Combining theoretical analysis & agent-based modeling *in the quest for* an inclusive, fair & efficient group decision method

Jobst Heitzig, joint work with Forest W Simmons
PIK, May 2020

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

A Social App for Group Decision Making

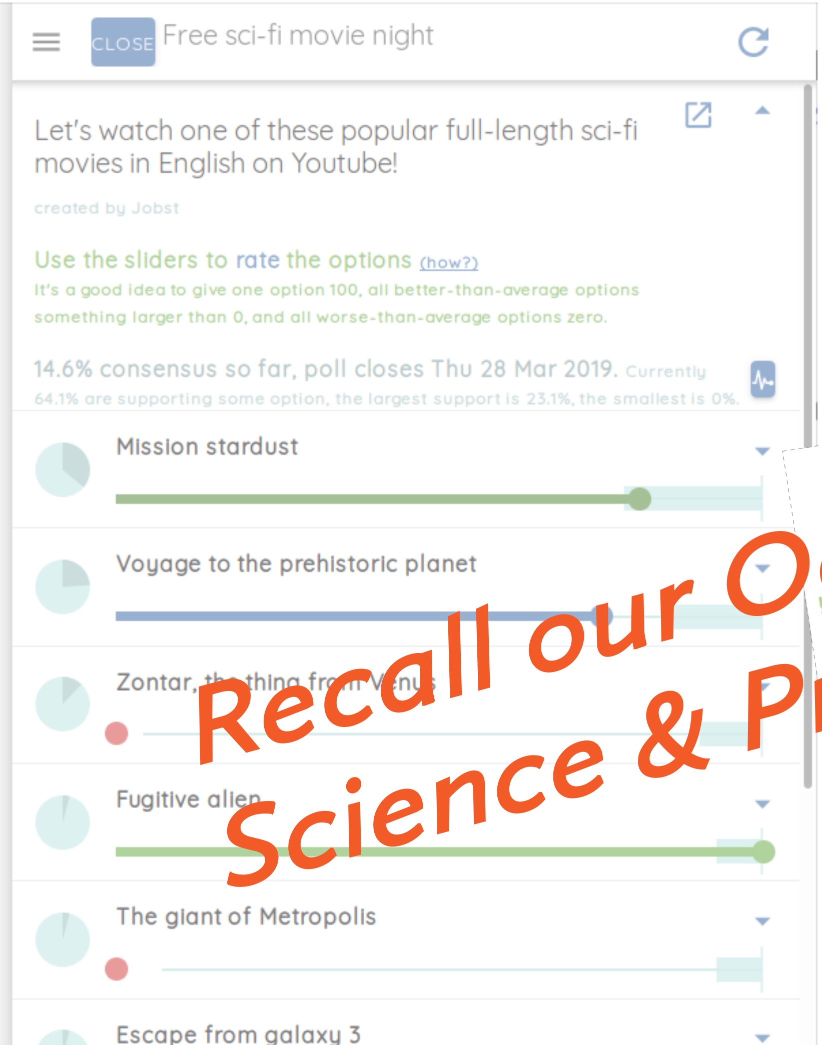

- based on conditional commitments
- uses chance to incentivize consensus

Jobst's pet project (*help needed!!*)

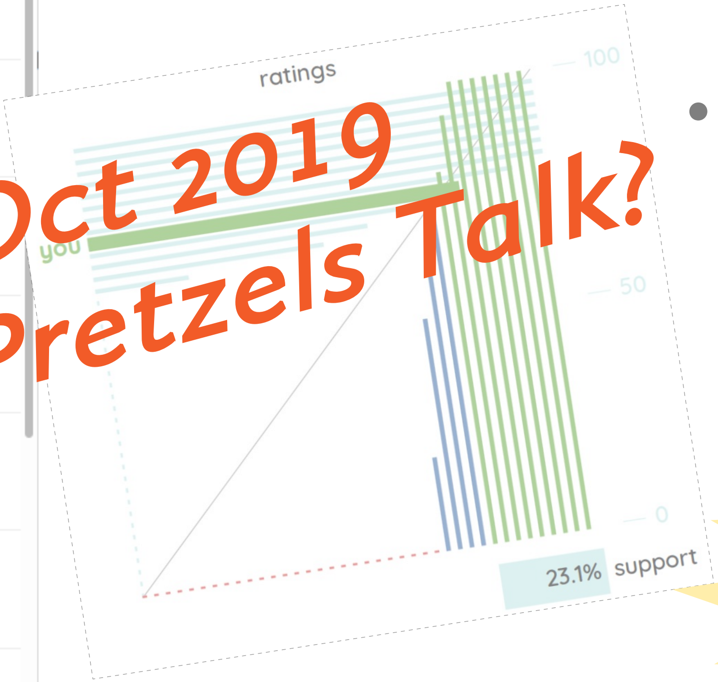

Recall our Oct 2019 Science & Pretzels Talk?

- distributes effective decision power proportionally (in contrast than majoritarian rules)

check it out on **vodle.it**

FutureLab Themes
& Today's Question

Mechanisms
for Cooperation

Multilevel
Governance

Learning &
Decision-
Making

Are there voting methods which

- *give everyone, including minorities, an equal share of effective power even if voters act strategically,*
- *promote consensus rather than polarization and inequality,*
- *do not favour the status quo or rely too much on chance?*

