Fair and efficient collective decisions via nondeterministic proportional consensus

Jobst Heitzig
FutureLab on Game Theory & Networks of Interacting Agents, Potsdam Institute for Climate Impact Research, Potsdam, Germany, heitzig@pik-potsdam.de
Forest W. Simmons
Liberal Arts & Mathematics Division, Portland Community College, Cascade Campus, Portland, Oregon, USA

Abstract

Are there collective decision methods which (i) give everyone, including minorities, an equal share of effective power even if voters act strategically, (ii) promote consensus and equality, rather than polarization and inequality, and (iii) do not favour the status quo or rely too much on chance? We show the answer is yes by describing two nondeterministic voting methods, one based on automatic bargaining, the other on conditional commitments to approve compromise options. Our theoretical analysis and agent-based simulations show that, when these group decision methods are used, majorities cannot consistently suppress minorities as with deterministic group decision methods, proponents of the status quo cannot block decisions as in other consensus-based approaches, the resulting aggregate welfare is comparable to that provided by other common voting methods, and the average amount of chance employed by the method is lower than for other nondeterministic methods. Our results suggest that the welfare costs of fairness and consensus are small compared to the inequality costs of majoritarianism.
Keywords: non-majoritarian voting, conditional commitments, strategic voting

1 Introduction

Picking the family's holiday destination, the movie to watch with friends, the new logo for a joint venture — groups and organizations make collective decisions all the time. Whether informally or by formal voting, such group decisions often follow the ideal of majority rule. This is not only true if the "standard" voting method of Plurality Voting is used. Also most other common voting methods, such as Approval Voting (Brams and Fishburn, 1978), essentially aim to identify "the will of the majority". But majority rule, even though it is considered a cornerstone of democracy, allows for the oppression of minorities (Lewis, 2013). On a large scale, this 'tyranny of the majority' may lead to separatism or violent conflict (Collier, 2004; Cederman et al., 2010) and ultimately lead to welfare losses and unnecessary suffering. To give just one example, Devotta (2005) documents the role of majority voting in the Sri Lankan separtist war. But also in smaller groups, it seems unlikely that, say, a clique of three boys and two girls will exist for very long if the three males always outvote or otherwise overrule the two females when choosing their leisure activities.

A natural idea for addressing the tyranny of the majority is to somehow employ the fundamental fairness principle of proportionality (Cohen, 1997; Cederman et al., 2010). Large groups often do so via proportional representation. But if proportionality is only used to elect a representative body which then still uses a group decision method based on majorities in their own decisions, the tyranny of the majority can be upheld anyway (Zakaria, 1997). Why? Proportional representation does not imply proportional power: even a 49 percent faction may not be able to influence any single decision. For example, given the strong polarization in the US Senate (McCarty et al., 2016), the Democratic Party currently would appear to have zero effective power in it (if it were not for such oddities as the "filibuster"). This can be seen, e.g., if one computes their Banzhaf and Shapley–Shubik power indices (Dubey and Shapley, 1979). But can effective power be distributed proportionally at all?

Smaller groups often try to overcome the majority problem by seeking various forms of deliberation aimed at consensus or consent, but that is difficult in strategic contexts (Davis, 1992) and may be perceived as less legitimate than formal voting (Persson et al., 2013). Supporters of the status quo may block consensus indefinitely. If the protocol uses a fallback method that is applied when no consensus is reached by some deadline, and if this fallback method is majoritarian, then any majority can simply wait for this to be invoked. Hence common consensus procedures are either not neutral about the options because of the special role of the status quo (Bouton et al., 2018), or are effectively majoritarian. The latter feature is shared by most common group decision methods, at least if voters may act strategically (Bouton, 2013; Kawai and Watanabe, 2013; Spenkuch et al., 2018).

Judging from social choice theory, such nonproportional effective power distribution seems unavoidable (May, 1952), no matter whether in a political or an everyday group decision making context. But this is only so if the employed decision methods are required to be essentially *deterministic,* only using chance to resolve ties and otherwise not using any randomization.

In fact, it is quite easy to distribute effective power in a completely proportional way with *nondeterministic* methods. In such methods, the winning option is at least sometimes determined using some amount of chance, not only to resolve ties but, e.g., to achieve fairness or provide incentives for cooperation. If potentially non-deterministic methods are considered, it is natural to measure the effective power of a group by the amount of winning probability that the group can guarantee their chosen option. The 'Random Ballot' (aka 'Lottery Voting') method (Amar, 1984), even though it also has many undesirable properties, quite easily distributes this kind of effective power in a perfectly proportional way. In that method, all voters mark a single option on their ballot, then one of the ballots is drawn at random and decides the winner. Variants of this method can also be used to distribute effective power in ways between the "majority-takes-all" approach of majoritarian methods and a perfectly proportional distribution. E.g., one could draw a sequence of standard ballots until one option's vote count is two; this method (which might be called "first to get two") would lead to the S-shaped power distribution shown in Fig. 1.

Not quite so trivial is the question of how to support consensus at the same time as distributing power proportionally, so that the group decision method will not only lead to fair but also efficient outcomes that avoid extreme results and foster social cohesion. 'Random Ballot' does the exact opposite: voters have no incentive whatsoever to mark a potential consensus option even if it was everybody's second-best choice and a very good compromise. Still, recent theoretical results show that the combination of fairness (proportional allocation of power) and efficiency (electing good compromise options as consensus) is possible when chance is used in some way that incentivizes consensus (Heitzig and Simmons, 2012; Börgers and Smith, 2014).

To many, it will seem outlandish to use a decision method that employs chance on a regular basis, thus producing uncertain outcomes, even though there is some highranking theoretical literature on such methods (e.g., Brandl et al. (2016)). But real-world problems typically involve quite some unavoidable stochastic risk and other forms of uncertainty anyway, e.g., due to lacking information, complexity, dependence on others (Carnap, 1947), or having been born into a certain voting district, gender, skin color, etc. Also, nondeterministic procedures are routinely used in contexts such as learning (Cross, 1973), optimization (Kingma and Ba, 2014), strategic interactions (Harsanyi, 1973), or the allocation of indivisible resources as in school choice (Troyan, 2012), and are also increasingly proposed for

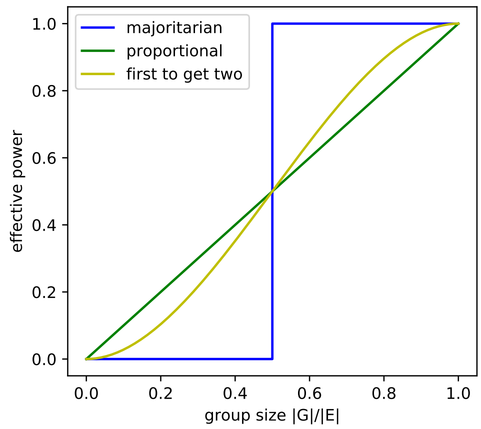

Figure 1: Distribution of effective power by type of group decision method. Effective power of a subgroup G of an electorate E, as a function of relative group size $|G|/|E|$, for different types of voting methods. Blue: majoritarian deterministic methods such as Plurality Voting, Approval Voting, Range Voting, Instant Runoff Voting, and Simpson–Kramer Condorcet. Green: proportional nondeterministic methods such as Random Ballot, Full Consensus/Random Ballot, Full Consensus/Random Ballot/Ratings, and our two novel methods Nash Lottery and MaxParC. Yellow: an example of a nonproportional nondeterministic method ("first to get two").

composing citizens' councils or even appointing officers as in ancient Athens. This shows that using chance can be quite beneficial, efficient, and acceptable. Those examples also demonstrate that carefully using chance must not be confused with outright randomness.

