PIK

GaNe
FutureLab on Game Theory
& Networks of Interacting Agents

MCC

Nondeterministic Proportional Consensus – theoretical and simulated properties of a novel nonmajoritarian single-winner voting method aiming at fairness and efficiency –

Jobst Heitzig, Forest W. Simmons

International Seminar on Social Choice, 8 December 2020

Leibniz-Gemeinschaft

related paper and presentation: https://www.pik-potsdam.de/members/heitzig/maxparc

Overview

- Intro
- Proportional Power Allocation
- Supporting Consensus
- MaxParC method
- Agent-Base Simulations

My Short History in Social Choice Theory

- 1998–2002 PhD pure maths, partial orders → P. Fishburn → Approval V.
- PhD supervisor → ballroom dancing → problems with "skating system"
- 1998 Nobel Memorial Prize for Amartya Sen → seminar at U. Hannover
  → **social choice became my "pet project"**
- 1998– "election methods mailing list" → met co-author Forest Simmons
- J.F. Laslier's "Tournament Solutions & Majority Voting"; S. Barbera's work
  → **interest in probabilistic methods**
- 2010 article in Social Choice and Welfare
- since 2010 at Potsdam Institute for Climate Impact Research
  → **ongoing interest in game th. & fair mechanisms for cooperation**

Problem Statement

- **Setting:** a more or less fixed group of people makes many single-issue decisions over time
- Permanent minorities might exist
  → avoid tyranny of majority
  → distribute "power" more fair than majoritarian methods
- **We focus on *one* such single-issue decision:**
  ***N* voters must choose one of *k* options**

→ *What "fair" and "efficient" single-winner voting method to use?*

Analytical Framework

(all this is challenging
for axiomatic treatment
→ *I seek collaborators!*)

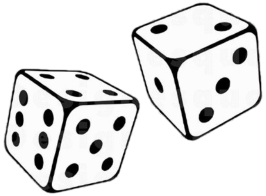

allow for probabilistic methods
→ **preferences over lotteries**
(expected-utility, prospect theory, etc.)

ballot design is part of the problem
→ **arbitrary types of ballots**
→ no canonical relationship to preferences
→ problem is *not* to "aggregate preferences"

voters may be strategic
→ **game-theoretic equilibrium concepts**
→ a voting method is a game form

bounded rationality → **agent-based modeling**

Definitions (verbal, see paper for formal)

Ballot = questionnaire asking voter for some type of data,
e.g. marking/ranking/rating one or several of the options

Voting method = function that maps
a profile of filled ballots
to a lottery of options

(effective) power of voter subgroup G
(in a certain decision problem under a certain voting method)
= largest winning probability G can guarantee any option X

Task 1 (easy): Distribute Power Proportionally

Method is *"fair"* iff power
is **proportional** to group size.

(→ In the <u>long run</u>, every voter
can get their will equally often)

Trivial solution: **Random Ballot**

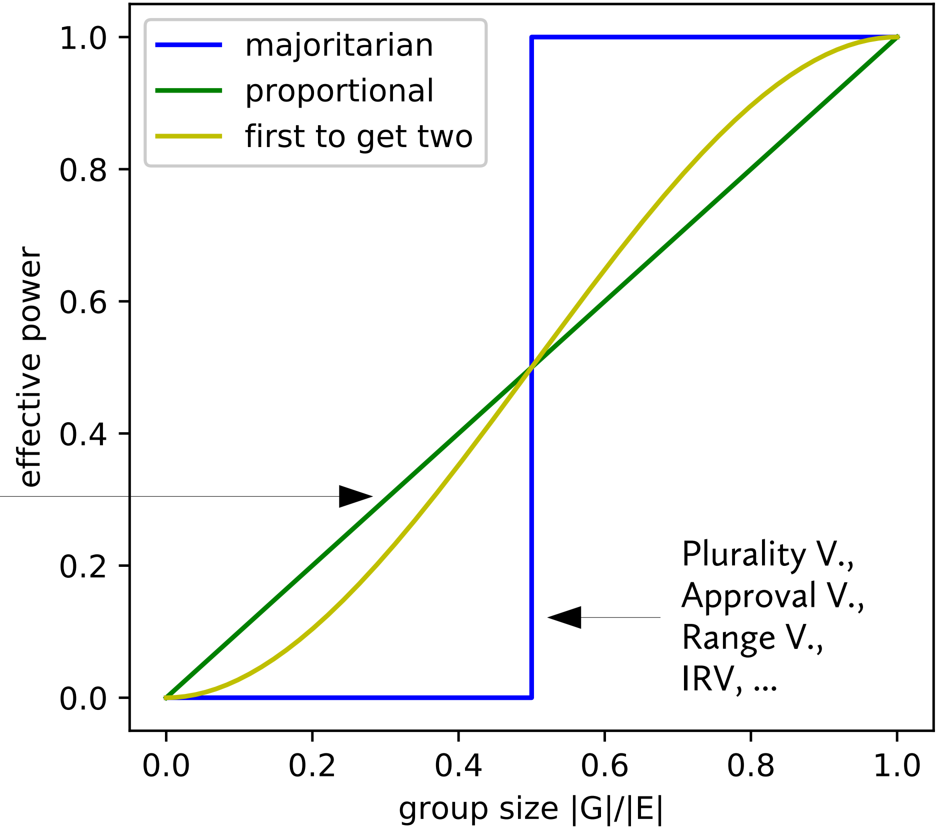

(Small exercise:
shape for Borda/Cusanus?)

Random Ballot

Voting: Each voter marks one option on their ballot
Tallying: One ballot is drawn uniformly at random, the option marked on that ballot wins

Some potentially desirable **properties**:

- anonymous & neutral
- monotonic (more marks → larger chance)
- Pareto-efficient (if all prefer Y to X, X will have zero probability)
- strategy-free (marking your favourite is a dominant strategy)
- ~~deterministic (use chance only in case of ties)~~
- simple to vote in and to tally
- distributes effective power proportionally
- ~~supports consensus~~
- ~~produces high "welfare"~~
- ~~reveals voters' detailed preferences~~

Game Theory

Welfare Theory, Behavioural Economics,
Agent-Based Modeling

Task 2 (still easy): Support Full Consensus

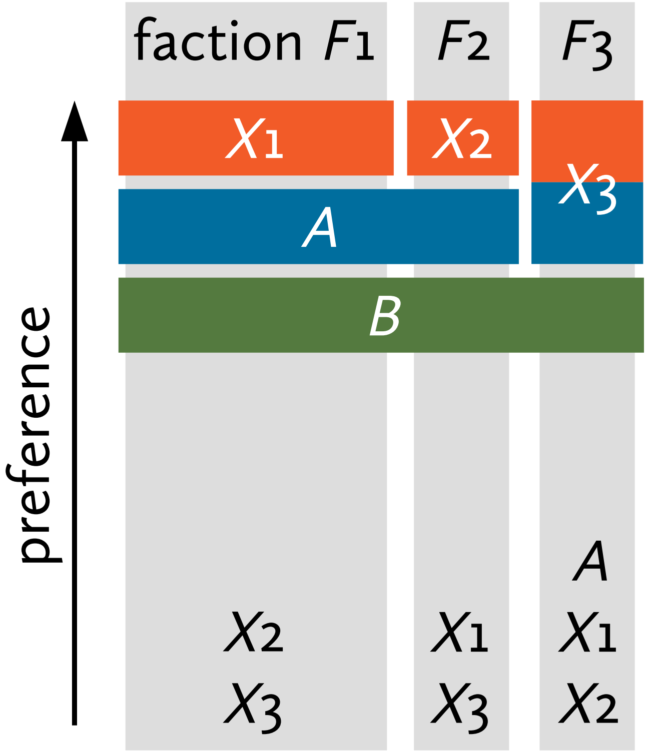

benchmark lottery = result of Random Ballot

potential full consensus = any option that is Pareto-better than benchmark lottery (here: B)

potential partial consensus for subgroup G = potential full consensus if problem restricted to G (here: A if $G=F1+F2$)