Social
& Ethical
Aspects

PIK MCC GaNe

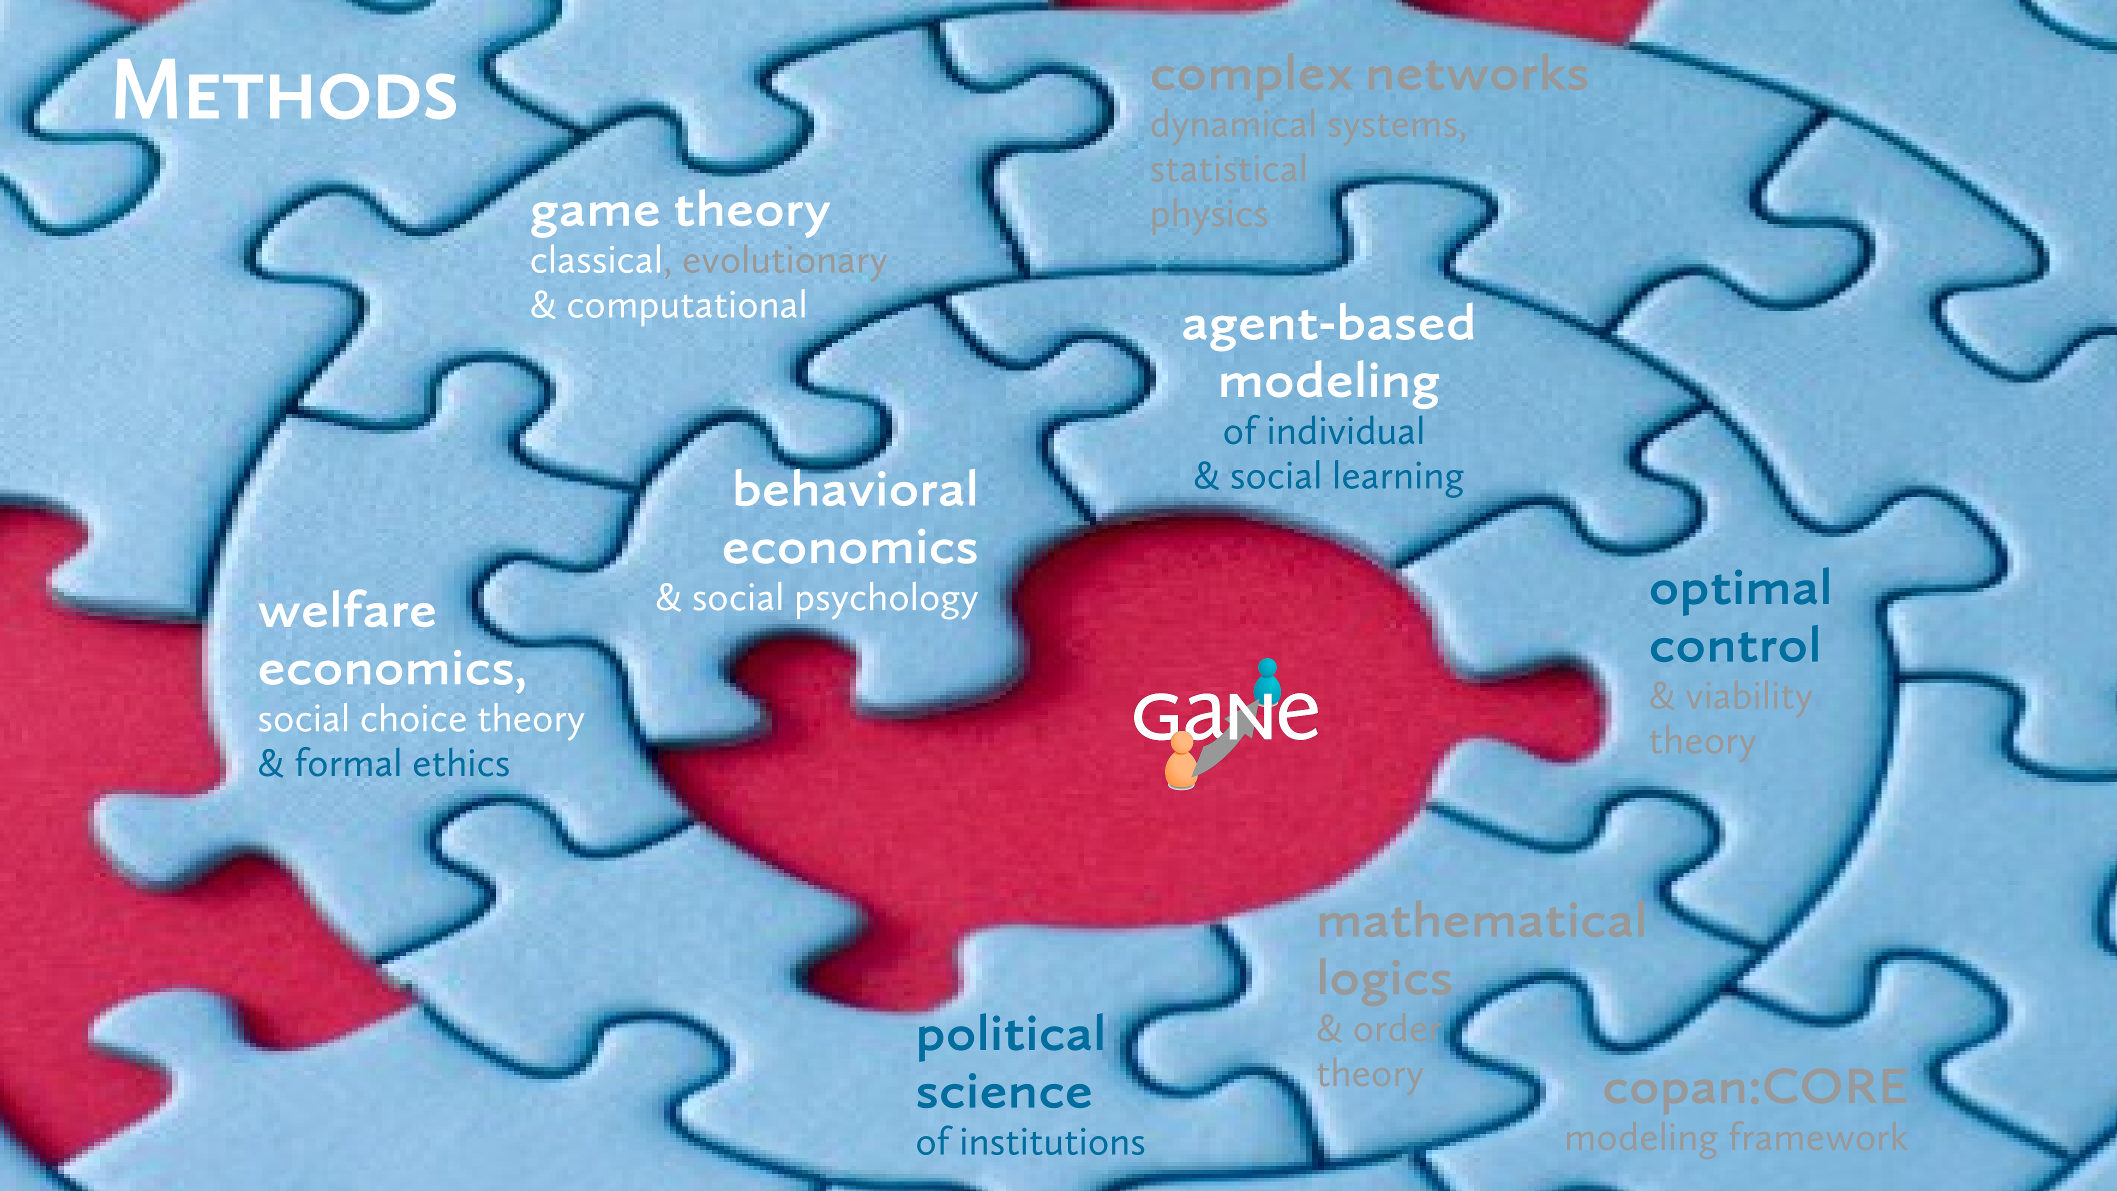

Methods
game theory
classical, evolutionary
& computational
complex networks
dynamical systems,
statistical
physics
agent-based
modeling
of individual
& social learning
behavioral
economics
& social psychology
welfare
economics,
social choice theory
& formal ethics
GaNe
optimal
control
& viability
theory
mathematical
logics
& order
theory
political
science
of institutions
copan:CORE
modeling framework

The Problem

Majority Rule (cornerstone of democracy?)

→ "Tyranny of the Majority" (Tocqueville, Lewis 2013)

→ separatism, violent conflict (e.g. Collier 2014, Cederman 2010)

Existing solutions?

- **Proportional representation?** (e.g. Cohen 1997, Cederman 2010)

  → If reps use majority rule to decide, problem remains (e.g. Zakaria 1997)

- **Consensus finding?**

  → Difficult in *strategic* contexts (e.g. Davis 1992)

  → blocking → majority's will (or status quo)

  → *effectively* majoritarian (like almost all voting methods)

Social Choice Theory?

May's Theorem (May 1952)

Majority Rule only method that satisfies some "natural" requirements

→ a mere 51% can make all decisions,
minorities have zero *effective* decision power

But: this applies only to *deterministic* methods
(which apply chance only to resolve rare ties)

- *non-deterministic* methods can distribute effective decision power differently, e.g. *proportionally*

Trivial Example: the "Random Ballot" Method

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" (treats all voters the same)
- neutral (treats all options the same)
- monotonic (more marks → larger chance)