Problem statement In this article, we adopt the hypothesis that at least in everyday group decision situations on individual issues in which people would say "let's have a vote", many groups might try a nondeterministic voting method as part of their regular decision-making processes — potentially in combination with some method for deliberation — if that had clear advantages in terms of ethical criteria such as fairness or rather economic criteria such as efficiency. In view of the increasing popularity of social apps for polling and scheduling and online tools for collaboration, we also assume that, aided by a suitable tool, groups will be able to use considerably more sophisticated methods than plain Plurality Voting. For this type of situation we study from a normative perspective two such group decision methods — one that we translated from a different context and one that we novelly designed — and show by theoretical analysis that they would indeed achieve fairness by distributing power proportionally and increase efficiency by supporting not just full but also partial consensus and compromise.

As a paradigmatic test case (Fig. 2), consider a group of three factions, F_1, F_2, F_3, with sizes $S_{1,2,3}$ (in percent of voters). Assume each faction has a favourite option, $X_{1,2,3}$, that is however not liked by the other two factions, respectively. Now suppose there is a fourth option, A, which is not liked by faction F_3, but which is liked by factions $F_{1,2}$ almost as much as their respective favourites, $X_{1,2}$. Let us call A a potential 'partial consensus' for the factions $F_{1,2}$ together. While efficiency (picking a "good" option) requires that A gets a considerable chance of winning, proportionality requires that also X_3 gets some chance of winning — namely exactly S_3 %. Accordingly, the two voting methods described in this

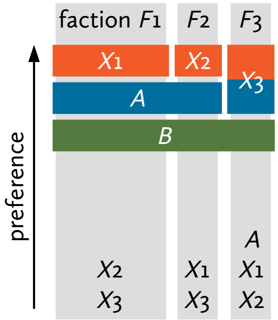

Figure 2: Archetypical group decision problem with potential for suppression of minorities, partial, or full consensus. Each of three factions of different size (column width) has a unique favourite (topmost). There might also be a potential 'partial consensus' option A and/or a 'full consensus' B. With strategic voters, common deterministic methods pick X_1 for sure. Our methods *Nash Lottery* and *MaxParC* pick B for sure if present (green); they pick one of $X_{1,2,3}$ with probabilities (colored area) proportional to faction size if neither A nor B is present (orange); and they pick A or X_3 with proportional probabilities if A but not B is present (blue).

paper will assign probabilities of $S_1 + S_2\,\%$ to A and $S_3\,\%$ to X_3. They will do so not only if all voters vote "sincerely" (i.e., honestly according to their true preferences) but even if some or all voters vote strategically (e.g., by misrepresenting their preferences in some way).

Even more so, if we add a fifth option, B, which factions $F_{1,2}$ like slightly less than A, and which also faction F_3 likes almost as much as their favourite X_3, then both our group decision methods will pick this potential 'full consensus' option B for sure rather than employing chance. In contrast, if the faction sizes fulfil $S_1 > S_2 + S_3$ and if voters act strategically, then virtually all existing group decision methods will either pick X_1 with certainty, or will assign probabilities of $S_{1,2,3}\,\%$ to $X_{1,2,3}$, respectively. In both cases, these methods would ignore the potential compromises A and B and thus produce much less overall welfare. In fact, no deterministic group decision method can let the two factions $F_{1,2}$ together make sure that A gets a chance without also allowing them to render faction F_3's votes completely irrelevant! This can only be achieved by employing a judicious amount of chance.

But how exactly should chance be used? That is, how to design a nondeterministic group decision method that is both efficient and proportional, even when voters act strategically, and which also fulfils other basic consistency requirements like those typically studied in social choice theory — such as anonymity, neutrality, monotonicity, and clone-proofness — that make it plausible and hard to manipulate? We will see that suitable combinations of ingredients from existing group decision methods such as Approval Voting and Random Ballot, game theoretical concepts such as the Nash Bargaining Solution, and a certain sociological model of social mobilisation can do the job.

Results

The Nash Lottery Our first method, the *Nash Lottery (NL)*, is basically what is known as 'Nash Max Product' or 'Maximum Nash Welfare' in the literature on fair division of resources. As suggested in Aziz et al. (2019), we translate it to our context by interpreting *winning probability* as a "resource" to be divided fairly, and study the strategic implications of this. The Nash Lottery can be interpreted as a form of *automatic bargaining* by means of the well-known Nash Bargaining Solution. Similar to score-based methods such as Range Voting (RV) (Laslier and Sanver, 2010), it asks each voter, i, to give a *rating,* $0 \leqslant r_{ix} \leqslant 100$, for each option x. Like Range Voting, the Nash Lottery then assigns winning probabilities, p_x, to all options x so that a certain objective quantity, $f(r,p)$, is maximized.

Range Voting maximizes the quantity $f(r,p) = \sum_i \sum_x r_{ix} p_x$, which is motivated by its formal similarity to a utilitarian welfare function. This results in a very efficient majoritarian method that is deterministic (i.e., we usually have $p_x = 1$ for some x except for ties). This determinism is because $f(r,p)$ is a linear function of p. But that method neither distributes power proportionally nor supports consensus when voters are strategic. In the example of Fig. 2, faction F_1 will quickly notice they can make X_1 win for sure by putting $r_{ix} = 0$ for all other options. Indeed, Range Voting is more or less strategically equivalent to the simpler Approval Voting (Dellis, 2010). As a consequence, strategic voters almost never have an incentive to make use of any other rating than either 0 or 100. In case there is no Condorcet winner (an option preferred to each other option by some majority), which is not too unlikely (Jones et al., 1995), there is not even any strategic equilibrium between factions, and thus the outcome is largely unpredictable.

The Nash Lottery instead maximizes the quantity

$$f(r,p) = \sum_i \log \left(\sum_x r_{ix} p_x \right), \tag{1}$$

which gives a nondeterministic method (i.e., usually several options have a positive probability p_x) that supports both full and partial consensus. In *Appendix 2.1.5*, we prove that in situations similar to Fig. 2, a full consensus will be the sure winner under the Nash Lottery, and a partial consensus would get a proportional share of the probability. This would be so both in the case where voters are sincere and in a certain strategic equilibrium. Even more so, we show in *Appendix 2.1.4* that using the logarithm rather than any other function of $\sum_x r_{ix} p_x$ is the only possible way to achieve a proportional distribution of power.