Goal: "make" B win for sure if available, otherwise "make" A win with probably $|F1+F2|/N$

Simple solution: the "Two Urns" method (Heitzig & Simmons 2010)

(but impractical)

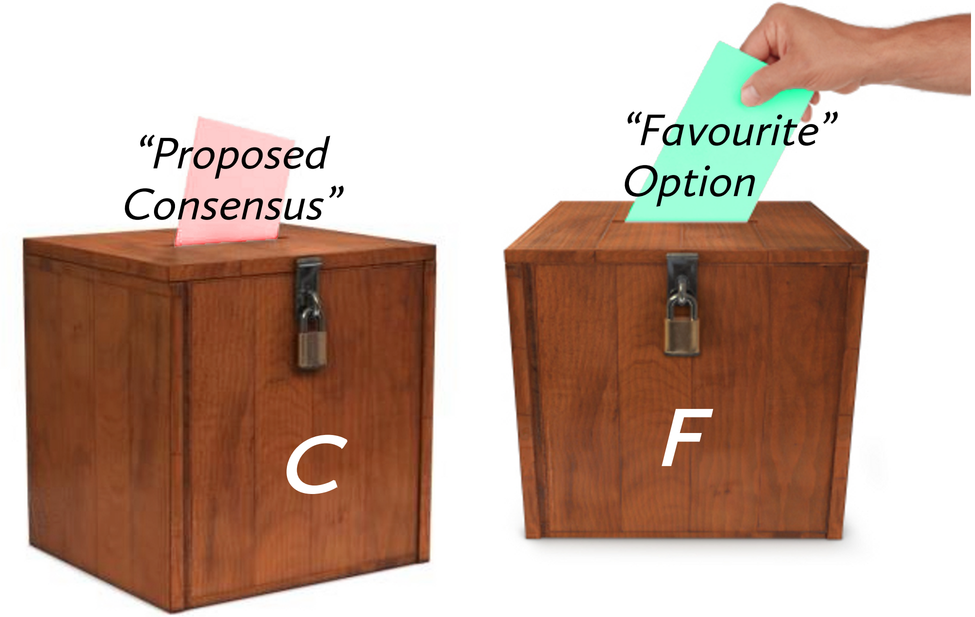

Voting: Each voter puts one standard ballot into urn C and one into urn F.

Tallying: If all ballots in urn C name the *same* option, that option wins; otherwise, the option named on a randomly drawn ballot from urn F wins.

Properties: anonymous, neutral, monotonic, Pareto-efficient, ~~strategy-free~~, simple, distributes power proportionally, **supports full** ~~(& partial)~~ **consensus**, ~~produces high "welfare"~~, ~~reveals detailed preferences~~.

(but a version with 3 urns does)

Task 3: Make it Work with Larger Electorates

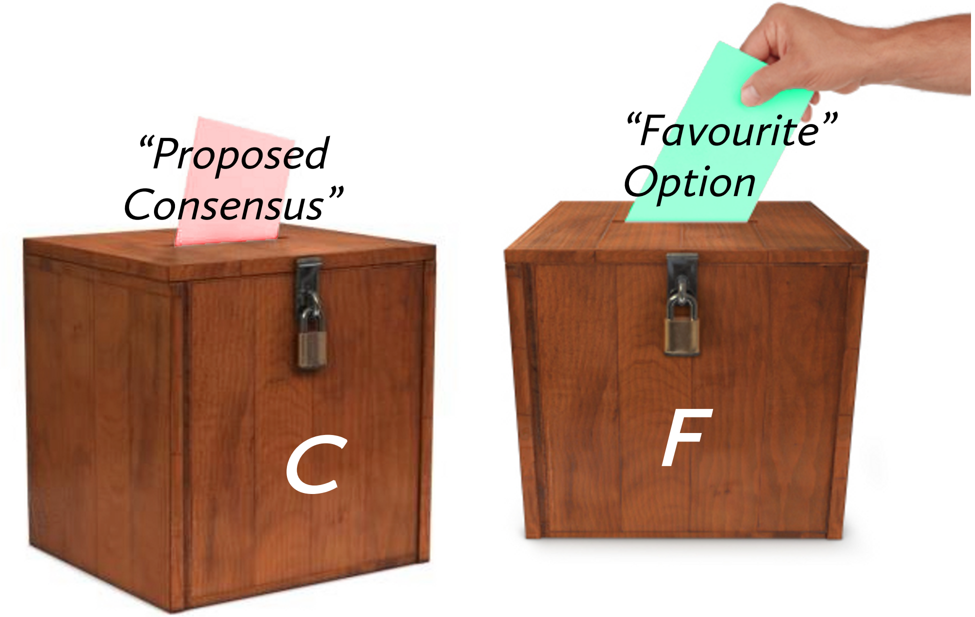

Problem 1: In large electorate, unanimity very unlikely

→ replace by large supermajority?

→ small minorities get zero power

→ violation of proportionality

Problem 2: Several competing potential consensus options

→ coordination problem

→ get help from Approval Voting?

The "Three Urns" Method
(still impractical)