- Pareto-efficient (if all prefer *Y* to *X*, *X* will have zero probability)
- strategy-free (it is always optimal to mark your favourite)
- ~~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~~

typically studied in Social Choice Theory

not so often studied in Social Choice Theory

Trivial Example: the "Random Ballot" Method

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" (treats all voters the same)
- neutral (treats all options the same)
- monotonic (more marks → larger chance)
- Pareto-efficient (if all prefer Y to X, X will have zero probability)
- strategy-free (it is always optimal to mark your favourite)
- ~~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~~$^{\Delta}$

typically studied in **Social Choice Theory**

Game Theory

Welfare Theory, Behavioural Economics, **Agent-Based Modeling**

Effective Power in single-winner voting

Situation: A group E of voters uses a formal, potentially non-deterministic voting rule R to make a single decision with a single winner (= pick one from a finite menu of distinct options).

(pre-filtered by a legal system)

Def.: The ***effective (ex-ante) decision power*** of a subgroup $G \subseteq E$ under rule R is the largest winning probability that G can guarantee any option of their choice, regardless of what the other voters do.

(Heitzig & Simmons 2010)

R is ***fair*** iff power is **proportional** to group size.

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

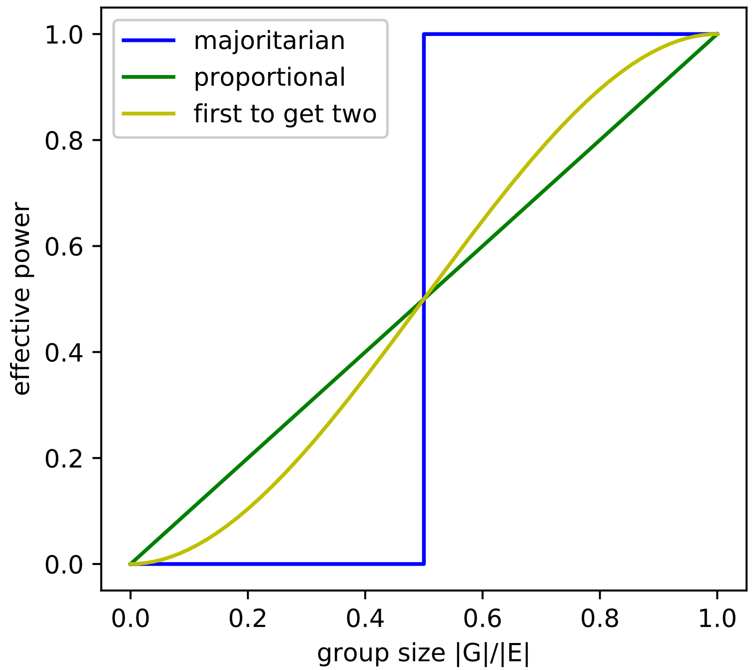

(Sideline: Power leads to Responsibility)

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 review (manuscript available upon request)

