United for Change: Deliberative Coalition Formation to Change the Status Quo

Edith Elkind

University of Oxford

joint work with

Davide Grossi (Groningen),

Udi Shapiro (Weizmann),

Nimrod Talmon (Ben Gurion)

Spatial Coalition Formation

- Traditional setup:
  - agents = political parties
  - parties have/adopt positions in a multidimensional proposal space
  - parties aim to form a winning coalition to govern
  - a coalition is typically associated with a position in the issue space
- Significant literature:
  - starting from Hotelling’29, see survey by de Vries’99, subsequent work e.g., by Rusinowska, de Swart and co-authors

This Work

- Proposal space: a metric space
  - e.g., can be a finite or infinite subset of $\mathbb{R}^d$
- Each agent has an ideal point (her opinion)
- Special point: the status quo (r)
  - goal: majority-supported change from the status quo
- Agents seek change and are open to compromise
  - approval preferences

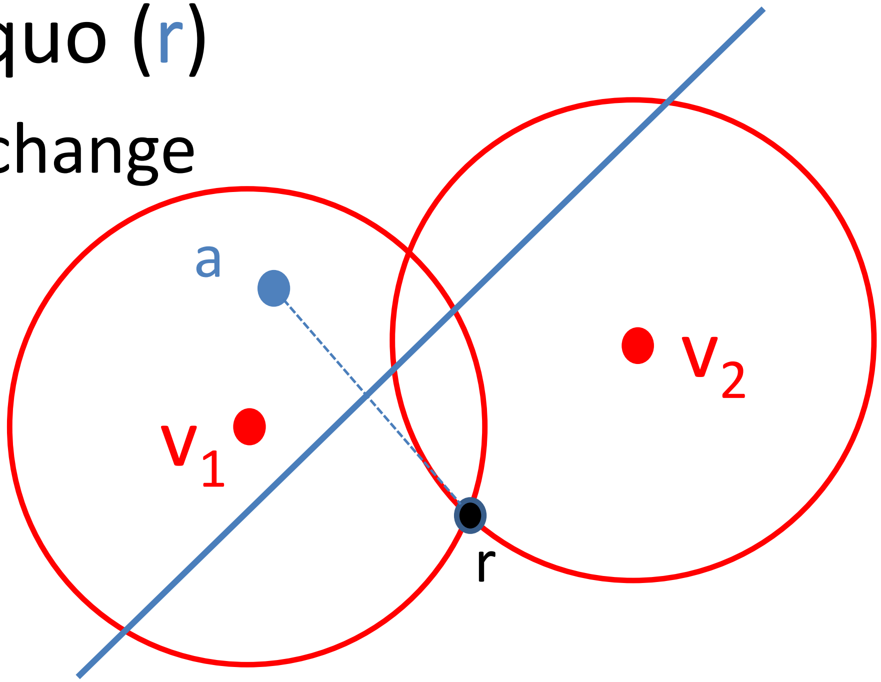

Example

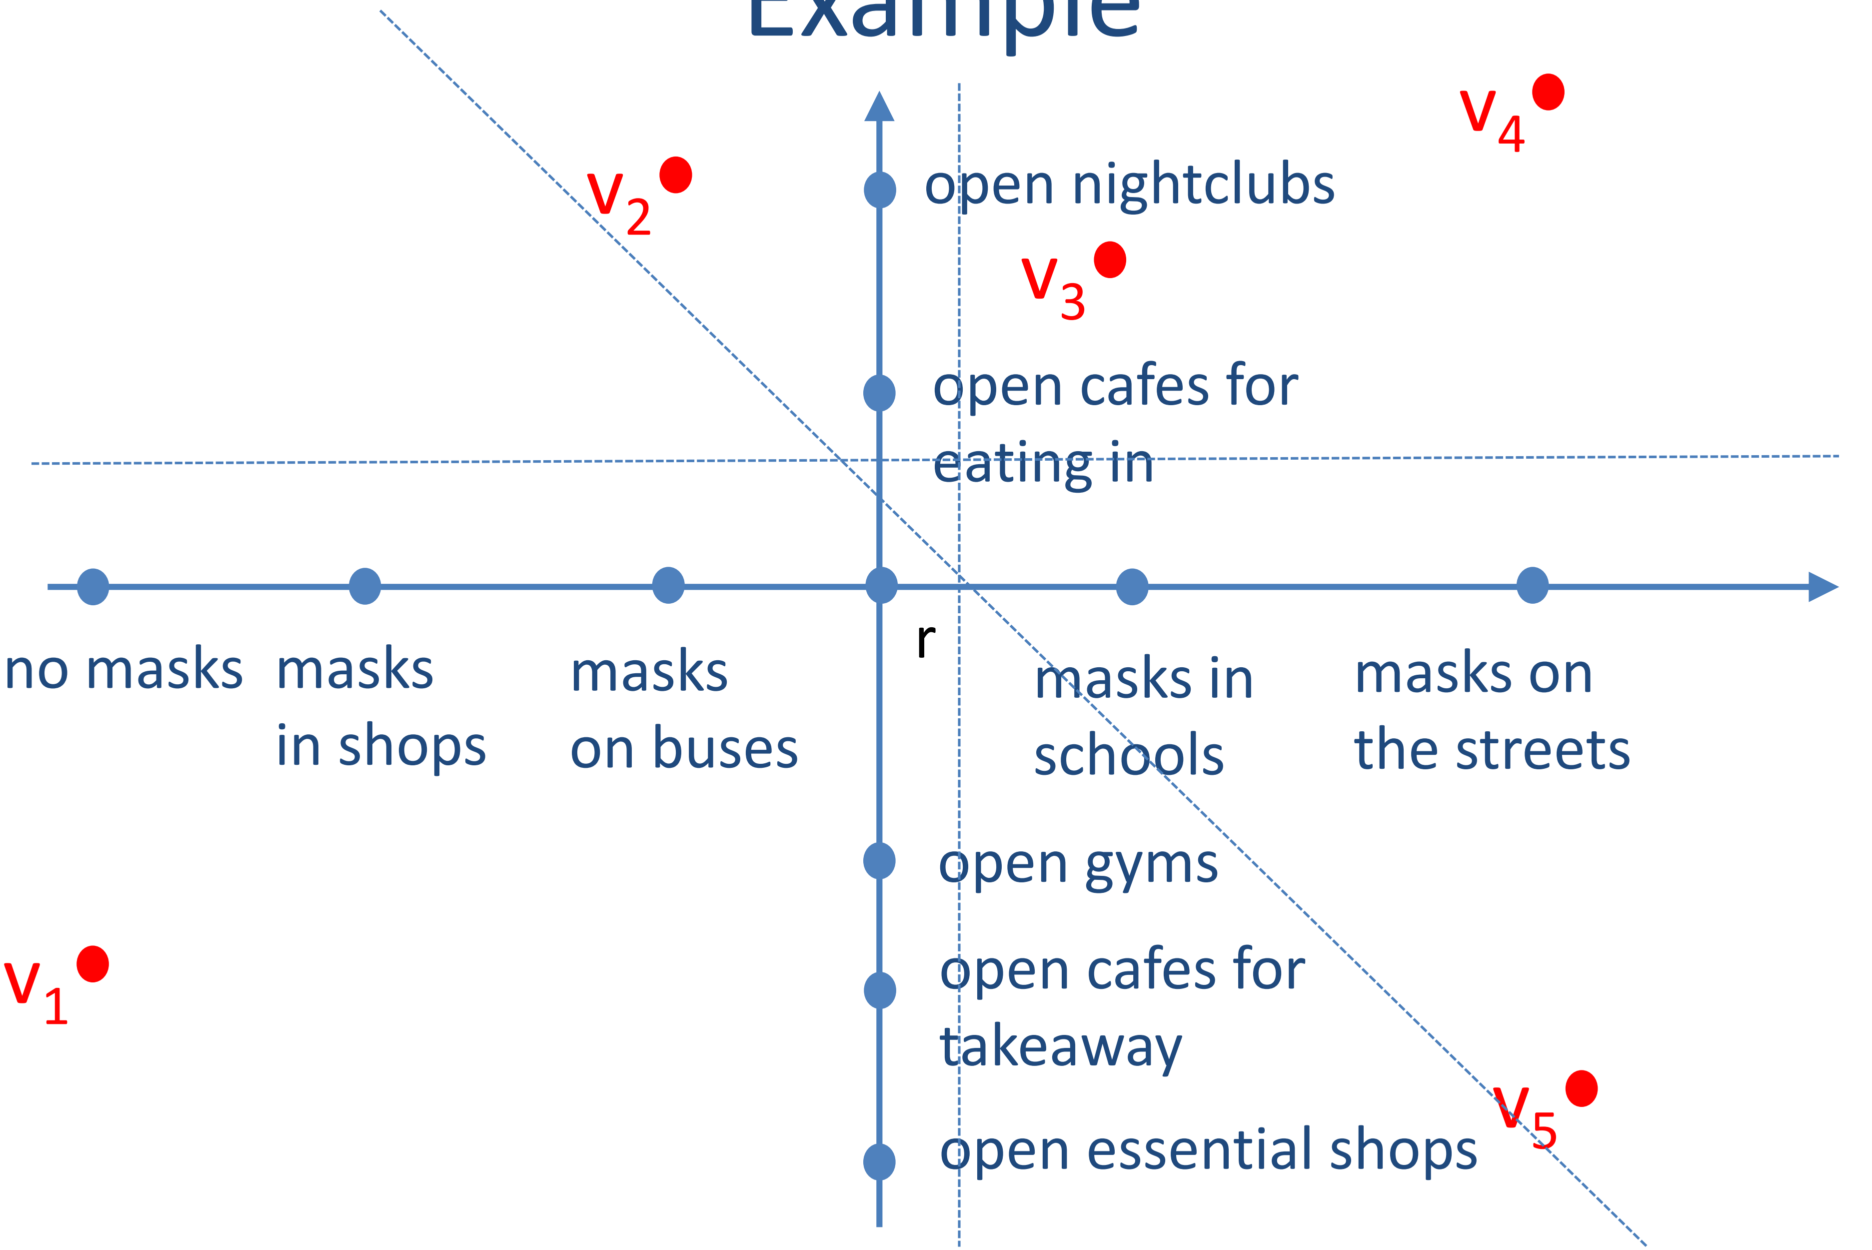

open nightclubs
open cafes for eating in
open gyms
open cafes for takeaway
open essential shops
no masks
masks in shops
masks on buses
masks in schools
masks on the streets
r
v1
v2
v3
v4
v5

Coalition Formation Dynamics

- Deliberative coalition (C, p) = set of agents + position
  - all agents in C prefer p to the status quo
- Coalition structure: partition of agents into deliberative coalitions: (C, p), (C’, p’), …
- Types of transitions:
  - single-agent deviations, position changes, merges, merges with some agents left behind
  - each transition involves a limited # of coalitions
- Agents favour larger coalitions
  - a transition is only feasible if it leads to formation of a larger coalition
  - but they do not distinguish among approved proposals

- Research question:
  which types of transitions guarantee emergence of a coalition around one of the most supported outcomes?
  - the answer may depend on the metric space
- Can we converge after polynomially many transitions?

Transitions

- Single-agent transitions:
  …, (C, p), (C', p'), … → …, (C+a, p), (C'-a, p'), …
  - permissible iff $|C| \geq |C'|$, a approves p
- Follow transitions:
  …, (C, p), (C', p'), … → …, (C U C', p), …
  - permissible if all members of C' approve p
- Merge transitions:
  …, (C, p), (C', p'), … → …, (C U C', p*), …
  - permissible if all members of C U C' approve p*
- Compromise transitions:
  …, (C, p), (C', p'), … →
  …, $(C \setminus C_{p^*}, p), (C' \setminus C'_{p^*}, p'), (C_{p^*} \cup C'_{p^*}, p^*)$, …
  - permissible if $C_{p^*} \cup C'_{p^*}$ approve p*, $|C_{p^*} \cup C'_{p^*}| > |C|, |C'|$

Example

3 deliberative coalitions:

$C_1 = \{1, 2, 3; a\}$

$C_2 = \{4, 5, 6; b\}$

$C_3 = \{7; c\}$

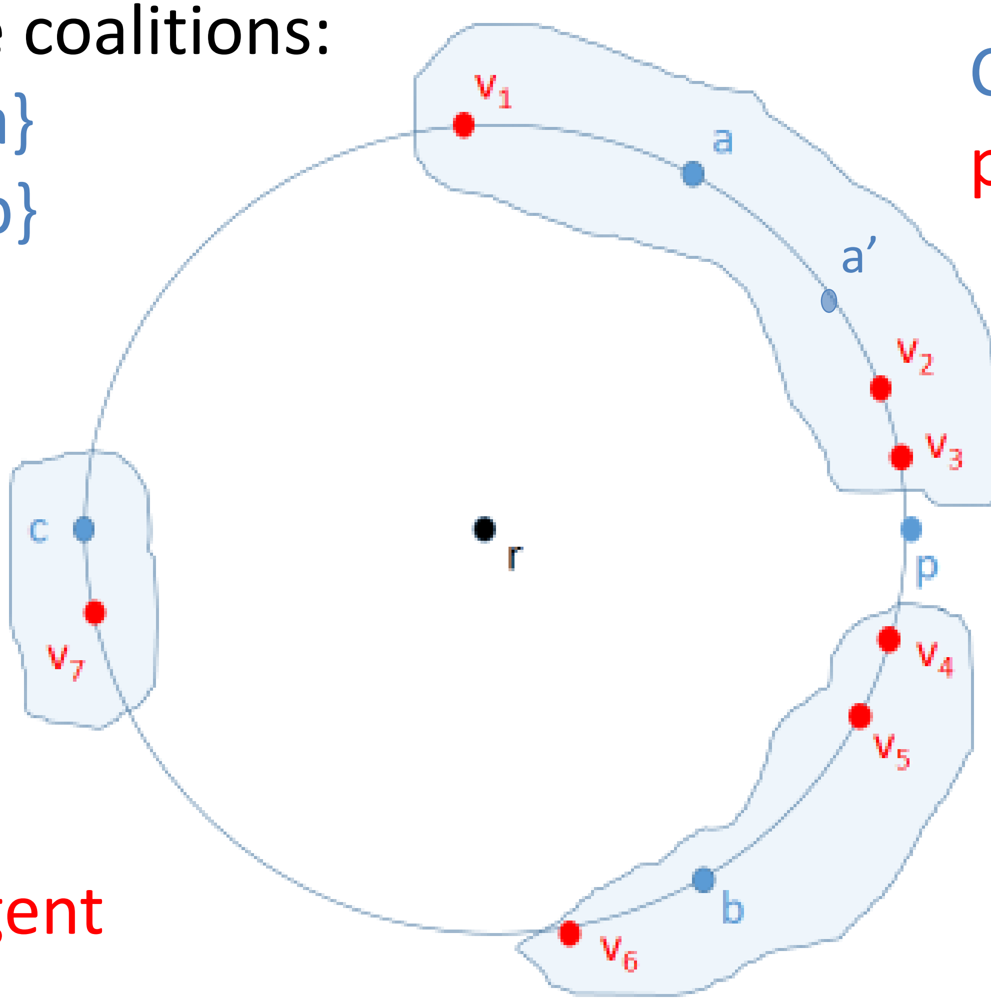

C_1 can change position to a'

no merge transitions

majority supports p

no single-agent transitions

Warm-Up: Convergence in 1D

- Observation: in $\mathbb{R}$, single-agent transitions may fail to succeed

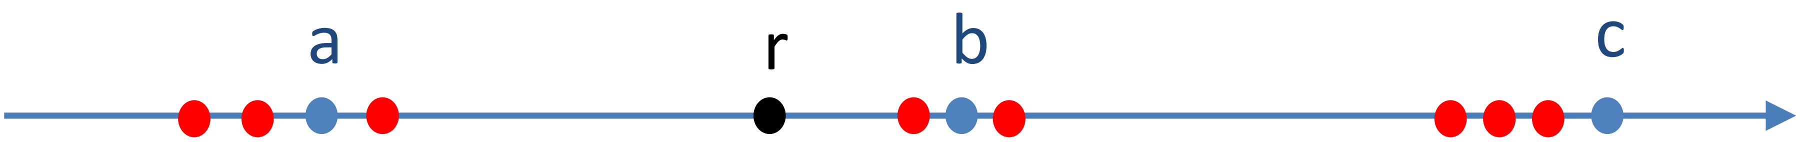

- Theorem: in $\mathbb{R}$, follow transitions converge
  - no coalition spans 0
  - if there are two "positive" coalitions (C, p), (C', p') with $p < p'$, then C' can follow C
  - so if no transitions are available, we have ≤2 coalitions (one +ve, one -ve)

Beyond One Dimension?