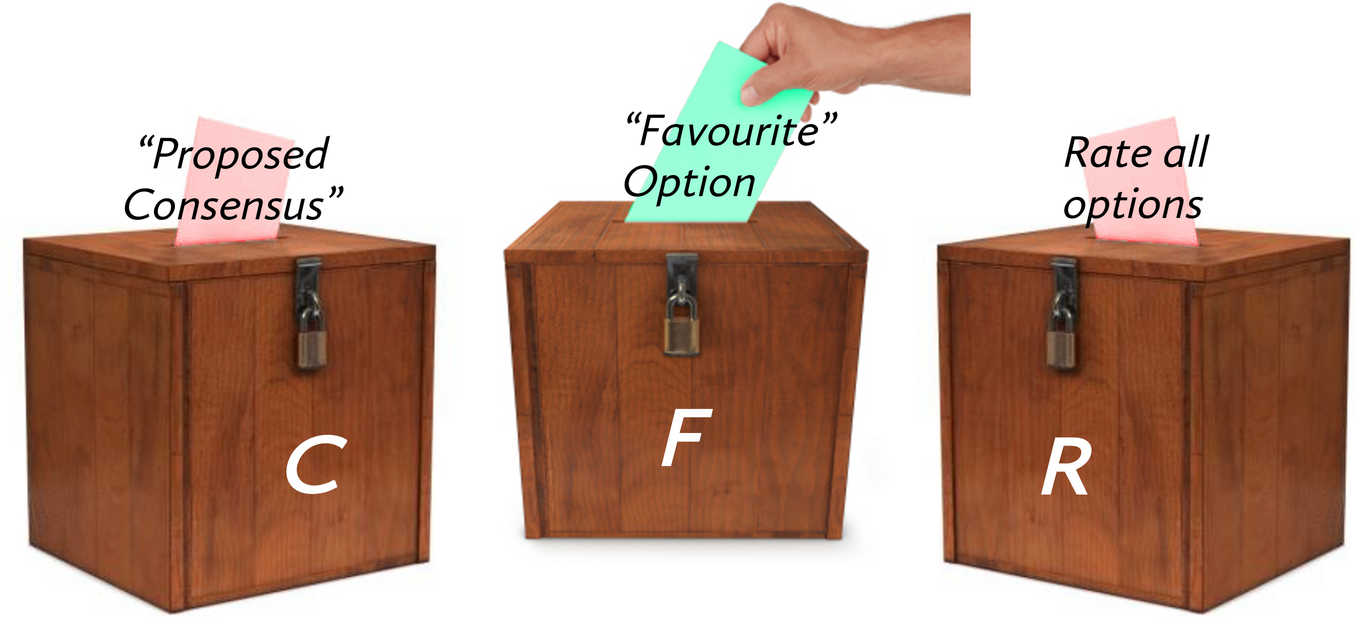

(Heitzig & Simmons 2010)

1) draw option X from C urn
2) let L be the lottery of drawing from F urn
3) if all ballots in R rate X above L, X wins, else apply L

- solves problem 2 (coordination): it is optimal to mark favourite potential consensus on your C ballot
- R reveals true preferences
- still does not solve problem 1 (unanimity)

Idea 1: Mix Random Ballot & Approval Voting

"Conditional Utilitarian" Method (Duddy 2015, Aziz et al. 2019)

- use approval ballots & tally the approval scores
- draw a ballot at random,
  and from the options approved on it, elect the highest-scoring one

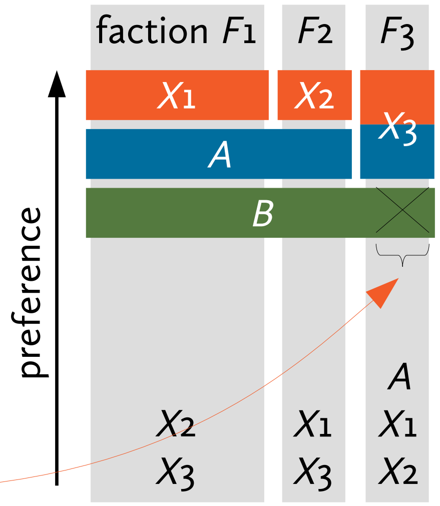

Pros:

- solves *coordination* problem for "sincere" voters, even for partial consensus
- fulfills a non-strategic version of proportionality

Problem: gives incentives to disapprove potential consensus option!
→ **will not elect potential consensus option for sure with strategic voters**
→ ***cooperation* problem**

Idea 2: Add some Conditional Commitments

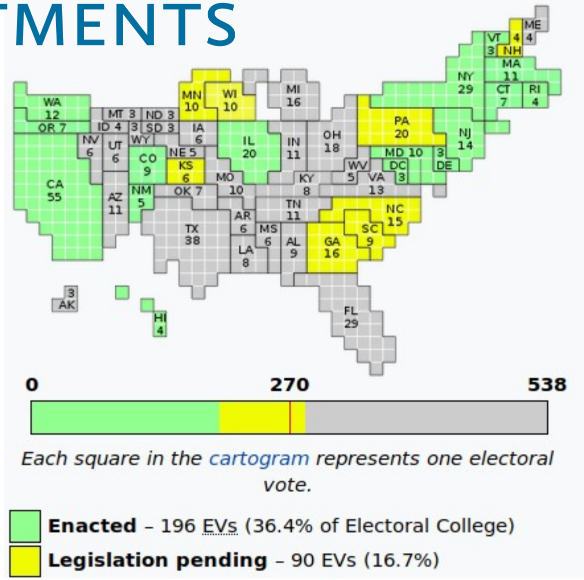

(Wikipedia)

Inspiration 1:
National Popular Vote Interstate Compact

Many US federal states have <u>committed</u> unilaterally to
'make electors elect winner of national popular vote
<u>if</u> enough other states do so as well
→ clear threshold of 270 electors

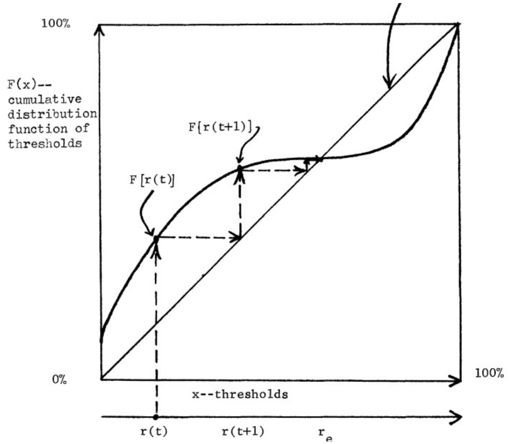

Inspiration 2:
Marc Granovetter's "threshold model" of social mobilisation (Granovetter 1978, famous in sociology)

Each person has an individual <u>threshold</u> for getting "active"
in terms of how many others are already active

Maximum Partial Consensus (MaxParC)

(Heitzig & Simmons 2020, arXiv:2006.06548)

- Voter basically says:
  if at least x% approve of A,
  I will approve of A
- Equivalent: specify rating *r*, automatically approve iff *r* + approval score > 100
- Solve this recursive, endogenous definition of "approval"
- Finally use Conditional Utilitarian rule