The Nash Lottery is conceptually simple and has some other desirable properties shown in Fig. 4 such as being immune to certain manipulations, e.g., cloning an option or adding a bad option to effect certain changes in the tallying process. But this group decision method also has three important drawbacks.

Its tallying procedure is intransparent because it requires numerical optimization that cannot be performed with pen and paper.

It also lacks a certain intuitive 'monotonicity' property: when a new option is added or a voter increases some existing option's rating, this might have the effect that some other option's winning probability increases rather than decreases. E.g., consider three options (A, B, C) and two voters one of which rates them $(6, 0, 2)$ and the other $(0, 12, 9)$. Then it turns out that the Nash Lottery assigns zero winning probability to B and gives all probability to A and C. But if the first voter increases her rating of A from 6 to 12, B's winning probability increases from 0 to 50 %.

Last but not least, the Nash Lottery often employs much more randomness than necessary.

Maximal Partial Consensus (MaxParC) All three drawbacks are overcome by our second method, the novel *Maximal Partial Consensus (MaxParC).* It is conceptually more complex, but is strongly monotonic and much easier to tally and uses less chance. Like with Random Ballot, each voter's "vote" represents an equal share, $1/N$, of the total amount of winning probability. But unlike Random Ballot, MaxParC lets each voter safely transfer

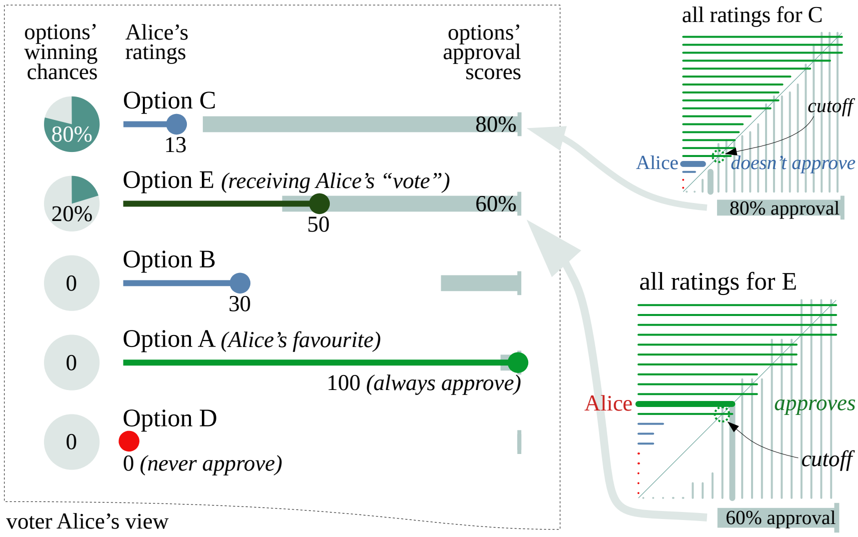

Figure 3: Group decision method MaxParC from the view of some voter Alice. Left box: Each rating value (thick colored needles coming in from the left) represents a conditional commitment by Alice to approve the respective option. Approval scores (nos. of voters approving an option) are represented by light bars coming in from the right, options are sorted by descending approval score. The rating of 13 for option C, for instance, is interpreted as saying that Alice approves C if and only if less than 13 percent of voters do *not* approve C. In other words, Alice is counted as approving an option if her rating needle overlaps with the approval score bar (green needles), and is otherwise counted as not approving the option (blue needles). Right diagrams: Approval scores can be determined graphically in a way similar to Granovetter (1978) by finding the leftmost intersection of the graph of ordered ratings with the main diagonal. Alice's "vote" (= share of $1/N$ of the winning probability) goes to the most-approved option approved by her (dark green needle). Winning probabilities are show to the left.

their vote from their favourite option to any potential consensus option. In this, "safely" means that the transfer only becomes effective if enough other voters transfer their votes as well so that this collaborative shift of support becomes desirable to all involved voters. This effect is achieved by a design element inspired by Granovetter's famous threshold model of social mobilisation (Granovetter, 1978; Wiedermann et al., 2020) which we use to implement a form of *conditional commitments*.

How is this done exactly? In MaxParC, voters again assign numerical ratings, $0 \leqslant r_{ix} \leqslant 100$. These are now interpreted as a 'willingness to approve', stating that "voter i will approve option x if strictly less than r_{ix} percent of all voters do *not* approve x." Or, equivalently, "voter i will approve option x if strictly more than $100 - r_{ix}$ percent of all voters *do* approve x." In particular, a zero rating means "don't approve no matter what" and a rating of 100 means "do approve for sure". All ratings together result in a set of mutually dependent constraints for the question of which voter ends up approving which options. It turns out that this set of equations can be solved quite easily in the same way as in Granovetter (1978). For each option x, one simply sorts the ballots ascendingly w.r.t. their rating of x. In this ordering, one seeks the first ballot, i, that is preceded by strictly less than r_{ix} percent of all ballots. This ballot i and all later ballots j (those with $r_{jx} \geqslant r_{ix}$)

are said to "approve" x. In other words, the rating of this "pivot" voter i serves as a "cutoff" value for everyone's approval x. Graphically, the cutoff can easily be identified by finding the first intersection of the sorted ratings graph with its main diagonal, as in Fig. 3 (right). After thus determining which voters approve which options, MaxParC then proceeds like the 'Conditional Utilitarian Rule' from Duddy (2015); Aziz et al. (2019): one ballot is drawn at random, and from the options approved by this ballot, the one with the largest overall no. of approving ballots wins. In case of remaining ties, the aggregated rating values decide.

Fig. 3 illustrates the MaxParC procedure, which can (at least in principle) be performed using pen and paper. It is easy to see that if at least one option is rated positive by everyone, then, because of the used tie-breaking rule, among all such options the one with the largest aggregate rating will win for sure. In that case, MaxParC is like Range Voting restricted to the set of universally approved options. Only if no option gets all positive ratings, chance will really play a role. For voters who are risk averse, the potential use of chance can work as a kind of threat against being too uncompromising, as in Heitzig and Simmons (2012). This gives voters a clear incentive to search for good compromise options and rate them positively in order to reduce the uncertainty. Indeed, MaxParC supports partial and full consensus just as well as the Nash Lottery does: in the test case of Fig. 2, all voters will give B a positive rating and thus make it win for sure; if B is not available, the $F_{1,2}$ voters will give A a rating slightly above S_3 and thus transfer their votes and the corresponding winning probability safely from $X_{1,2}$ to A. In both cases, no voter has an incentive to reduce their rating of A or B in order to keep their vote for their favourite option, because that would cause all others' votes to go back to their favourites as well. The *Appendix* contains a formal proof of this claim in the form of several game-theoretic equilibrium results, as well as a formal analysis of the other properties discussed here.