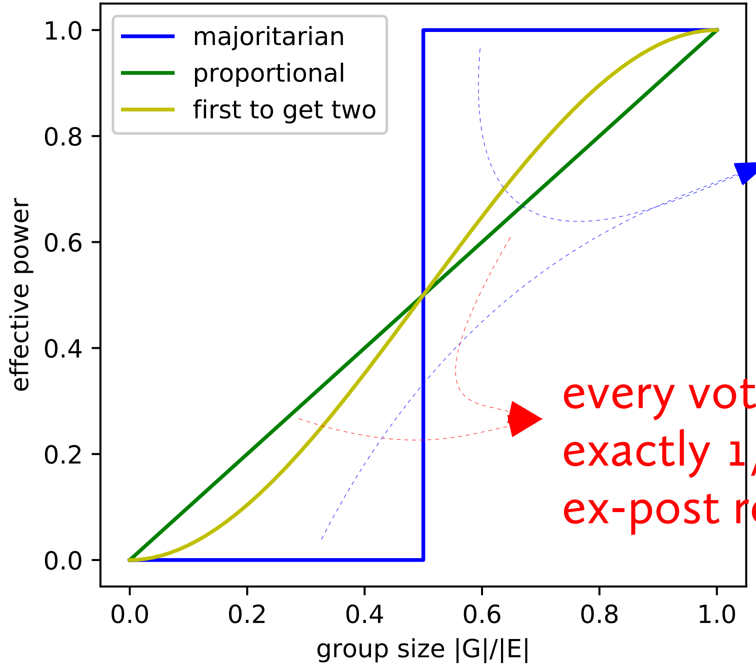

the majority has full responsibility as a group, no single voter has any ex-post responsibility unless the decision was ~fifty-fifty

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

"Supporting Consensus" in formal voting methods

(Here: no distinction between *consensus*, *consent*, and *accepted compromise*)

Def. (pragmatic): (Heitzig & Simmons 2010)

A ***potential consensus*** for a group is an option or lottery that all group members prefer over using Random Ballot inside the group.

example:

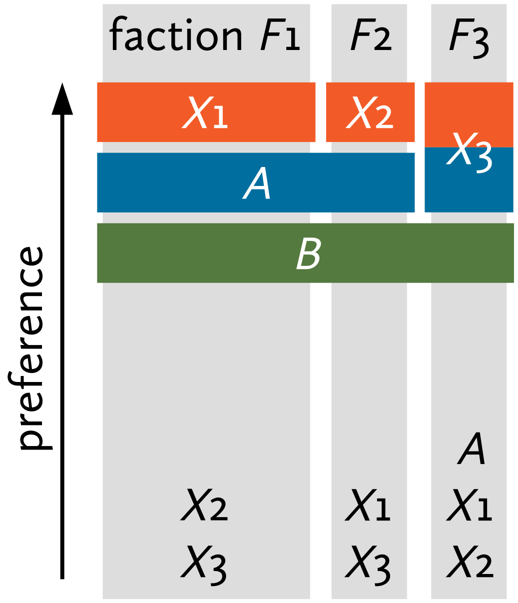

Option *A* is a potential ***partial*** consensus (for F1+F2)

Option *B* is a potential ***full*** consensus (for F1+F2+F3)

The lottery $75\%A + 25\%X_3$ is also a potential *full* consensus

“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.

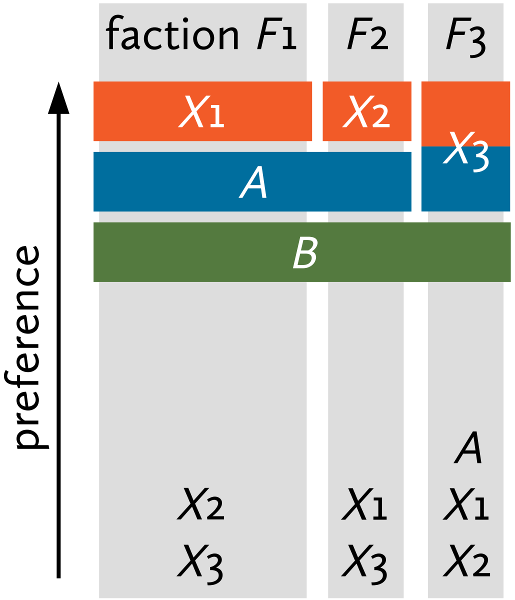

In the example:
Option *B* must be chosen in equilibrium

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

Simple solution: THE “TWO URNS” METHOD (Heitzig & Simmons 2010)
(but impractical)

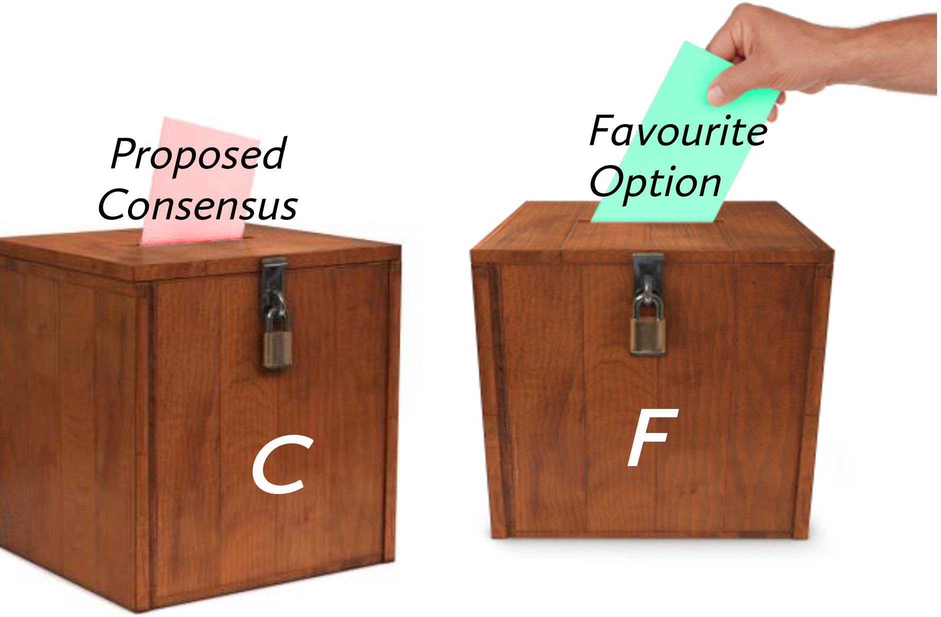

Voting: Each voter puts one standard ballot into urn *C* and one into urn *F*.

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~~, unavoidable (Gibbard/ Satterthwaite/Hylland)
simple, distributes power proportionally,
supports full ~~(& partial)~~ consensus,
~~produces high “welfare”~~,
~~reveals detailed preferences~~.