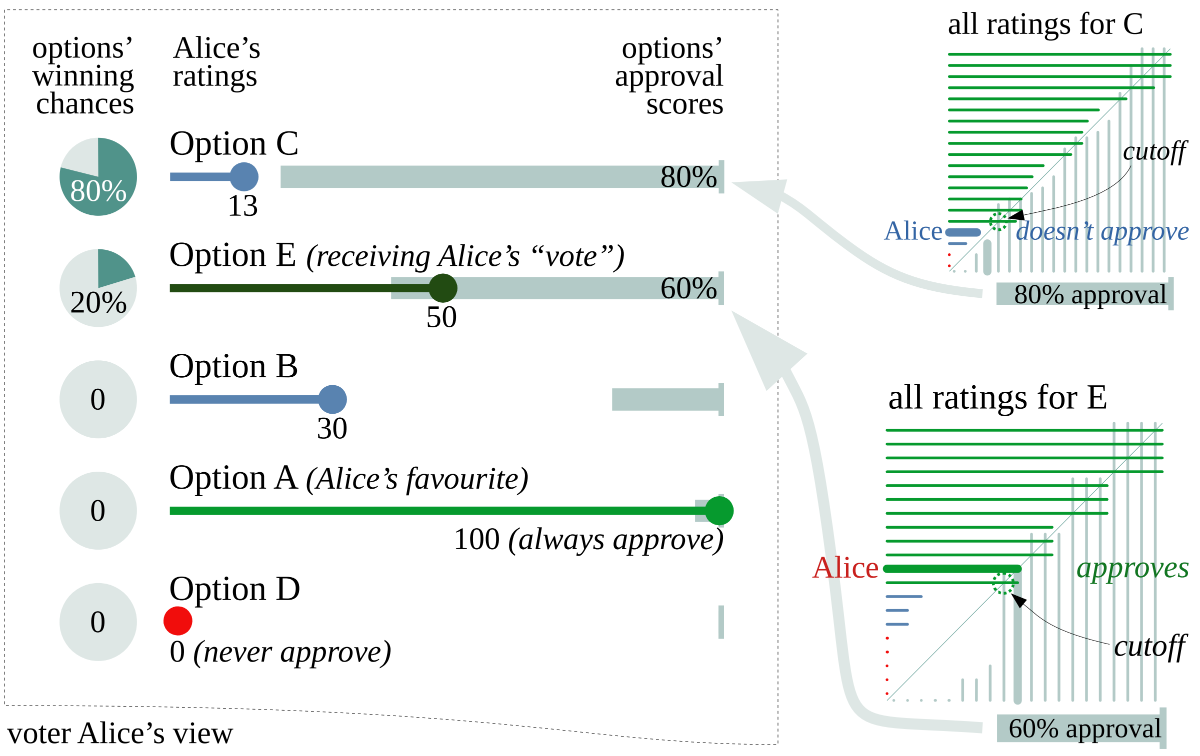

Maximum Partial Consensus (MaxParC)

- "Random ballot" ingredient guarantees proportionality
- "Approval" ingredient solves coordination problem
- Conditional commitment ingredient solves cooperation problem
  → full and partial consensus in strong forms of game-theoretic equilibrium
  (see paper for details)

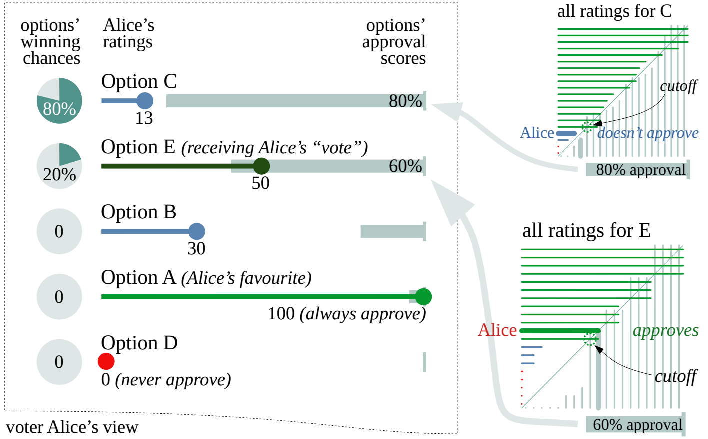

Mission accomplished? What about:

- Ease of voting / tallying?
- Resulting randomization & "efficiency"?
- Monotonicity, clone-proofness, manipulability, preference revelation, …?

OTHER *"Nondeterministic Proportional Consensus"* METHODS?

MaxParC's approach: conditional commitments to approve

Alternative: **"automatic" bargaining about winning probabilities**
→ Nash bargaining solution → the "**Nash Lottery**"
(= use "Nash product rule" to distribute winning prob.)

- allocates power proportionally
- supports full/partial consensus in equilibrium
- other nice properties

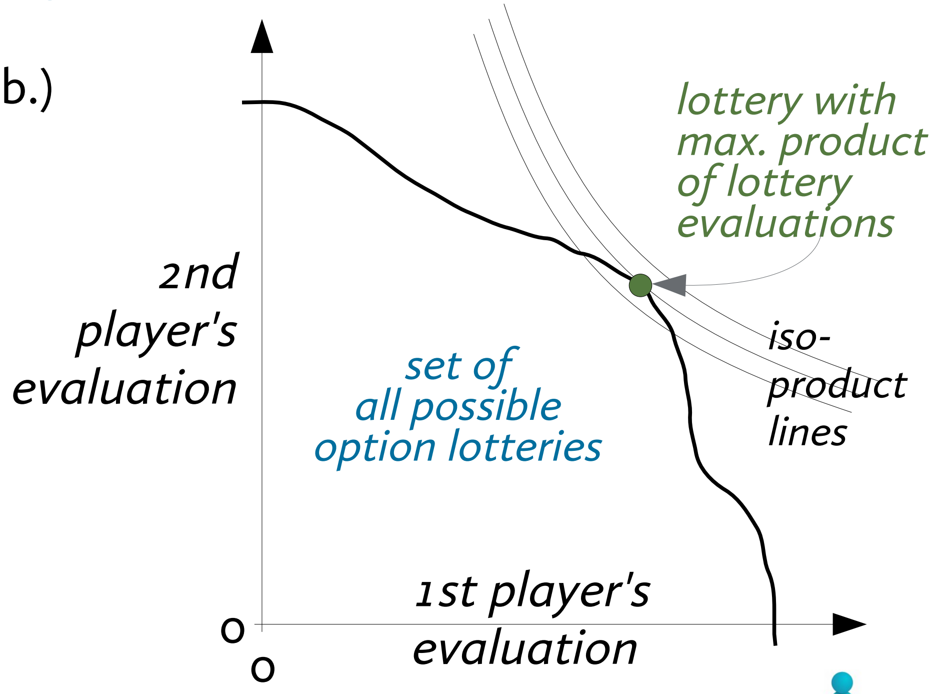

Which would make a better "standard" decision method, MaxParC or Nash Lottery?

Comparison of "Qualitative" Properties:

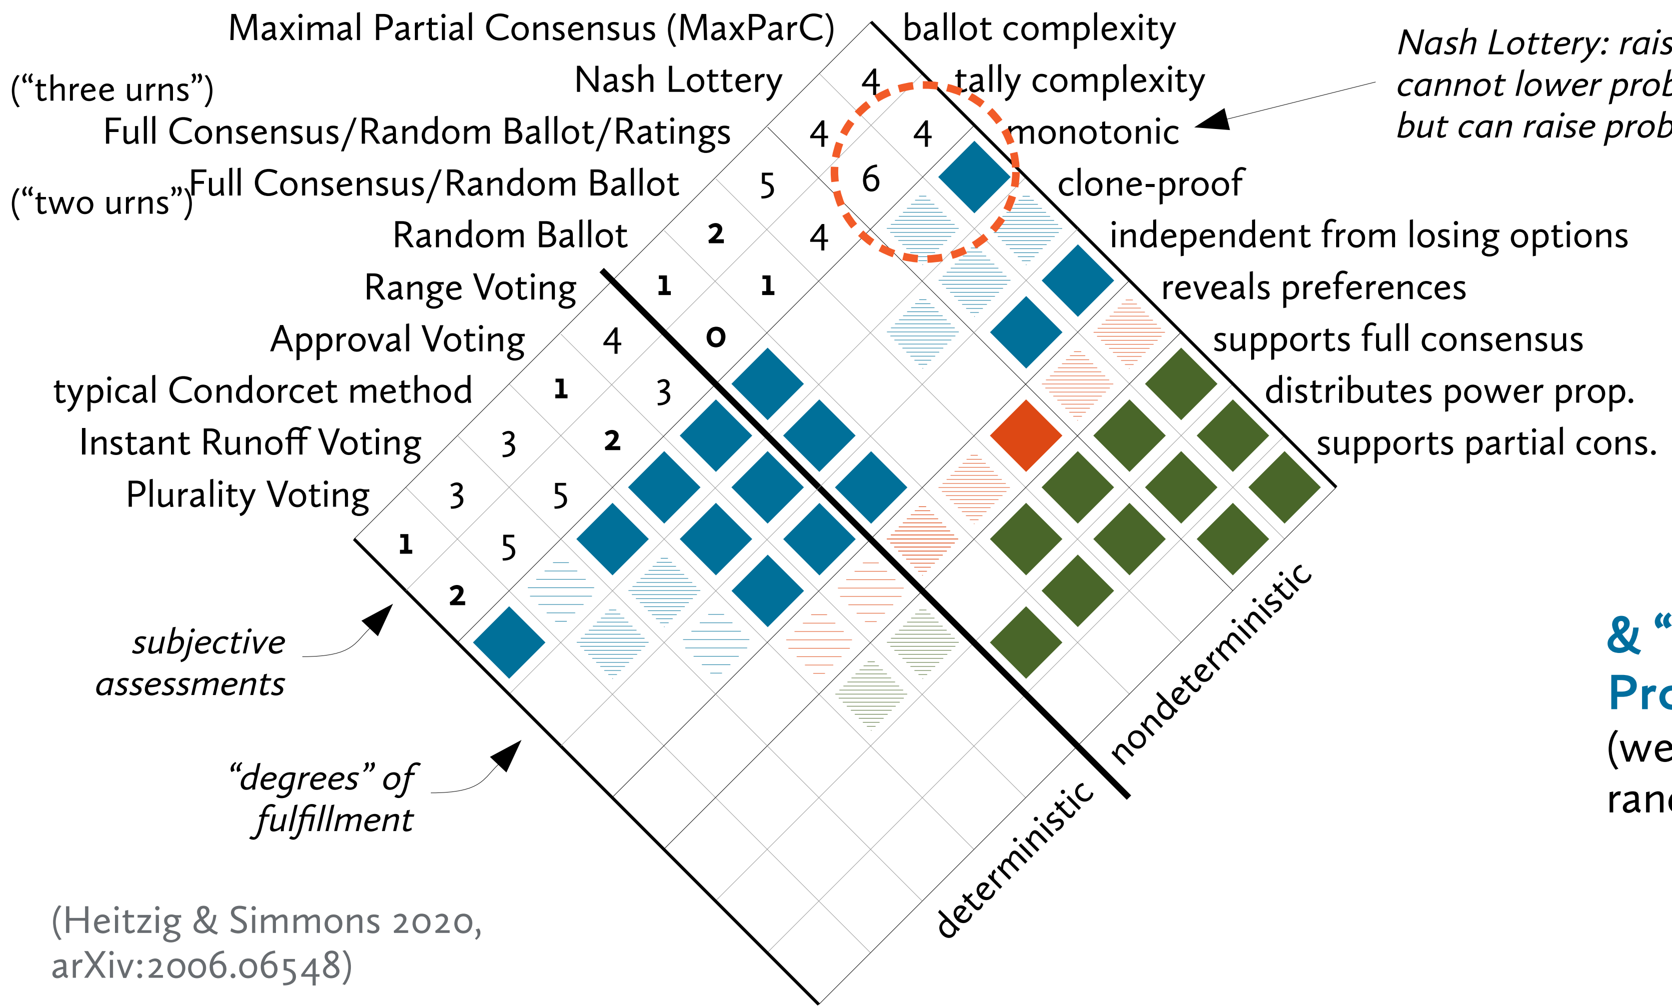

Nash Lottery: raising rating for X cannot lower prob. for X but can raise prob. for Y

& "Quantitative" Properties?
(welfare, inequality, randomness, ...)

(Heitzig & Simmons 2020, arXiv:2006.06548)

Ex-Ante Approach to Measuring “Welfare” Effects

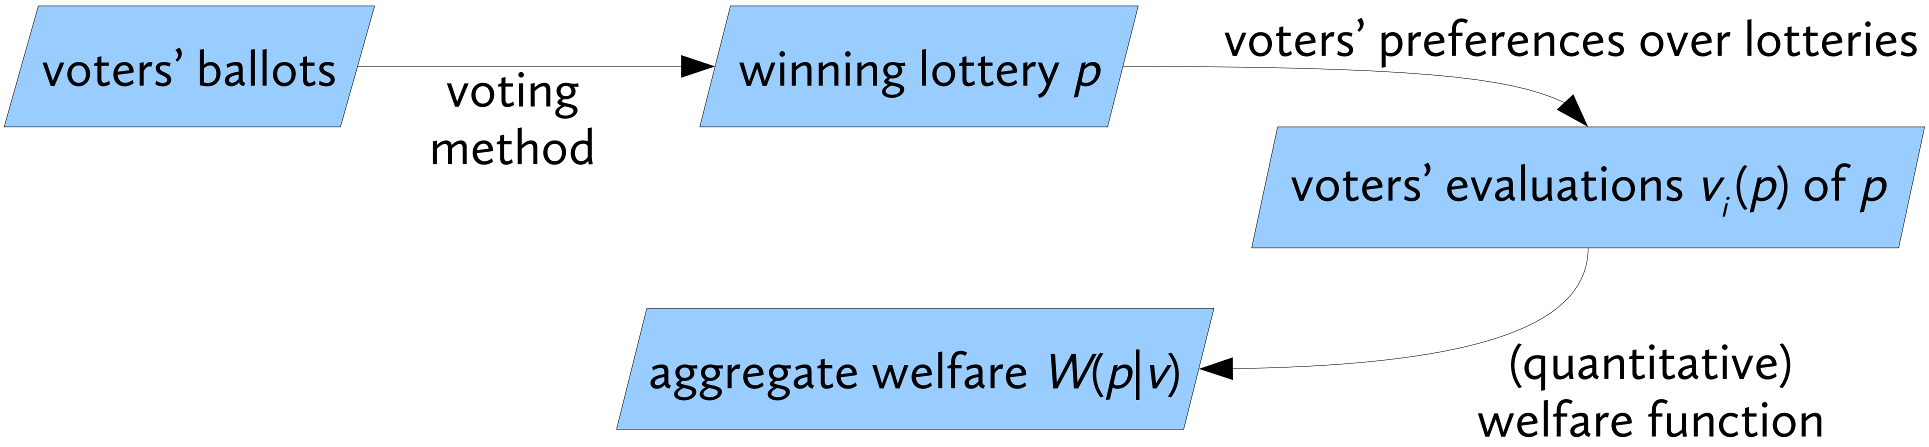

here: $W(p|v)$ =

- $\Sigma_i\, v_i(p)\;/\;N$ (Utilitarian welfare function)
- $\Sigma_i \Sigma_j \min[v_i(p), v_j(p)]\;/\;N^2$ (Gini-Sen)
- $\min_i u_i(p)$ (Egalitarian)

Voter Heterogeneity

Voters evaluate *options* according to their **preferences over options**

- spatial theory of voting (→ political science, e.g. Carroll et al. 2013)