Our theoretical analysis of the formal properties of the Nash Lottery and MaxParC as compared to typical group decision methods from the literature are summarized in Fig. 4, which verifies that these two methods perform well in terms of the considered qualitative criteria. It remains to study how they perform in more quantitative respects.

Performance in agent-based simulations To assess the potential costs of achieving fairness and supporting consensus in more quantitative terms of welfare, voter satisfaction, and amount of randomization, we complement our theoretical analyses by some first empirical evidence. Because a suitable lab experiment appeared prohibitively costly and will pose unique methodological challenges, we start the empirical study of the methods here by reporting on a large agent-based simulation we performed, which can be interpreted as a complex, computer-assisted thought experiment. In a very diverse sample of over 2.5 million hypothetical group decision problems, we compared the Nash Lottery's and MaxParC's performance to that of eight other methods. We chose five popular deterministic majoritarian group decision methods: Plurality Voting (PV, aka 'first past the post'), Approval Voting (AV), Range Voting (RV), Instant Runoff Voting (IRV, aka 'ranked choice voting'), and the 'Simpson–Kramer' method (aka 'Simple Condorcet', SC). As the three nondeterministic proportional methods, we used 'Random Ballot' (RB, aka 'Lottery Voting') and two methods from Heitzig and Simmons (2012): 'Full Consensus/Random Ballot' (FC) and 'Full Consensus/Random Ballot/Ratings' (RFC). To generate the decision problems, we used random combinations of the number and compromise potential of options, and the number, individual preference distributions, and risk attitudes of voters. We used various preference models from behavioural economics and the spatial theory of voting. For each combination of decision problem and group decision method, we simulated several opinion polls, a main voting round, and an interactive phase where ballots could be modified continuously for strategic reasons. In this, we assumed various mixtures of behavioural types of voters: lazy voting, sincere voting, individual heuristics, trial and error, and coordinated

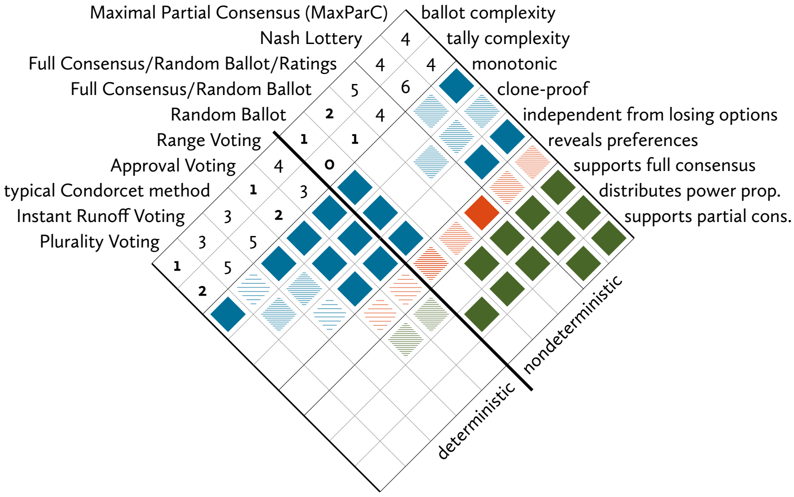

Figure 4: Properties of common group decision methods, the Nash Lottery, and MaxParC. Solid, densely dashed, and thinly dashed squares indicate full, strong partial, and weak partial fulfillment, respectively. Numbers are qualitative complexity assessments by the authors. Color is used to group criteria. See the *Appendix* for details and proofs.

strategic voting. For each decision problem, we computed several metrics of social welfare, randomness, and voter satisfaction for all group decision methods, and which voters would prefer which group decision methods (see the *Appendix* for details).

As can be expected from the very definition of 'majoritarian method', typically some majority of the simulated voters (namely that which was getting their will under the majoritarian methods) preferred the results of the majoritarian methods over those of the proportional ones. Within the group of proportional methods, voters preferred MaxParC over the other methods on average. Among the majoritarian methods, there was no such predominant method preference. Individual voters' satisfaction — normalized to zero for their least-preferred option and to unity for their favourite — averaged around 67 % for PV, AV, RV, and IRV; 61 % for SC, NL, and MaxParC; and still 57 % for RB, FC, and RFC. So MaxParC still produced 91 % of the voter satisfaction of the best deterministic methods, and about the same as typical Condorcet methods.

Regarding randomness, MaxParC produced only about 60 % of the entropy that RB produced, while NL still produced about 80 % of RB's entropy. In MaxParC, the largest winning probability was about 65 % on average, in NL only about 53 %. So MaxParC used significantly less randomization than NL.

The deterministic methods produced somewhat higher welfare on average, but for some preference models and welfare metrics, the nondeterministic methods matched or outperformed them. In two of the preference models ('BM' and 'unif'), the majoritarian methods generated slightly larger absolute utilitarian welfare and slightly smaller absolute egalitarian welfare values than the proportional methods. On the intermediate Gini–Sen welfare metric, the proportional methods beat the majoritarian ones in the 'unif' preference model, but were beaten in two other preference models ('QA' and 'LA'), see Fig. 5.

When fitting certain regression models for the considered welfare metrics, we saw that, generally and independently of the group decision method applied, a larger policy space dimension and larger voter heterogeneity were decreasing welfare. As expected, having more options increased welfare on average.

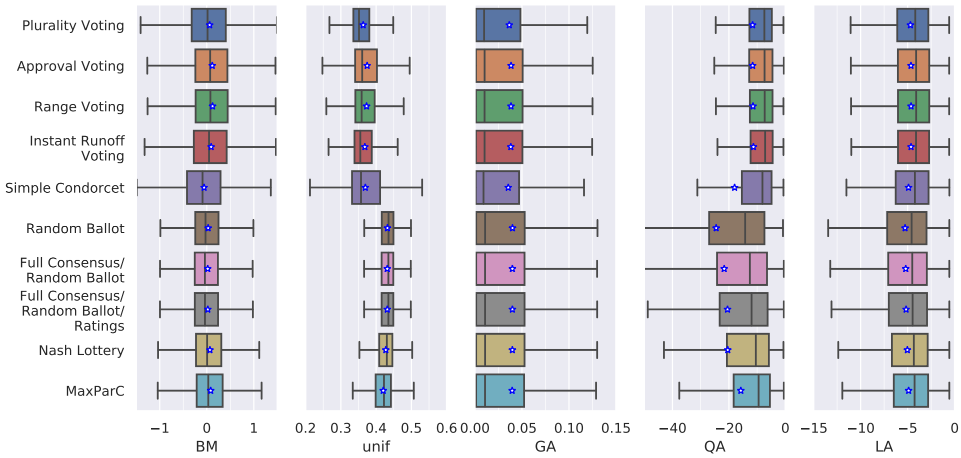

Figure 5: Simulated welfare effects of the considered group decision methods. Distribution of final Gini–Sen welfare across 2.5 mio. agent-based simulations by group decision method (rows), for five different models of how voter preferences might be distributed (columns). See *Appendix* for definitions and detailed additional results.