- <u>Observation</u>: in $\mathbb{R}^2$, single-agent and follow transitions may fail to succeed

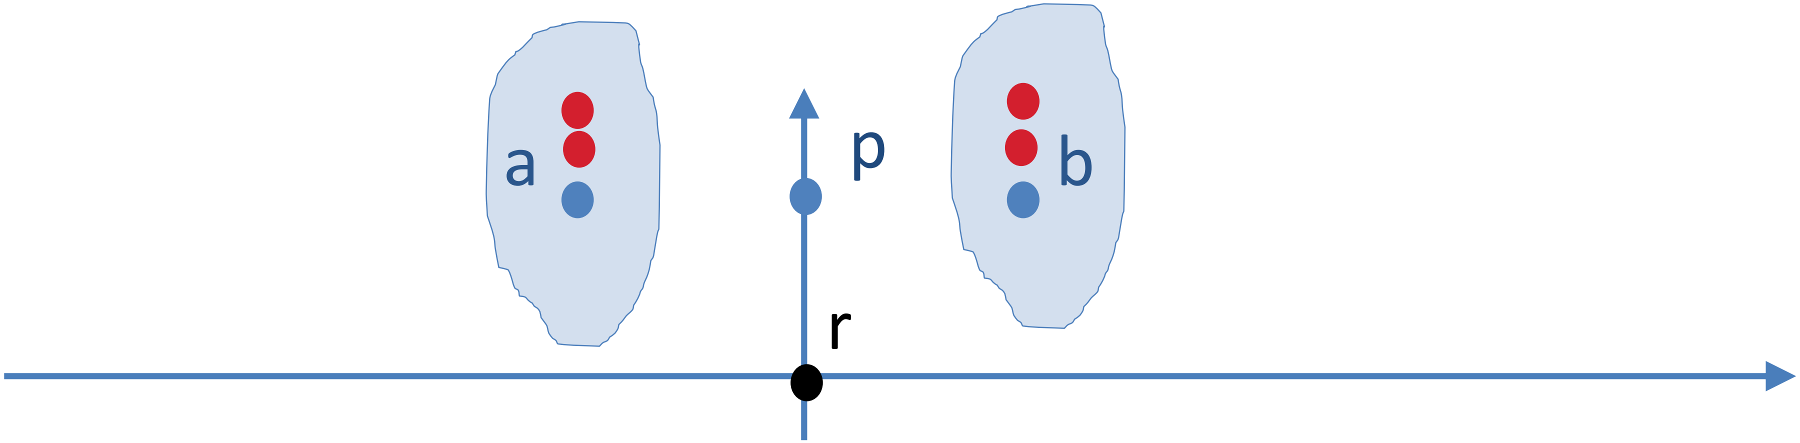

Do Merges Help?

- Theorem: if the metric space is a tree, merge transitions succeed
- Proof: if there is a “good” outcome, one of root’s children is “good”

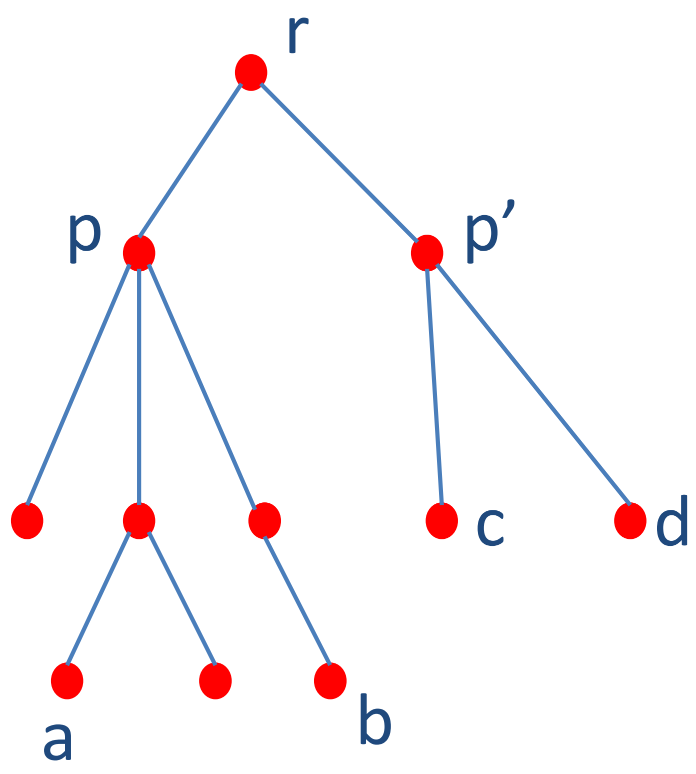

But Not In General?