Voters evaluate *lotteries* depending on their **risk-attitudes**

- ~20% rather conform to expected utility theory,
  ~80% rather conform to cumulative prospect theory (→ behavioural economics, e.g. Bruhin et al. 2010)

Voters have different (boundedly rational?) **voting behaviours**

- "sincere", fully strategic, heuristic, using trial and error, "lazy", …

These heterogeneities call for behavioural experiments (I couldn't do that yet)
or ***Agent-Based Modeling*** (aka In-Silico Voting Experiments, e.g. Laslier 2010)

Agent-Based Modeling

- Represent decision makers by individual **agents** with heterogeneous **attributes**
- Simulate what they do from **time step** to time step by programming individual behavioural **rules**

Here:

- agent = voter
- attributes: preferences, risk-attitude type, behavioural type
- rule = how the agent votes, maybe depending on others' attributes and observed earlier behaviour

Simulated Decision Procedure

1. Agents are told what the options are and form their preferences
2. In several polling rounds, they can express approval and support for options and see the poll's results
3. In a main voting round, they all vote simultaneously

Optionally:

4. In an interactive phase until some deadline, they can iteratively adjust their votes in reaction to others to improve the result

   (since such an interactive phase may become a crucial design element of online voting systems aiming at consensus)

Preference Models used

voter option

- "Uniform" (similar to "impartial culture") $u_i(x) \sim \mathrm{Unif}([0,1])$

"individuality" parameter

- "Block model" (BM) $$u_i(x) = \underbrace{U_{J(i)}(x)} + \iota \underbrace{\varepsilon_i(x)}$$

i's "voter block"

iid ~ N(0,1)

- Spatial models:
  - "Gaussian allotment" (GA) $u_i(x) = e^{-\|\eta_i - \xi_x\|_2^2/2\sigma_x^2} / (\sqrt{2\pi}\sigma_x)^d$
  - "Quadratic allotment" (QA) $u_i(x) = -\|\eta_i - \xi_x\|_2^2/2\sigma_x^2 - d\ln(\sqrt{2\pi}\sigma_x)$
  - "Linear allotment" (LA) $u_i(x) = -\|\eta_i - \xi_x\|_1/\sigma_x - d\ln(2\sigma_x)$

voter position

option position

option "broadness"