Finally, we compared for each decision problem (i) the difference between individual voters' utilities when a particular group decision method was used, and (ii) the difference between voters' average utility under different group decision methods. We found that in more than 75 % of all decision problems, the utility difference between the average and the worst-off voter when using Range Voting (the best deterministic method considered) was at least seven times as large as the difference in average voter utility between the results of Range Voting and MaxParC. This can be interpreted as saying that the welfare costs of fairness and consensus are small compared to the inequality costs of majoritarianism.

On most of these simulation results, preference distributions had a larger effect than behavioural type or the amount of interaction. Surprisingly, strategic voters gained no clear advantage over voters that used a simple heuristic, and also risk attitudes played a minor role. The good performance of heuristic voters also suggests that in MaxParC, a voter does not need to know others complete preferences in order to vote effectively, but can easily do so on the basis of simple approval polls. The *Appendix* has more detailed results.

Discussion

In 2007, one of the authors asked the Election Methods Electronic Mailing List (Lanphier et al., 1996) what method would elect the compromise (C) rather than the majority option (X_1) in a situation similar to Fig. 2, even when voters acted strategically. Soon it became obvious that no deterministic method would do, but several 'lottery' methods were quickly found that elected the compromise with certainty. So why do election methods experts show little enthusiasm for nondeterministic methods? Perhaps because their primary interest is in high-stakes public elections every several years. The proportional fairness of nondeterministic methods however arises from their average proportionality over many individual decisions. For instance, few would recommend that two newlyweds should resolve their future diagreements by electing a permanent household dictator by flipping a coin. Using coin flips for their many individual everyday decisions would be somewhat better — because stakes are lower and advantages level out over time. But that would still not lead to a single consensus. Much better still would be using one of the two methods presented in this article, which would likely make them agree on some compromise in most situations and toss a coin only rarely. Both the splitting-up into many decisions and the incentives for agreement lower the resulting overall randomness. Indeed, in view of our results, not much speaks against the recommendation to routinely use MaxParC in all everyday group decision problems where currently Plurality Voting is used. After all, voters can easily reduce or completely avoid unwanted randomization in MaxParC by rating compromise options positively. Thus, the more risk-averse the voters, the more likely consensus will become with MaxParC. Of course, one consequence of radical proportionalism is that individual voters who won't compromise can still have their potentially extreme will with a small probability. So the options on the menu of a group decision must be guaranteed to not violate any basic rights. But history has demonstrated that this is always the case, even with majoritarian methods.

Ideally, one would perform large-scale behavioural experiments to test our findings in various real-world group decision contexts, and indeed this should be the natural next step in the research agenda on improved group decision methods. However, such experiments will likely pose unique methodological challenges to the experimenter since in everyday group decisions, preferences can be complex and much less easily related to the monetary payouts typically used in behavioral experiments, and voters will likely be influenced as well by procedural preferences, fairness perceptions, and framing issues. On another angle, also the sets of strategic equilibria emerging from the discussed methods under rational choice conditions should be explored in an in-depth game-theoretic study.

Finally, while it may seem that consensus-supporting proportional methods of the type discussed here are best used for everyday decisions only, we'd like to point out by way of an outlook that they might also be applied to larger decisions such as allocating some budget or electing a parliament. In that case they will be deterministic rather than probabilistic. This is because the asset distributed by a proportional method need not be the winning probability in a single-outcome decision as in this article, but could be some other resource to be shared. For example, the asset might be parliamentary seats: suppose that instead of one of the common simple proportional methods, the Nash Lottery or MaxParC was used for allocating parliamentary seats to party lists, based on voters' ratings of all parties, and converting winning probability shares into seat shares using some suitable rounding method. Would not this method take better advantage of opportunities for consensus without sacrificing proportionality? Since the seat distribution would on average have a lower entropy than usual, would it not avoid unnecessary fragmentation of parliament without sacrificing representation of minorities?

We hope that this discussion serves to stimulate the reader's imagination to some of the possibilities of application, as well as avenues for further exploration. Teams, organizations, and communities might thus learn to avoid knife-edge decisions that satisfy only a mere half of their members, as in so many recent political elections and referendums, first in small-scale everyday decisions, but eventually also in large-scale contexts.

Methods

Overall study design Based on theoretical considerations, we designed two candidate group decision methods, analysed their formal properties within a suitable theoretical framework, and assessed their quantitative performance aspects in a large agent-based simulation experiment.

Voting method design For designing the MaxParC method, we combined ingredients from Granovetter's threshold model of social mobilisation (Granovetter, 1978), Approval Voting (Brams and Fishburn, 1978), the nondeterministic group decision methods from ref. (Heitzig and Simmons, 2012), and the 'Conditional Utilitarian Rule' from refs. (Duddy, 2015; Aziz et al., 2019).

Theoretical framework For the theoretical analysis of the properties presented in Fig. 4, we used a formal mathematical framework similar to that used in other works from social choice theory, described in detail in *Appendix 1.2 and 2.1.* To determine which methods support partial or full consensus, we performed a game-theoretical analysis of archetypical decision situations similar to Fig. 2, using pure-strategy Nash equilibrium between individual voters or between factions of voters with similar preferences as the solution concept, see *Appendix 2.1.5.*

Analysed properties of group decision methods If a group decision method employs no chance at all, or only to resolve rare ties, we call it *deterministic,* otherwise *nondeterministic. Anonymity and neutrality* require a method to treat all voters and options alike, respectively. For *Monotonicity,* we use two variants of 'mono-raise' monotonicity (Woodall, 1997), and a property related to 'mono-add-plump' monotonicity (Woodall, 1997).

Independence from Pareto-dominated alternatives, losing options, or cloned options. The essence of these criteria is that the removal of (i) a Pareto-dominated option, (ii) any option receiving zero winning probability, or (iii) an option that is indistinguishable from another option from all ballots should have no effect on the winning probabilities (of options other than the clone).

Revelation of preferences. This non-crisp criterion is about how much of the true preferences of strategically acting voters can be derived from their voting behaviour.

Proportional allocation of effective power. This criterion requires that for every option x and group of voters G from the whole electorate E, there must be a way of voting β_G for G so that for all ways of voting β_{-G} of the other voters, the winning probability of x is at least as large as G's relative size: $p_x \geqslant |G|/|E|$. A related criterion was discussed for the special case of 'dichotomous preferences' under the name 'Core Fair Share' in ref.(Aziz et al., 2019).

Supporting of full or partial consensus. These non-crisp criteria demand that in archetypical decision situations similar to Fig. 2, the 'natural' strategic equilibria of the resulting voting game should lead to a full consensus being elected with certainty, or a partial consensus being elected with a probability at least proportional to its supporting voter group's size.