- Observation: in $\mathbb{R}^2$, single-agent, follow, and merge transitions may fail to succeed

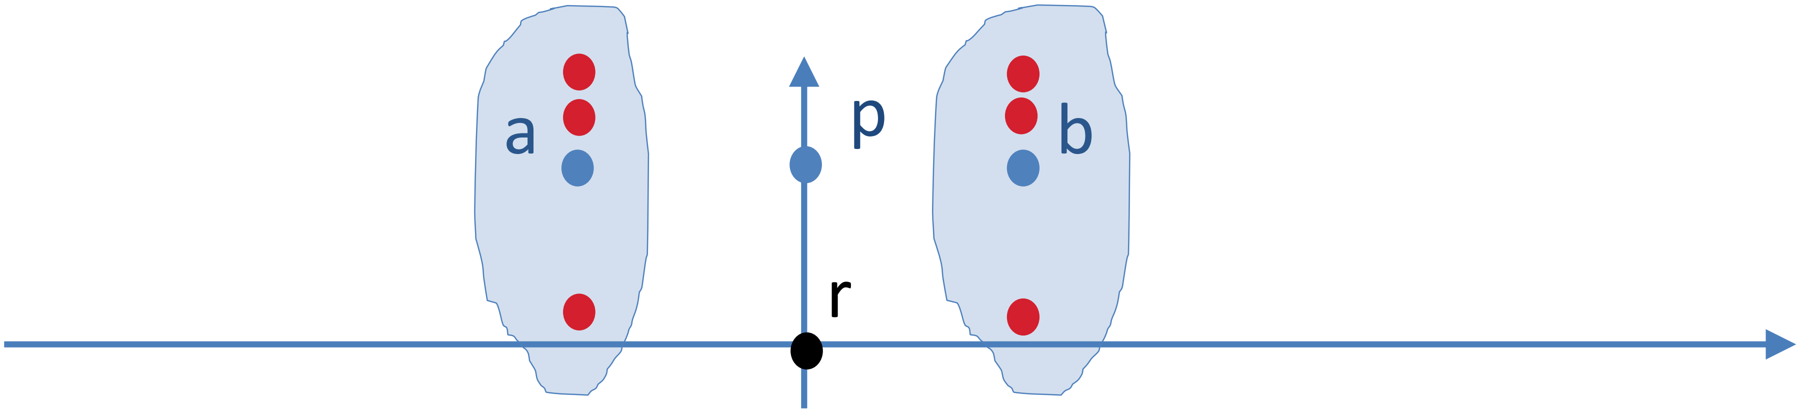

Do Compromises Help?

- Theorem: if the proposal space is $\mathbb{R}^d$, then compromise transitions succeed
- Proof idea:
  - 2 coalitions can always compromise
  - if 3 coalitions are left, either
    - someone can join the largest coalition, or
    - two coalitions can merge