policy space dimension

Example in a two-dimensional policy space

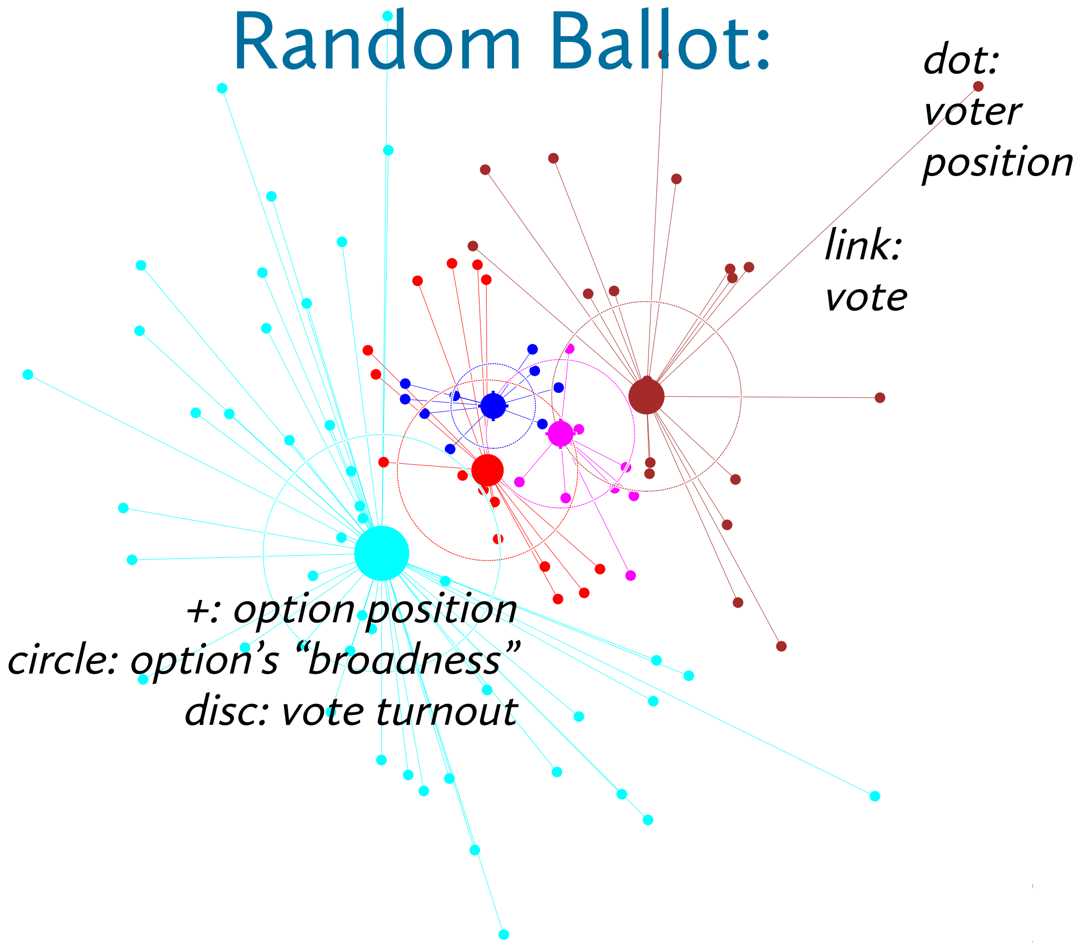

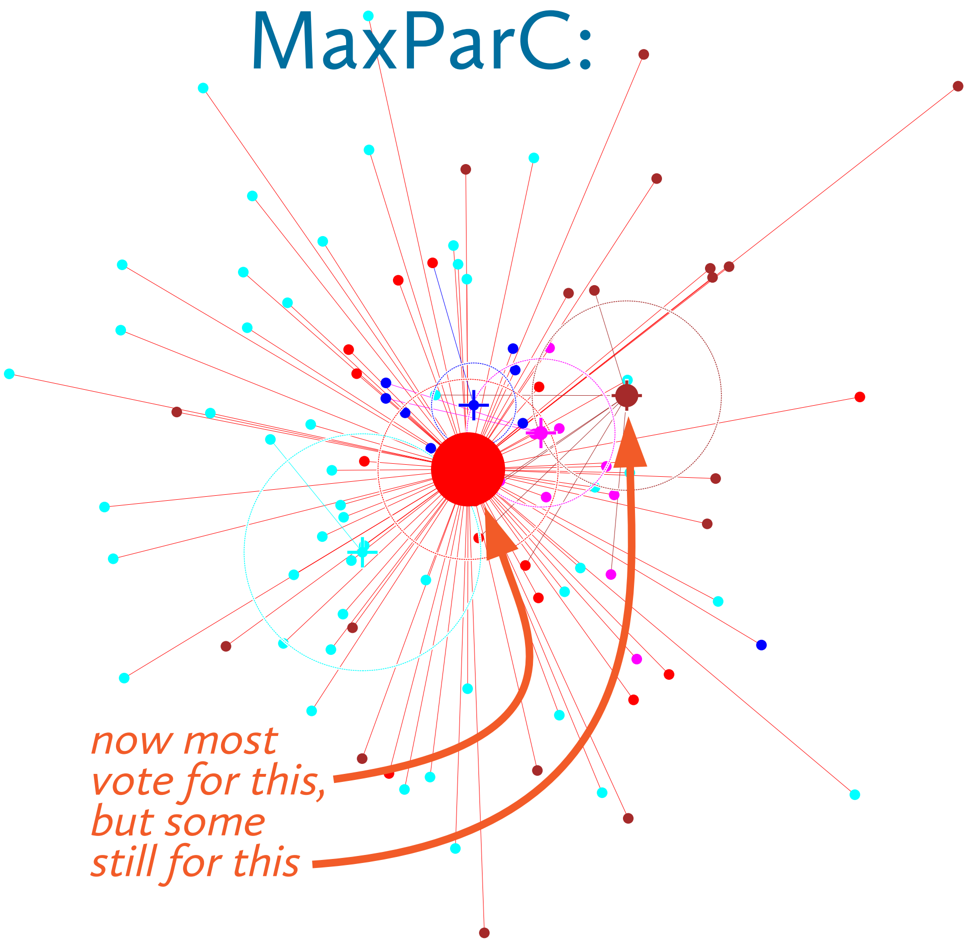

Simulated Behavioural Patterns

E.g. for voting method = Approval Voting:

- **Lazy voters:** *approve of favourite and no other option*
- **Sincere voters:** *approve of what you prefer to the benchmark lottery (as estimated by polling results)*
- **Heuristic voters:** *approve of all options you prefer to the option leading the polls, & approve of that one if you prefer it to the runner-up*
- **Trial-and-error:** *start like heuristic voter; during interactive phase, always pick a random option, then change your approval of it if you profit from that change*
- **Factionally strategic:** *start like heuristic voter; during interactive phase, always switch to your faction's best response to the other factions' current votes*

Similar for other voting methods (details differ due to ballot & tallying differences)

Monte-Carlo Experiment Design

Simulations: Large ensemble (>2.5 mio. runs)
with **broadly varying parameters**:

- no. of: options 3–9, voters 9–999, polling rounds 1–10
- preference model, 2–9 blocks / 1–3 policy space dimensions, varying voter position heterogeneity, option broadness heterogeneity, distance-to-utility conversions
- varying population mixtures of
  - risk-attitudes (expected utility + two forms of cumulative prospect theory)
  - behavioural types (lazy, sincere, heuristic, trial and error, factionally strategic)
- 5 deterministic majoritarian + 5 probabilistic proportional methods, with or without interactive phase

Monte-Carlo Experiment Design

Simulations: Large ensemble (>2.5 mio. runs) with **broadly varying parameters**

Output:

- Several **aggregate welfare/satisfaction/entropy** metrics

Analysis:

- **Descriptive statistics** for these metrics (overall, grouped by single parameters)
- **Multivariate regression** analysis to identify influence of parameters and voting method

Selected Results

- Welfare costs of achieving fairness and supporting consensus exist but are much smaller than the inequality produced by majoritarianism
- MaxParC clearly outperforms the other four proportional methods and under some conditions also the majoritarian methods
- All proportional methods lead to considerable entropy, MaxParC the least
- Strategic voters have only negligible advantage over lazy & heuristic voters
- Among all parameters, the preference model has the strongest effect most results

final Gini-Sen welfare by method and preference model

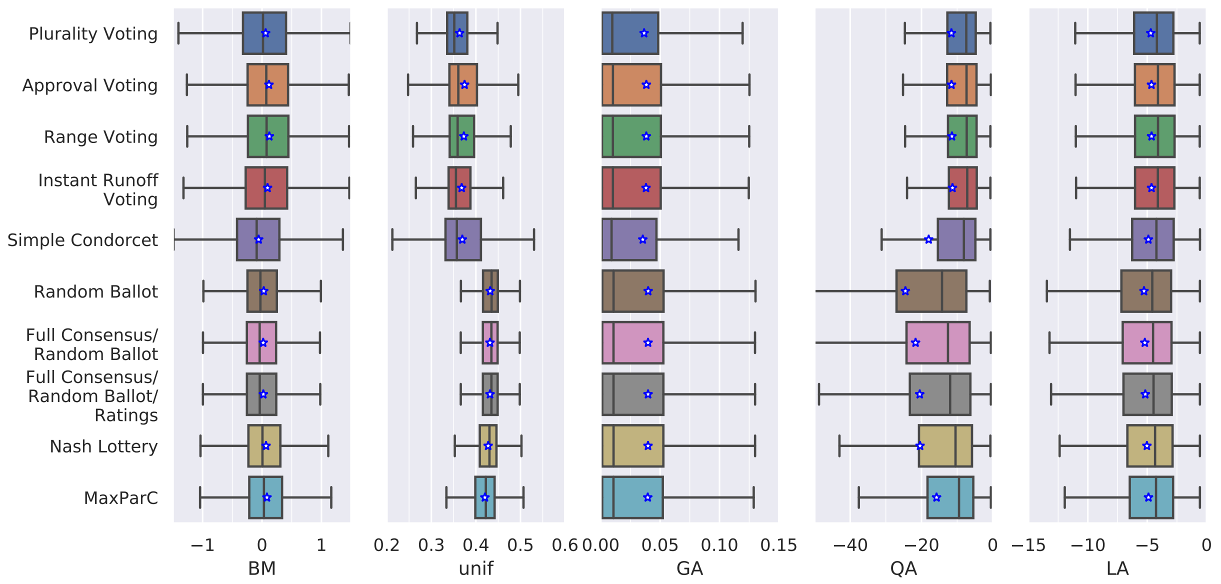

Summary

- Nondeterministic proportional voting methods
are fairer than deterministic majoritarian methods
and can support full and partial consensus
even in strategic contexts

- In theoretical analyses, they perform well
in terms of other desirable qualitative properties

- In agent-based simulations,
they do not systematically perform worse
than deterministic methods
in terms of quantitative properties

Next Steps

- Get this into review
  → *journal suggestions?*
- Develop a social voting app
  for mobile phones: www.vodle.it
- In-depth game-theoretical analysis
  for generic preference profiles
- Lab experiments w/ Elke Weber &
  Sara Constantino (Princeton U.)
- Axiomatic characterizations?
  → *anyone interested?*