(but a version with 3 urns does)

“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 not option B exists,
option A must be chosen with at least 75% probability
in equilibrium.

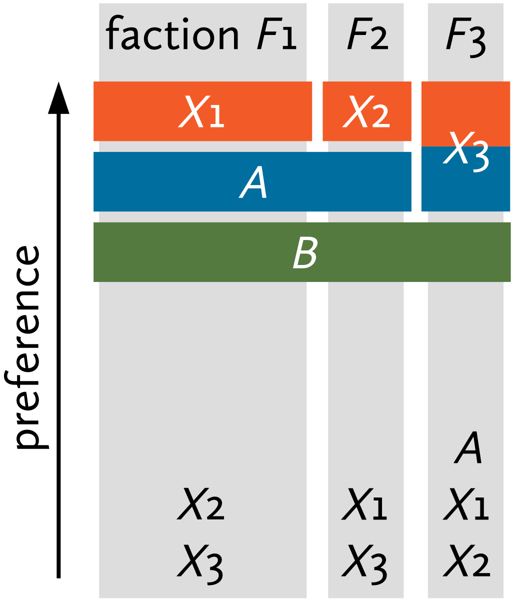

From Theory to Actual Method Design

Goal: ***Design a voting method for everyday group decisions that distributes power proportionally, supports full & partial consensus, and produces high "welfare"!***

Ingredients & Inspirations:

- Random Ballot (drawing ballots gives proportionality)
- Approval Voting (approval information helps finding potential consensus)
- Range Voting (numerical ratings help fine-tuning choices)
- Conditional Commitments (makes cooperation safe)
- Granovetter's threshold model of social mobilisation
- the Nash Bargaining Solution

New method 1: the "Nash Lottery"

Range Voting:

- Voting: each voter i gives each option x a **rating** r_{ix}
- Tallying: the option x that maximizes $\Sigma_i\, r_{ix}$ wins for sure

Nash Bargaining Solution of a bargaining problem:

- Choose the agreement a that maximizes $\Sigma_i \log(u_{ia} - u_{id})$,
  where u_{ia} [u_{id}] is the utility to i when a [or nothing] is agreed

→ "Nash Lottery" voting method:

- Voting: each voter i gives each option x a **rating** r_{ix}
- Tallying: find the lottery p that maximizes $\boldsymbol{\Sigma_i \log(\Sigma_x\, p_x r_{ix})}$,
  then draw an option x from that lottery (i.e. with probabilities p_x)

Interpretation: **automatic bargaining over lotteries**

New method 1: the "Nash Lottery" (2)

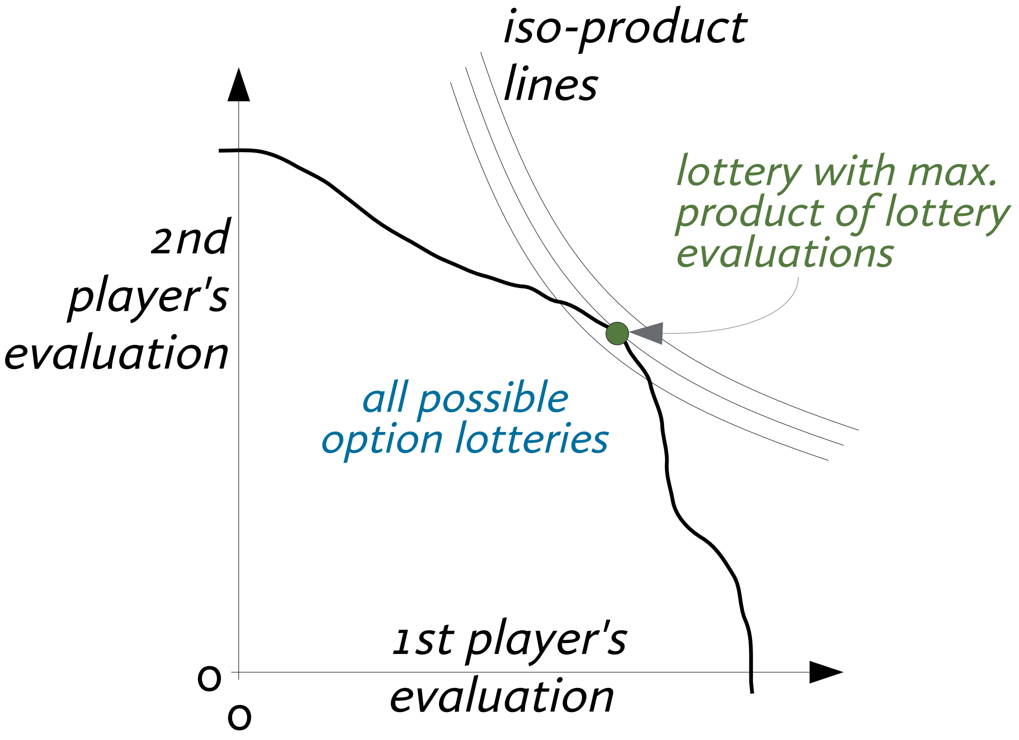

Properties:

anonymous

neutral

~~monotonic~~ (increasing a rating of one option may increase the chances of *another* option)

Pareto-efficient

~~strategy-free~~

~~simple~~ (requires numerical optimization; result is hard to interpret)

distributes power proportionally

supports full **& partial** consensus

~~reveals detailed preferences~~.

What about "welfare"?

New method 2: "Maximum Partial Consensus"

Idea: Each voter "owns" an equal share of the winning probability and the method provides a simple way by which voters can agree to jointly shift their shares from their various favourites to a potential consensus option.

New method 2: "Maximum Partial Consensus"

Idea: Each voter "owns" an equal share of the winning probability and the method provides a simple way by which voters can agree to jointly shift their shares from their various favourites to a potential consensus option.

Voting: Each voter gives each option x a rating $0 \le r_{ix} \le 100$

- Interpretation: i conditionally commits to approve of x iff less than r_{ix} percent of all voters do *not* approve of x)

New method 2: "Maximum Partial Consensus"

Idea: Each voter "owns" an equal share of the winning probability and the method provides a simple way by which voters can agree to jointly shift their shares from their various favourites to a potential consensus option.