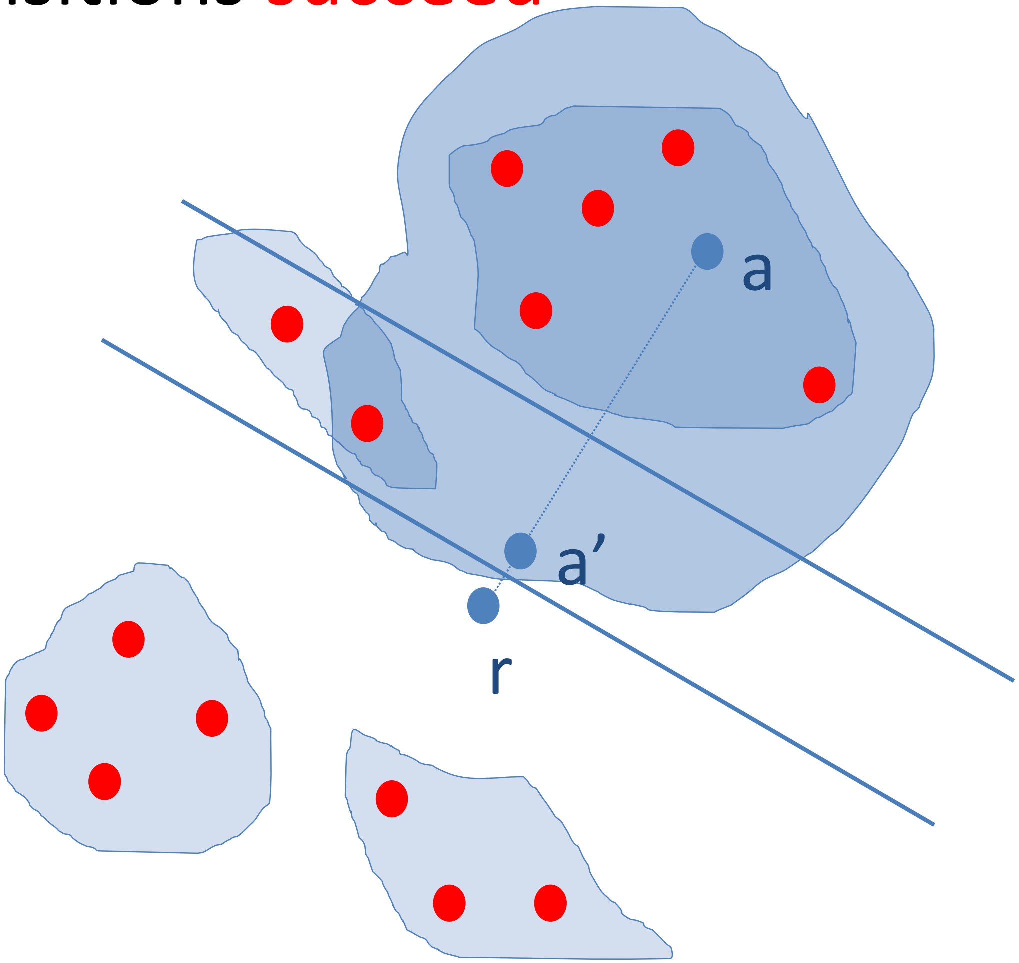

When Compromises Fail...

- The theorem holds if the proposal space is a dense subset of $\mathbb{R}^d$
- ... but not if it is an arbitrary subset of $\mathbb{R}^d$