See *Appendix 2.1* for more formal definitions and analysis of these properties leading to the assessment in Fig. 4.

Agent-based simulation We simulate voter preferences by using one of several different models to generate a profile of individual utility functions over options, u_i, and then derive individual utility functions over lotteries of options depending on each voter's risk attitude type. We use eight different *utility models,* the five reported on in Fig. 5 and three special cases of them. In the uniform model (unif), all utilities are uniformly i.i.d. similar to the 'impartial culture' model (Laslier, 2010). In the block model (BM), we assign voters to blocks and model a voter's utility as the sum of a block-specific and an individual term, both normally distributed. The spatial models LA, QA, GA are variants of those used in the spatial theory of voting (Carroll et al., 2013) in political science, with utility depending on distance in policy space either linearly (LA), quadratically (QA) or in a Gaussian fashion

(GA), see *Appendix 1.3.1.3*. Roughly following ref.(Bruhin et al., 2010), one fifth of all simulated voters had expected utility theory attitudes towards risky prospects, while the rest had risk attitudes rather conforming to one of two variants of cumulative prospect theory, see *Appendix 1.3.1.4*. Agents' behaviour in polling and voting rounds and the interactive phase was assumed to follow either a 'sincere', 'lazy', 'heuristic', 'trial and error', or 'factional' behavioural rule. The 'trial and error' rule is based on stochastic local optimization, the 'factional' one on best response dynamics, see *Appendix 1.3.2 and 2.3*.

Experimental design We generated $1,293,906$ many independent group decision problems, drawing their various parameters (no. of voters, options, polling rounds, preference model, risk attitude scenario, and distribution of behavioural types) independently from certain probability distributions, see *Appendix 1.3.3*. For each decision problem, we generated a second one with one option replaced by a special compromise option, giving >2.5 mio. problems in total. For each of these, we first simulated several rounds of polling and then, for each of the ten group decision methods independently, an initial voting round and an interactive voting phase, all based on the same polling results.

Statistical analysis We use a set of welfare metrics based on three social welfare measures (utilitarian, Gini–Sen, and egalitarian welfare), taken either on an absolute or a relative utility scale, and either taken before or after the interactive phase of the simulations. These measures are aggregating voters' individual utility, $u_i(\ell)$, from the resulting lottery, ℓ, as modelled by the preference and risk models described above. In applications where there is only a single decision taken, these metrics can be interpreted as measuring the 'ex ante' efficiency of the group decision method, as opposed to the 'ex post' efficiency that would be based on the utilities of the actual options chosen by the resulting lottery. In applications where we imagine a sequence of decisions, the metric can be interpreted as measuring the long-run efficiency of the group decision method over the whole sequence of decisions, see *Appendix 1.3.4*. In addition, we also compare welfare differences between group decision methods with utility differences within the electorate to assess the influence of method choice on welfare. In analogy to the notion of a "price of anarchy" (Koutsoupias and Papadimitriou, 1999), we define a "relative cost of fairness" for this, see *Appendix 1.3.4.5*.

We assess the degree of randomization a group decision method actually applies by means of the resulting Shannon entropy, Rényi entropy of degree two, and the maximal probability any option gets. We also computed each voter's "satisfaction level"

$$\frac{u_i(\ell) - \min_{x\in C} u_i(x)}{\max_{x\in C} u_i(x) - \min_{x\in C} u_i(x)} \in [0,1], \tag{2}$$

where C is the set of options and ℓ the winning lottery produced by the group decision method. This would be zero if voter i's least preferred option won for sure, and unity if i's favourite won for sure. Based on these, we analyse average satisfaction levels in the whole electorate and, to assess possible advantages of strategic behaviour, by behavioural type.

Finally, to get an idea of which methods voters would chose if that choice was itself performed by majority rule, we counted for each decision problem how many voters would prefer the lottery resulting from some method A to that resulting from some method B. For detailed results, see *Appendix 2.2*.

References

Akhil Reed Amar. 1984. Notes Choosing Representatives by Lottery Voting. *The Yale Law Journal* 93, 7 (1984), 1283–1308.

Haris Aziz, Anna Bogomolnaia, and Hervé Moulin. 2019. Fair mixing: The case of dichotomous preferences. *ACM EC 2019 – Proceedings of the 2019 ACM Conference on Economics and Computation* (2019), 753–781. https://doi.org/10.1145/3328526.3329552 arXiv:1712.02542

Tilman Börgers and Doug Smith. 2014. Robust Mechanism Design and Dominant Strategy Voting Rules. *Theoretical Economics* 9 (2014), 339–360. https://doi.org/10.3982/te1111

Laurent Bouton. 2013. A theory of strategic voting in runoff elections. *American Economic Review* 103, 4 (2013), 1248–1288. https://doi.org/10.1257/aer.103.4.1248

Laurent Bouton, Aniol Llorente-Saguer, and Frédéric Malherbe. 2018. Get rid of unanimity rule: The superiority of majority rules with veto power. *Journal of Political Economy* 126, 1 (2018), 107–149. https://doi.org/10.1086/695477

Steven J Brams and Peter C Fishburn. 1978. Approval Voting. *American Political Science Review* 72, 3 (1978), 831–847. https://doi.org/10.7551/mitpress/9780262015134.003.0018

Florian Brandl, Felix Brandt, and Hans Georg Seedig. 2016. Consistent Probabilistic Social Choice. *Econometrica* 84, 5 (2016), 1839–1880. https://doi.org/10.3982/ecta13337 arXiv:1503.00694

Adrian Bruhin, Helga Fehr-Duda, and Thomas Epper. 2010. Risk and Rationality: Uncovering Heterogeneity in Probability Distortion. *Econometrica* 78, 4 (2010), 1375–1412. https://doi.org/10.3982/ECTA7139

Rudolf Carnap. 1947. Probability as a Guide in Life. *The Journal of Philosophy* 44, 6 (mar 1947), 141. https://doi.org/10.2307/2020041

Royce Carroll, Jeffrey B. Lewis, James Lo, Keith T. Poole, and Howard Rosenthal. 2013. The Structure of Utility in Spatial Models of Voting. *American Journal of Political Science* 57, 4 (2013), 1008–1028. https://doi.org/10.1111/ajps.12029

LE Cederman, Andreas Wimmer, and Brian Min. 2010. Why do ethnic groups rebel? New data and analysis. *World Politics* 1, 1 (2010). http://journals.cambridge.org/abstract{_}S0043887109990219

Frank S. Cohen. 1997. Proportional Versus Majoritarian Ethnic Conflict Management in Democracies. *Comparative Political Studies* 30, 5 (oct 1997), 607–630. https://doi.org/10.1177/0010414097030005004

P. Collier. 2004. Greed and grievance in civil war. *Oxford Economic Papers* 56, 4 (jun 2004), 563–595. https://doi.org/10.1093/oep/gpf064

John G Cross. 1973. A stochastic learning model of economic behavior. *The Quarterly Journal of Economics* 87, 2 (1973), 239–266.