Voting: Each voter gives each option x a rating $0 \leq r_{ix} \leq 100$

- Interpretation: i conditionally commits to approve of x iff less than r_{ix} percent of all voters do *not* approve of x)

Tallying:

- Determine who approves of what according to that interpretation
  (as in Granovetter's threshold model from sociology)

this is the only nontrivial part!

- Draw one ballot at random
- Among those options approved on that ballot, the one with the largest overall approval wins

New method 2: “MaxParC” (2)

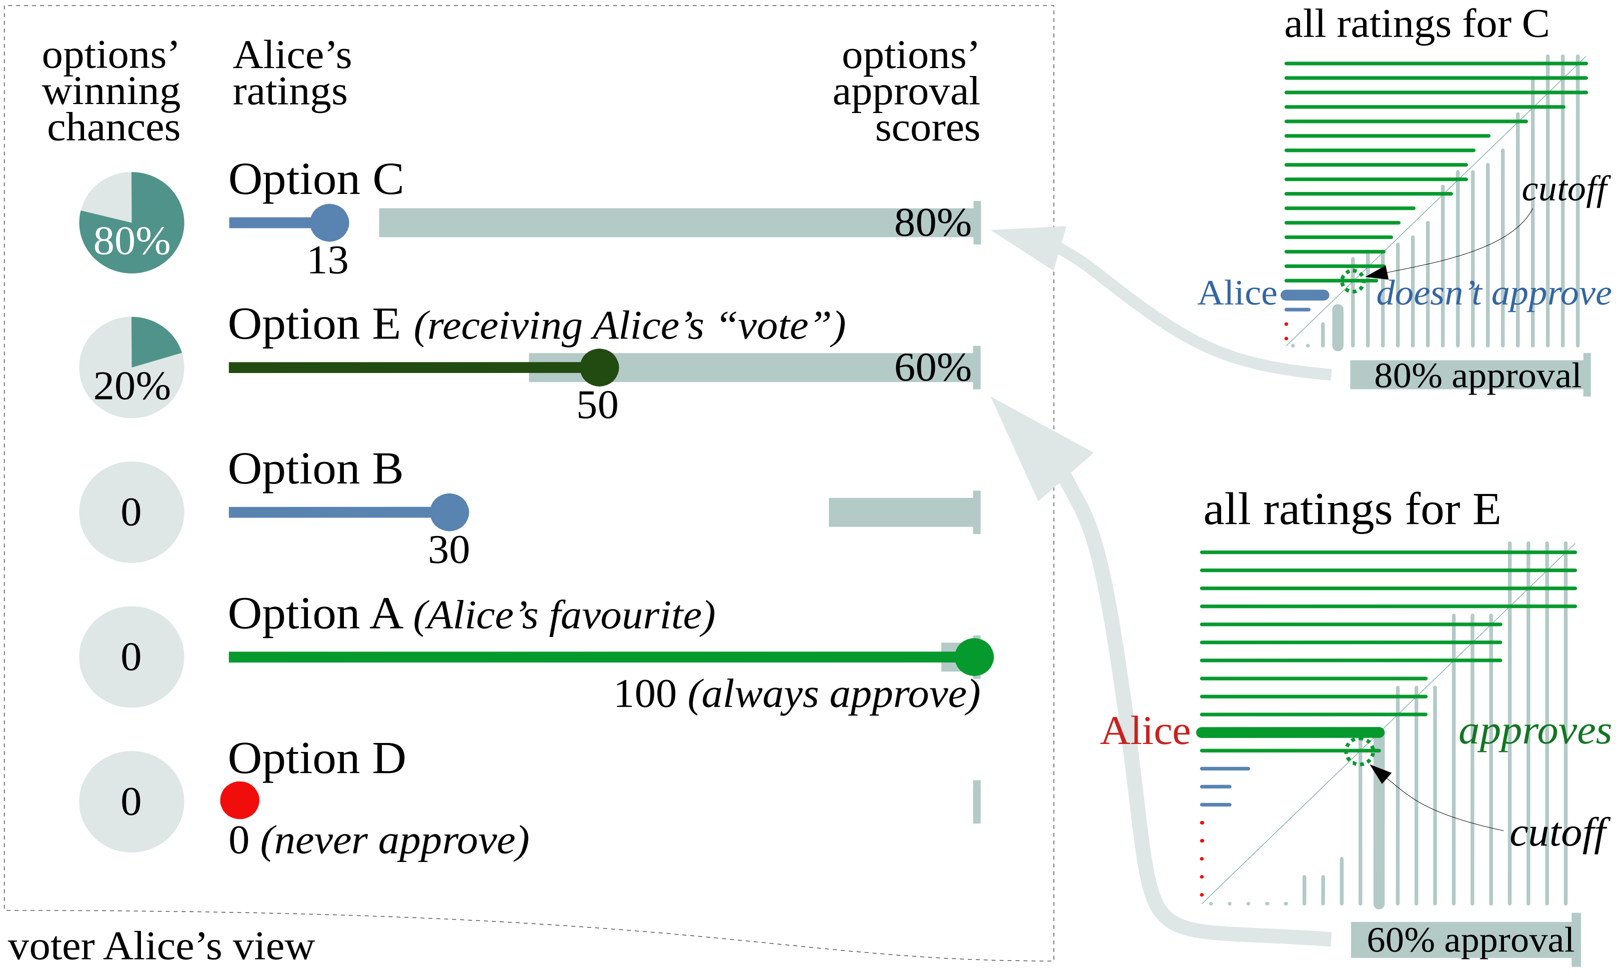

New method 2: "MaxParC" (3)

Properties:

anonymous,
neutral,
monotonic,
Pareto-efficient,
~~strategy-free~~,
simple(r),
distributes power proportionally,
supports full & partial consensus,
~~reveals detailed preferences~~.