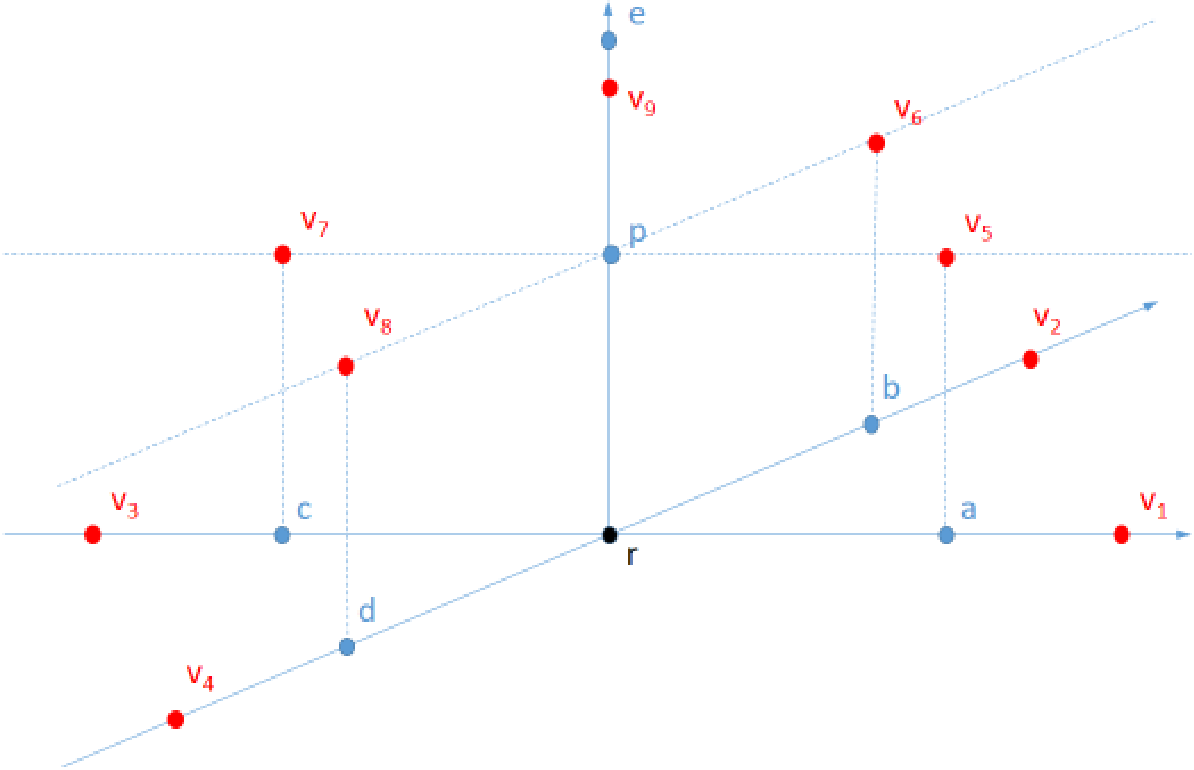

Speed of convergence (1/2)

- Claim: any sequence of single-agent, merge and follow transitions terminates in $O(n^2)$ steps (where n is the number of agents)
- Proof:
  - given a coalition structure $(C_1, p_1), \dots, (C_k, p_k)$, consider $Z = |C_1|^2 + \dots + |C_k|^2$
    - quadratic potential function
  - Z takes values between 0 and n^2
  - every transition increases Z

Speed of Convergence (2/2)

- Observation: a compromise transition may fail to increase Z
- Theorem: every sequence of compromise transitions terminates after at most n^n steps
  - a lexicographic potential function
- Observation: in $\mathbb{R}^d$ there is a sequence of compromise transitions that converges in at most n^2 steps
  - if there are 3 coalitions, there is a merge or single-agent transition

d-Hypercube

- Metric space: $\{0, 1\}^d$ with Hamming distance
  - r = (0, 0, …, 0)
- For d = 3 compromises may fail

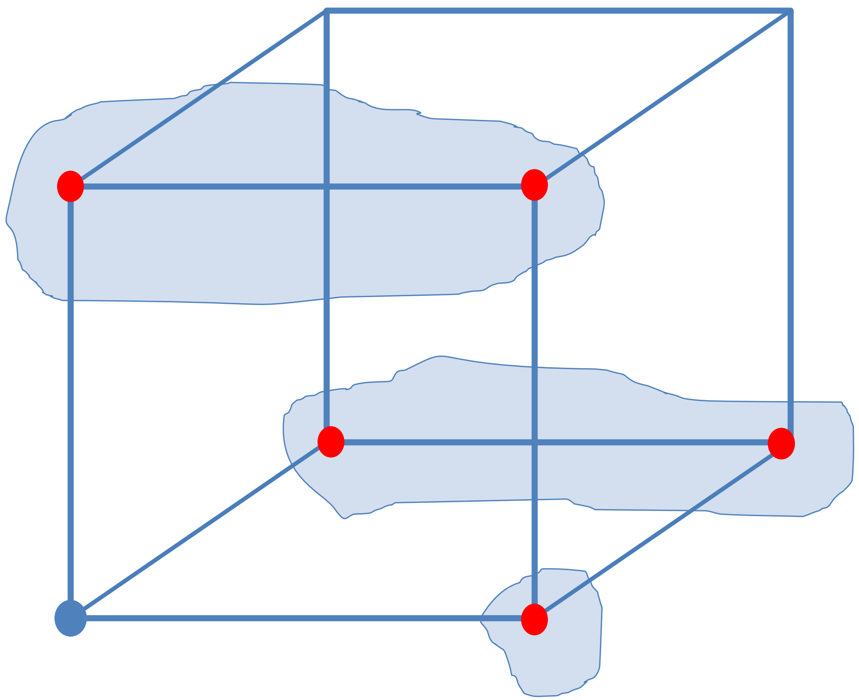

Beyond 2-Compromises?

- Suppose we allow compromises involving t coalitions ($t > 2$)
- What is the smallest value of t that guarantees success in the d-hypercube?
- $t^*(d) \leq 2^d - 1$
- For $d = 3$, we have $t^*(d) = 3$
- For $d = 4$, we have $t^*(d) = 5$
- Lower bound: $t^*(d) \geq d$
- Upper bound: $t^*(d) \leq 2^{d-1} + (d+1)/2$

Open problem: close the gap

Further Open Questions

- Are there "simple" transitions that ensure convergence when proposal space is a subset of $\mathbb{R}^d$?
- How "rich" should a space be for compromise transitions to succeed?
- Is there an explicit sequence of compromise transitions that is exponentially long?