James H Davis. 1992. Some compelling intuitions about group consensus decisions, theoretical and empirical research, and interpersonal aggregation phenomena: Selected examples 1950–1990. *Organizational Behavior and Human Decision Processes* 52, 1 (1992), 3–38.

A Dellis. 2010. Policy moderation and endogenous candidacy in Approval Voting elections. In *Handbook on Approval Voting*, Jean-François Laslier and M. Remzi Sanver (Eds.). Springer, Chapter 18, 431–454. http://link.springer.com/chapter/10.1007/978-3-642-02839-7{_}18

Neil Devotta. 2005. From ethnic outbidding to ethnic conflict: The institutional bases for Sri Lanka's separatist war. *Nations and Nationalism* 11, 1 (2005), 141–159. https://doi.org/10.1111/j.1354-5078.2005.00196.x

Pradeep Dubey and Lloyd S Shapley. 1979. Mathematical properties of the Banzhaf power index. *Mathematics of Operations Research* 4, 2 (1979), 99–131.

Conal Duddy. 2015. Fair sharing under dichotomous preferences. *Mathematical Social Sciences* 73 (2015), 1–5. https://doi.org/10.1016/j.mathsocsci.2014.10.005

Mark S Granovetter. 1978. Threshold Models of Collective Behavior. *Amer. J. Sociology* 83, 6 (1978), 1420–1443. https://doi.org/10.1086/226707

John C Harsanyi. 1973. Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *International journal of game theory* 2, 1 (1973), 1–23.

Jobst Heitzig and Forest W. Simmons. 2012. Some chance for consensus: Voting methods for which consensus is an equilibrium. *Social Choice and Welfare* 38, 1 (nov 2012), 43–57. https://doi.org/10.1007/s00355-010-0517-y

Bradford Jones, Benjamin Radcliff, Charles Taber, and Richard Timpone. 1995. Condorcet Winners and the Paradox of Voting: Probability Calculations for Weak Preference Orders. *American Political Science Review* 89, 1 (1995), 137–144. https://doi.org/10.2307/2083080

Kei Kawai and Yasutora Watanabe. 2013. Inferring strategic voting. *American Economic Review* 103, 2 (2013), 624–662. https://doi.org/10.1257/aer.103.2.624

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case equilibria. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 1563. Springer Verlag, 404–413. https://doi.org/10.1007/3-540-49116-3_38

Rob Lanphier et al. 1996. Election-methods electronic mailing list. *https://electorama.com/em, last accessed 9 October 2020* (1996). https://electorama.com/em/

Jean-François Laslier. 2010. In Silico Voting Experiments. In *Handbook on Approval Voting.* 311–335. https://doi.org/10.1007/978-3-642-02839-7_13

Jean-François Laslier and M Remzi Sanver. 2010. *Handbook on approval voting.* Springer Science & Business Media.

Daniel Lewis. 2013. *Direct democracy and minority rights: A critical assessment of the tyranny of the majority in the American states.* Routledge.

Kenneth O May. 1952. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica: Journal of the Econometric Society* (1952), 680–684.

Nolan McCarty, Keith T Poole, and Howard Rosenthal. 2016. *Polarized America: The dance of ideology and unequal riches.* MIT Press.

Mikael Persson, Peter Esaiasson, and Mikael Gilljam. 2013. The effects of direct voting and deliberation on legitimacy beliefs: An experimental study of small group decision-making. *European Political Science Review* 5, 3 (2013), 381–399. https://doi.org/10.1017/S1755773912000173

Jörg L. Spenkuch, B. Pablo Montagnes, and Daniel B. Magleby. 2018. Backward induction in the wild? evidence from sequential voting in the US senate. *American Economic Review* 108, 7 (2018), 1971–2013. https://doi.org/10.1257/aer.20150993

Peter Troyan. 2012. Comparing school choice mechanisms by interim and ex-ante welfare. *Games and Economic Behavior* 75, 2 (2012), 936–947.

Marc Wiedermann, E. Keith Smith, Jobst Heitzig, and Jonathan F. Donges. 2020. A network-based microfoundation of Granovetter's threshold model for social tipping. *Scientific Reports* 10, 1 (2020). https://doi.org/10.1038/s41598-020-67102-6 arXiv:1911.04126

Douglas R. Woodall. 1997. Monotonicity of single-seat preferential election rules. *Discrete Applied Mathematics* 77, 1 (1997), 81–98. https://doi.org/10.1016/S0166-218X(96)00100-X

F Zakaria. 1997. The rise of illiberal democracy. *Foreign affairs* 76, 6 (1997), 22–43. http://www.jstor.org/stable/20048274

Acknowledgements: We'd like to thank Marius Amrhein, Markus Brill, Pascal Führlich, Anne-Marie George, Ulrike Kornek, Fabrizio Kuruc, Adrian Lison, E. Keith Smith, Lea Tamberg, and the members of the election-methods list (Lanphier et al., 1996) for fruitful discussions and comments, and countless journal editors for their beautifully crafted desk rejection letters.

Funding: This work received no external funding.

Author contributions: J.H. and F.W.S. conceptualized the study, developed theory and methodology, performed the formal analysis, and wrote the manuscript. J.H. developed the software and performed the simulations. The authors declare no competing or conflicting interests.

Data and materials availability: All data and software used is available in the manuscript, the online appendix, or from the authors upon reasonable request. An online voting tool based on the MaxParC method presented here is under open-source development at github.com/mensch72/maxparc-ionic and will be available at vodle.it.

The attached Appendix contains: More Detailed Methods and Results, References (1–23), Table S1–S3, Fig S1–10.

Appendix for paper "Fair and efficient collective decisions via nondeterministic proportional consensus"

Jobst Heitzig[1*] and Forest W Simmons[2]

[1]FutureLab on Game Theory and Networks of Interacting Agents,
Potsdam Institute for Climate Impact Research,
PO Box 60 12 03, D-14412 Potsdam, Germany
[2]Liberal Arts & Mathematics Division, Portland Community College,
Cascade Campus, 705 N. Killingsworth Street, TH 220, Portland, OR 97217, USA

[*]To whom correspondence should be addressed; E-mail: heitzig@pik-potsdam.de.