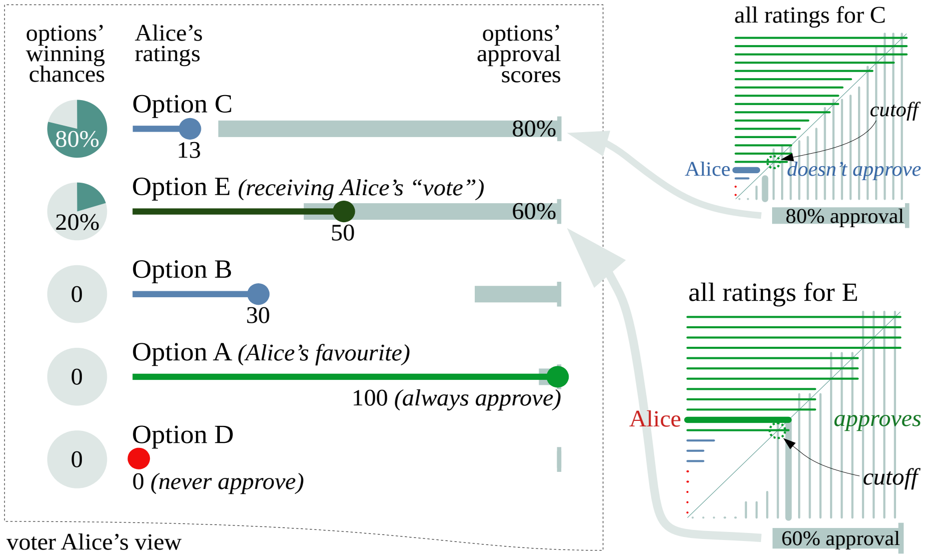

What about "welfare"?

Measuring "Welfare" in voting methods

Each voting method results in some lottery p of the options (maybe a "sure-thing" lottery that picks some x for sure).

Given all voters' *evaluations* $v_i(p)$ of this lottery p, one can compute a ***welfare function*** (→ welfare economics)

$W(p) =$

$\Sigma_i\, v_i(p)\ /\ N$	(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**

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

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

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

Voters have different **voting behaviours**

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

This type of heterogeneity calls for behavioural experiments (I cannot do that) or for ***agent-based modeling***

Agent-Based Modelling

- Represent decision makers by individual **agents** with heterogeneous **attributes**
- Simulate what they do from **time step** to time step by programming individual behavioural **rules**
  object-oriented: agent type → class, behavioural rule → class method

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 an major 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

Simulated Behavioural Patterns

crucially informed by theory!

If voting method = Approval Voting:

- **Lazy** voters: *approve of favourite and no other option*
- **Sincere voters:** *approve of what you prefer to the Random Ballot lottery according to 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 heuristic; during interactive phase, pick a random option, then change your approval of it if you profit from that change*
- **Factionally strategic:** *start heuristic; during interactive phase, switch to your faction's best response to the other factions' current votes*

heavy numerical optimization

Similar for other voting methods (details differ considerably)

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: uniform, block, and several spatial models; 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)
- 10 different voting methods, with or without interactive phase

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

Example in a two-dimensional policy space

Random Ballot:

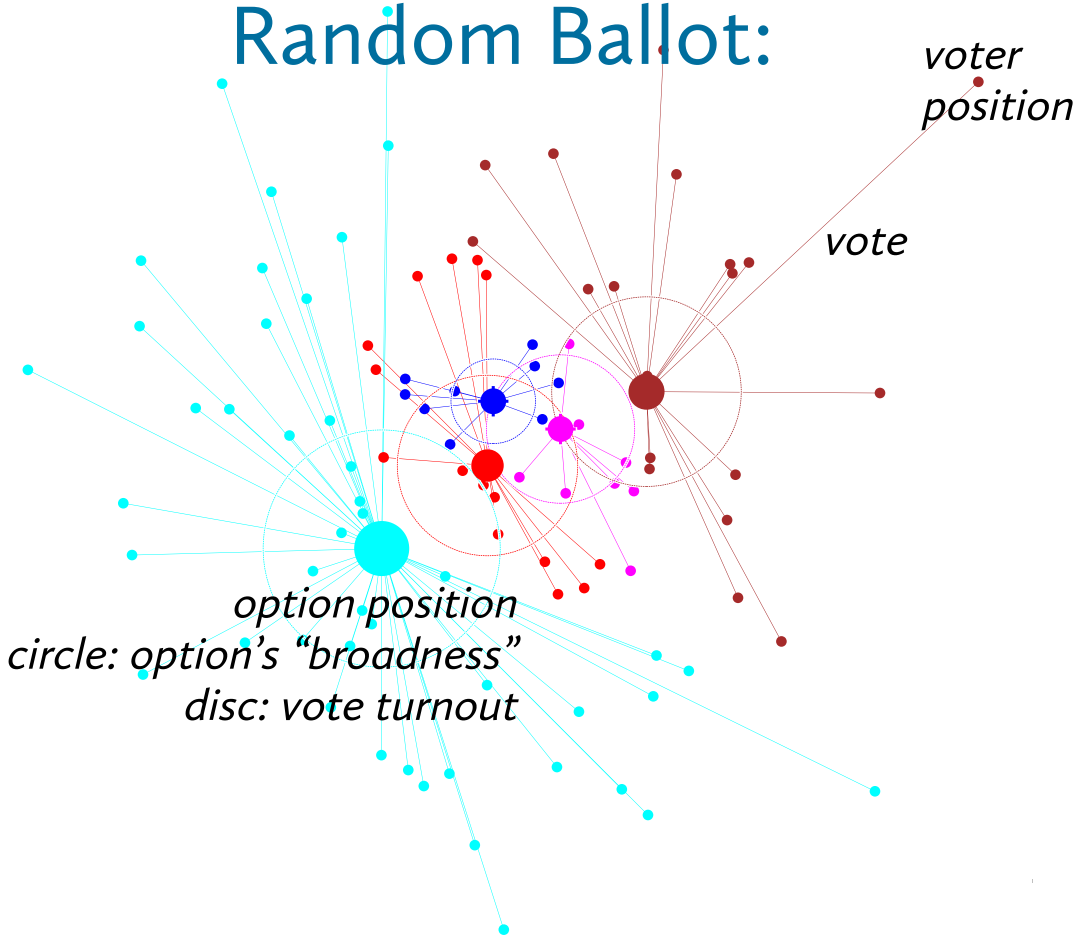

MaxParC:

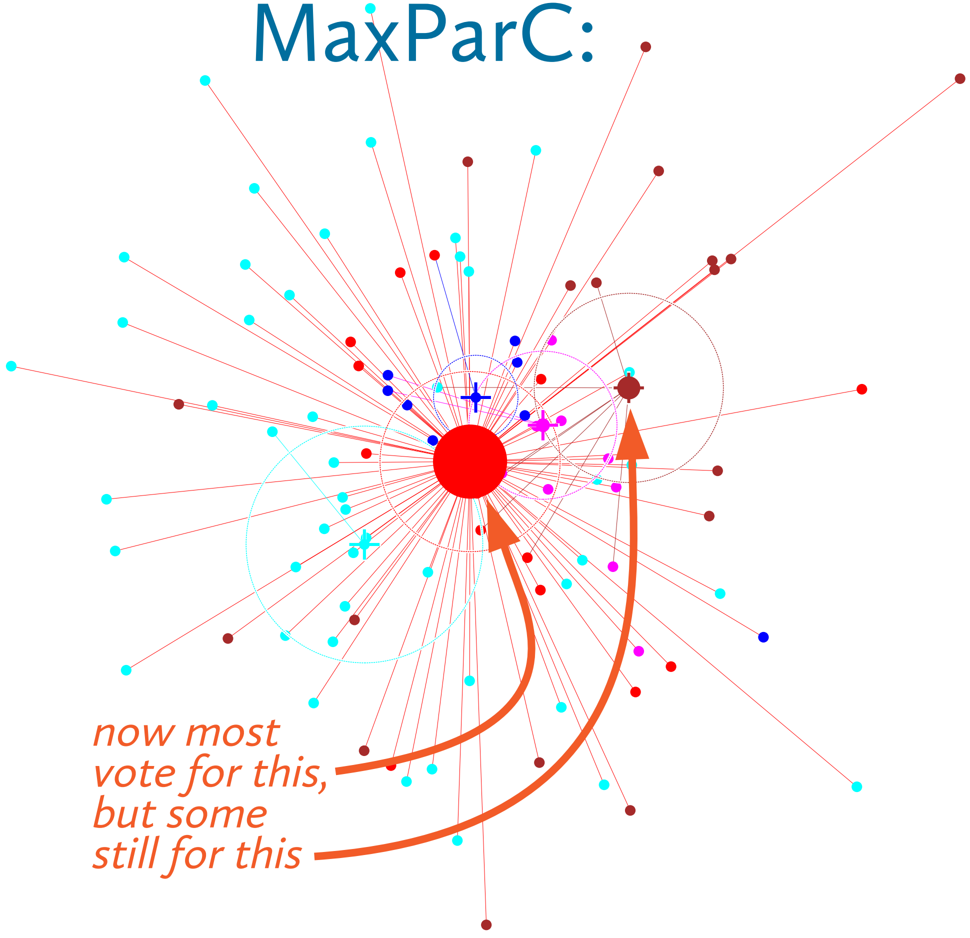

Selected Results

- 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 lead to considerable entropy
- Strategic voters have negligible advantage over lazy voters
- Among parameters, preference model has strongest effect on all this

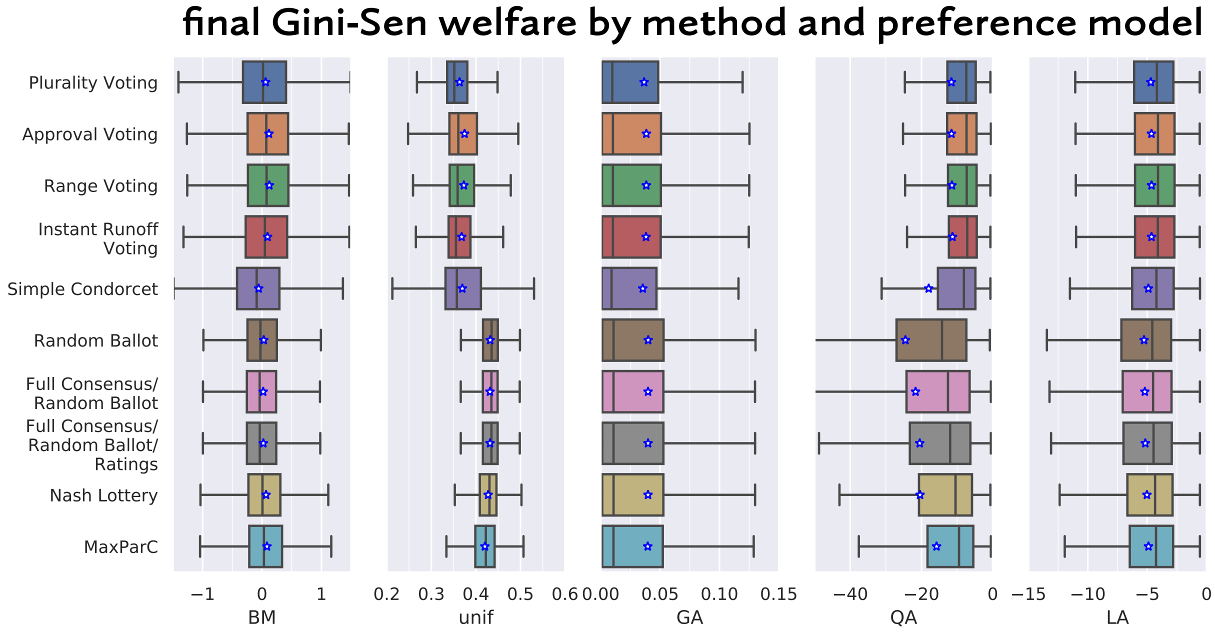

Summary

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

- *Both theoretical analysis and agent-based simulations are needed to assess the formal, qualitative, and quantitative properties of voting methods*

- Proper agent-based studies crucially depend on…
  - theory-guided specification of behavioural rules
  - careful treatment and analysis of uncertain parameters

Thank you for your attention!
→ Questions? Comments? Potentials for collaboration?

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

slides and paper: www.pik-potsdam.de/members/heitzig/maxparc

prototype of related voting app: www.vodle.it

for developers: github.com/mensch72/maxparc-ionic