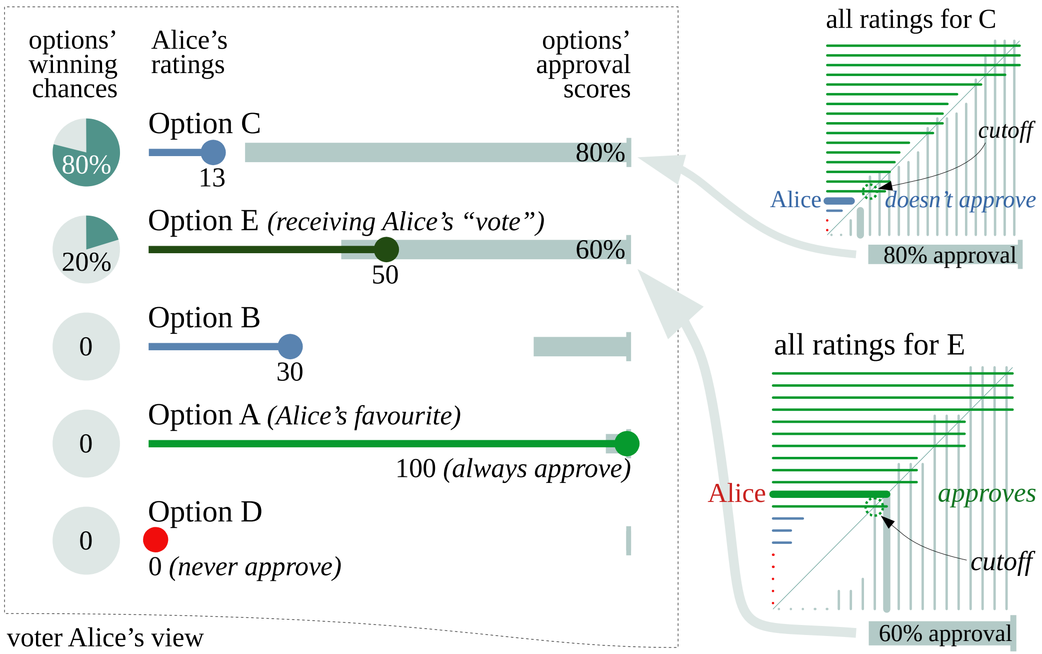

Thank you!
→ *Questions? Comments?*

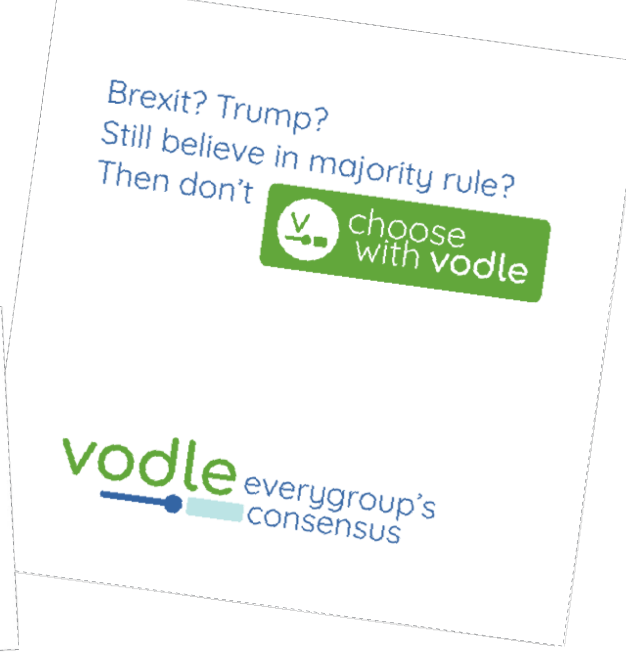

www.pik-potsdam.de/research/futurelabs/gane

related material: www.pik-potsdam.de/members/heitzig/maxparc

prototype of related voting app: www.vodle.it

Backup Slides

(Sideline: Power leads to Responsibility)

Sarah Hiller's (hiller@pik-potsdam.de) PhD project on
formalizing ethical responsibility in multi-agent situations with uncertainty

→ Joint paper on responsibility in social choice situations:
Heitzig & Hiller 2020, in revision, arXiv: XXX

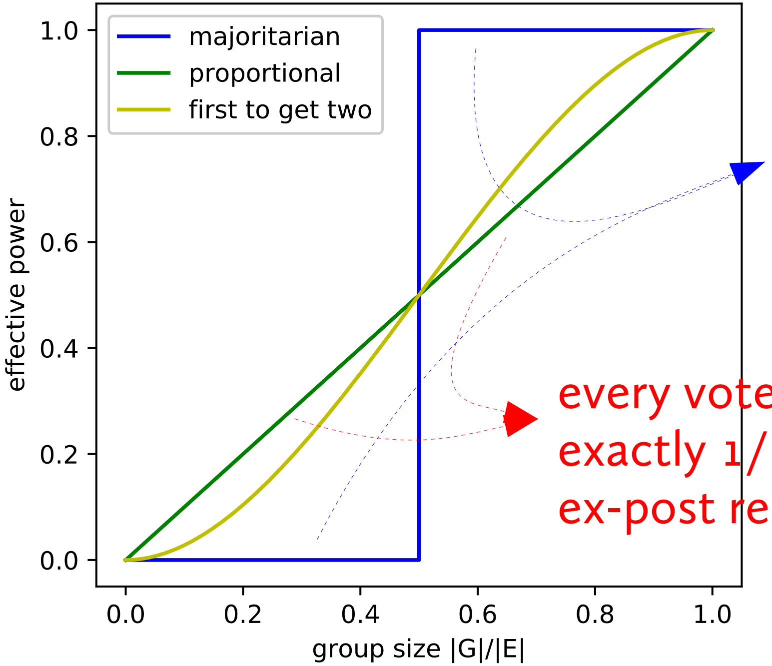

the majority has full responsibility as a group, no single voter has any ex-post responsibility unless the decision was ~fifty-fifty

every voter has always exactly 1/N ex-ante and ex-post responsibility

“Supporting Consensus” in formal voting methods (2)

Def. (vague): (Heitzig & Simmons 2010)
A method ***supports full consensus*** iff in “typical” situations where a potential full consensus exists, the “natural” strategic equilibria of the resulting voting game will result in such a full consensus being chosen for sure.

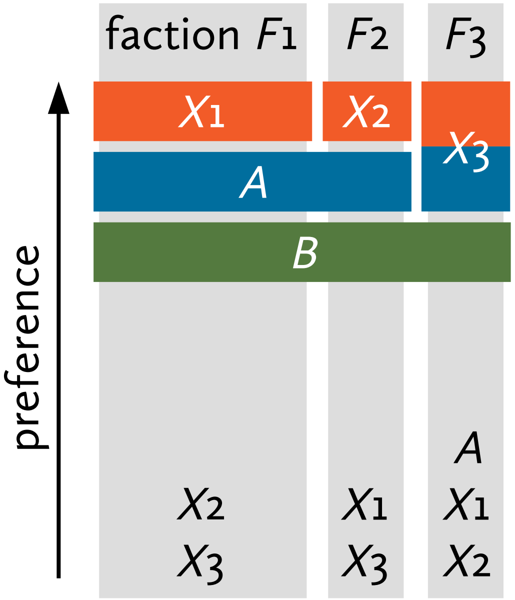

In the example:
Option *B* must be chosen in equilibrium

Note that for some voting rules (e.g. Approval Voting), sometimes not even a single equilibrium exists!

“Supporting Consensus” in formal voting methods (3)

Def. (vague): (new paper Heitzig & Simmons 2020, about to be submitted)
A method ***supports partial consensus*** iff in “typical” situations where a potential partial consensus for some group G exists, the “natural” strategic equilibria of the resulting voting game will result in such a partial consensus being chosen with probability at least $|G|/N$.

In the example:
If option A but <u>not</u> option B exists,
option A must be chosen with at least 75% probability
in equilibrium.