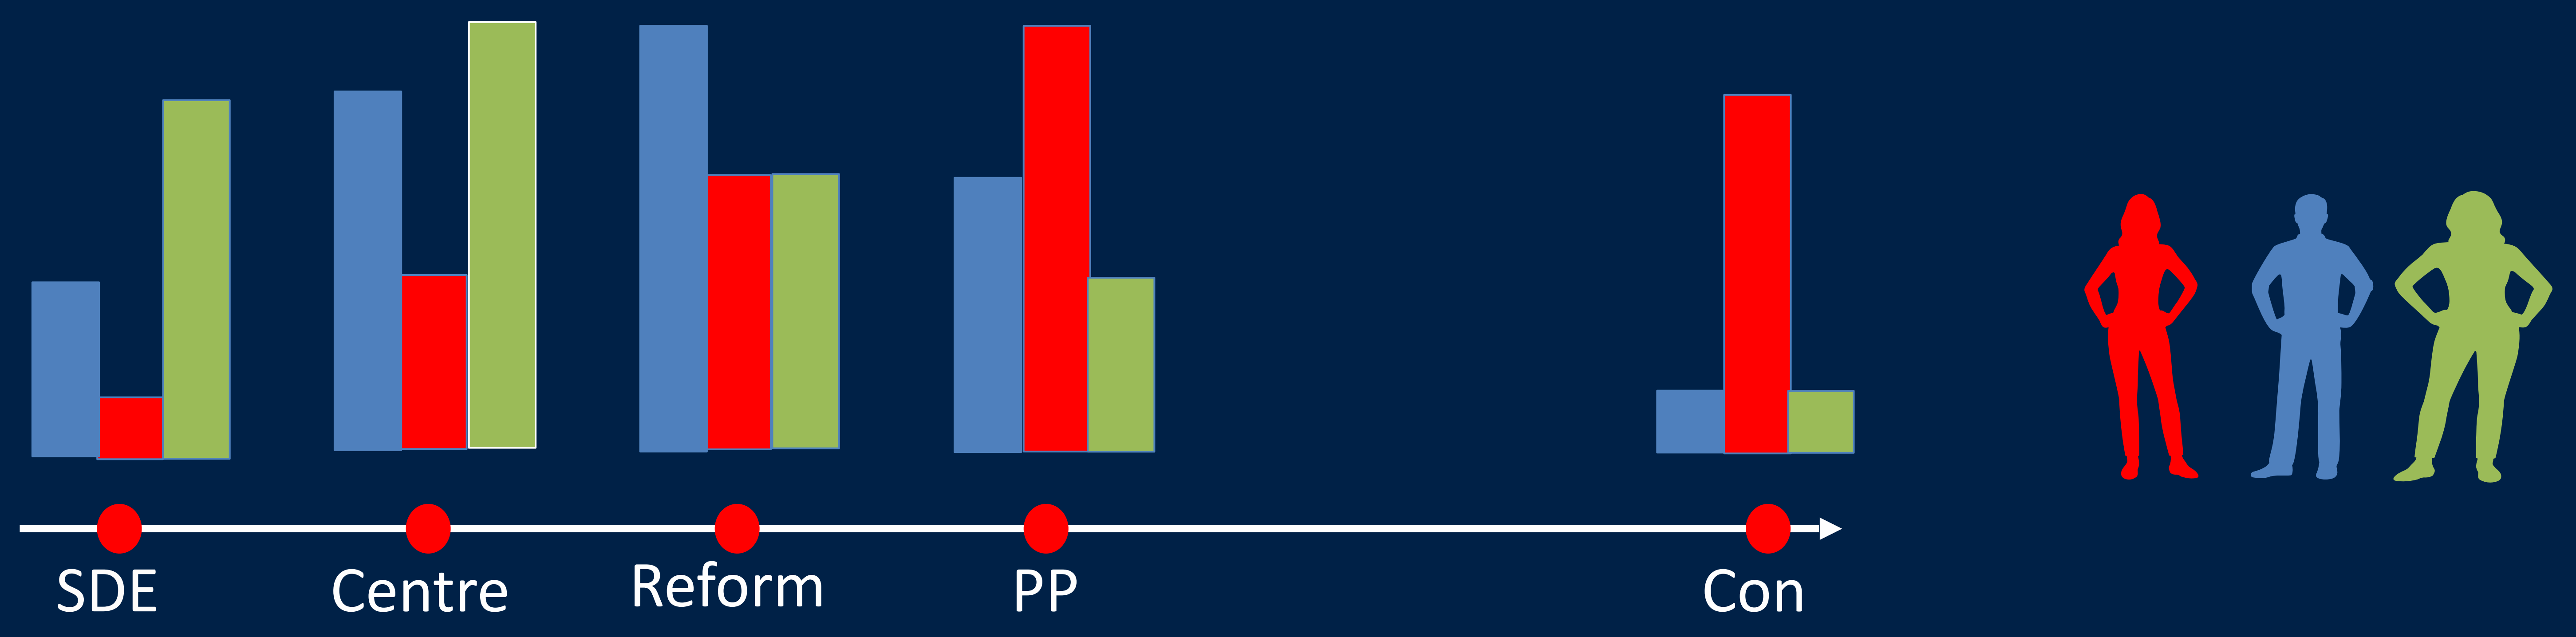

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

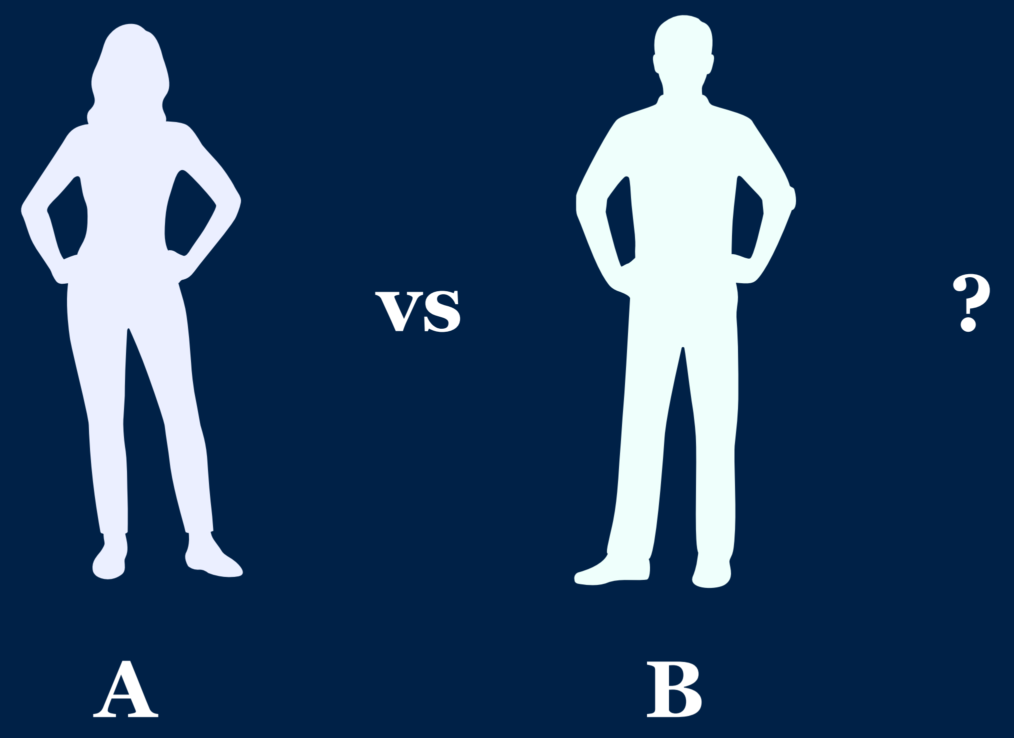

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
    A. By sampling voters' preferences from common distributions
    B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   - A. By sampling voters' preferences from common distributions
   - B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on “learning as much as possible” by sampling voter's preferences

Our Contributions

Learning a single axis:

1. Tight bounds on uniquely identifying the underlying axis
   A. By sampling voters' preferences from common distributions
   B. In the equivalence query model
2. Upper bounds on learning the underlying axis from binary comparisons from common distributions

Learning two axes:

Upper and lower bounds on "learning as much as possible" by sampling voter's preferences