Contents

1 Method Details

1.1 Summary

1.1.1 Context

A finite group of *voters* must collectively pick exactly one *winning option* out of a given finite number of *options* of any kind (e.g., certain kinds of objects, places, time-points, actions, strategies, people, etc.). The menu of options is already *given* at the beginning of the situation we consider, and all options are mutually exclusive and *feasible,* i.e., each one could be implemented without violating any relevant constraint (e.g., budgets, applicable laws, basic rights, time constraints etc.). Although in the real world, the option menu might sometimes change during a group decision procedure, we consider the composition of the option menu as a separate process here which has been completed before the situation we consider. We assume the voters will apply some *formalized* method to pick one option that may be described as some form of *protocol* or *game form* which requires the voters to provide some kind of information in a step we term *voting* and then determines a winning option from this information in a step we term *tallying,* using some kind of *algorithm* that may or may not involve some form of randomization. We call the information a voter provides this voter's *ballot* and the method that turns sets of ballots into winning options a *voting method.*

In addition, we assume each voter possesses some form of *preferences* regarding the options and regarding possible probability distributions of options. Again, even though in the real world, preferences might sometimes change during a group decision procedure, in particular in certain forms of deliberation or consensus finding, we consider here also the formation of preferences as a separate process which has been completed before the voting. This assumption is in line with the established approach taken in social choice theory. We do *not,* however, assume that there is a simple deterministic relationship between voters' preferences and their ballots, but rather assume that voters may use different kinds of heuristics or strategies to decide how to fill in their ballots.

1.1.2 Problem statement

In this context, we aim at finding a voting method that fulfills certain consistency, fairness, and efficiency *criteria* as stated in the main text and detailed further in Section 2 of this document.

1.1.3 Transdisciplinary methodological approach

To this end, we study the qualitative and quantitative properties of different voting methods, some well-known from the social choice literature, one adapted from the theory of fair budget allocation, and one designed newly. To study the *qualitative properties* of voting methods, we apply a mixture of methods from social choice theory and game theory. For the *quantitative properties* we apply large-scale (Monte-Carlo) numerical simulations of an agent-based model whose assumptions are partially based on the spatial theory of voting from political science and on insights from the study of risk attitudes and bounded rationality in behavioural economics. We analyze simulation results by means of metrics adapted from welfare economics and information theory.

Note on terminology. Because we assume preferences have been determined before voting, we do not in this study distinguish linguistically between the terms *'consensus', 'consent',* and *'compromise',* but rather use a very pragmatic working definition of *potential consensus* here. In this study, consensus does not mean that all voters consider the exact same option from the option menu their favourite option. Informally, we rather say there is (full or partial) *potential consensus* whenever there is some option which (all or some group of) voters would prefer to having one randomly drawn voter make the choice.

1.2 Ballot types and voting methods

In this section we introduce a formal mathematical framework for comparing quite different voting methods and then use it to define our versions of a number of common group decision methods, at which point the motivations for the various abstract notions should become clear.

We assume an infinite universe of potential voters $\mathcal{I}$ and an infinite universe of potential options $\mathcal{X}$, leading to a universe of potential finite electorates $\mathcal{E} = \{E \subset \mathcal{I} : 1 \leq |E| < \infty\}$ and a universe of potential finite choice sets $\mathcal{C} = \{C \subset \mathcal{X} : 1 \leq |C| < \infty\}$. For each choice set

$C \in \mathcal{C}$, let $L(C) = \{\ell \in [0,1]^C : \sum_{x \in C} \ell(x) = 1\}$ be the set of all *lotteries* on C, and let $\ell_x \in L(C)$ with $\ell_x(x) = 1$ be the *sure-thing* lottery that picks $x \in C$ for sure. In most of what follows we deal with fixed sets E and C and denote their sizes by N and k, but for some proofs we have to consider all of $\mathcal{E}$ and $\mathcal{C}$.

1.2.1 Ballot types

Since we will deal with voting methods using quite different types of ballots, some only letting the voter mark a single option, others many, still others requiring a strict ranking or asking for quantitative ratings or the like, we need a formal framework general enough to cover all relevant cases and make them comparable in those respects important for the assessment of the method's properties. In particular, we need to make clear for each ballot type what it means when we say that a ballot states a preference for one option over another or a ballot results from another ballot by advancing an option to a certain degree. The following abstract definition will prove useful in these tasks:

A *ballot type* is a tuple (B, P, Q) with the following properties:

- B is a function such that for all choice sets $C \in \mathcal{C}$, $B(C)$ is a nonempty set representing the different ways b in which a voter might fill in a ballot for the choice set C, and $B(C) \cap B(C') = \emptyset$ if $C \neq C'$.
- P is a function such that for all $C \in \mathcal{C}$ and all filled-in ballots $b \in B(C)$, P_b is a strict partial ordering relation on C (i.e., irreflexive, asymmetric, and transitive, but not necessarily complete) representing that part of ballot b that will be interpreted as *stated preferences*, with $x\,P_b\,y$ meaning that x is put in a strictly "better place" (ranking, rating, threshold, etc.) on b than y is.
- Q is a function such that for all $C \in \mathcal{C}$ and all options $x \in C$, Q^C_x is a strict partial ordering relation on $B(C)$ with $b\,Q^C_x\,b'$ meaning that the two filled-in ballots $b, b' \in B(C)$ only differ in the fact that b puts x in a strictly "better place" than b' while each other option is treated the same on b and b', so that b can be seen as resulting from b' by "advancing" x in some way and changing nothing else.

Note that $P_b = P_{b'}$ does not imply $b = b'$ in general since some ballot types (e.g., ratings ballots) also contain other information than just a binary preference relation.

If $C \in \mathcal{C}$, $b \in B(C)$, and $x\,P_b\,y$ for all $y \in C \setminus \{x\}$, we call x the *stated favourite* on b and write $F(b) = x$. Since P_b may be incomplete but is asymmetric, a ballot contains either no stated favourite (in which case we write $F(b) = \emptyset$) or exactly one. If $F(b) = x$ and $\neg y\,P_b\,z$ for any $y, z \neq x$, we say that b is a *bullet vote* for x.

1.2.2 Ballot profiles and voting methods.

A *ballot profile* of type (B, P, Q) for electorate $E \in \mathcal{E}$ and choice set $C \in \mathcal{C}$ is a function $\beta : E \to B(C)$ specifying a filled-in ballot β_i for each voter $i \in E$.

A (potentially probabilistic) *voting method* is a tuple (B, P, Q, M) such that (B, P, Q) is a ballot type and M is a function such that for all $E \in \mathcal{E}$ and $C \in \mathcal{C}$, and all ballot profiles β of type (B, P, Q) for E and C, it specifies a *winning lottery* $M(\beta) \in L(C)$. Of course, $M(\beta)$ may be a sure-thing lottery with $M(\beta)_x = 1$ for some $x \in C$.

1.2.2.1 Plurality Voting (PV). We formalize the ballot type of *Plurality Ballot* as follows:

- $B(C) = C$, i.e., each voter has to *vote for* exactly one option $x \in C$ by putting $b = x$.
- $x\, P_b\, y$ iff $b = x \neq y$, i.e., voting for x is interpreted as stating a strict preference for x over all other options.
- $b\, Q^C_x\, b'$ iff $b = x \neq b'$, i.e., "advancing" x means converting a vote for a different option into a vote for x.

Using this ballot type, the voting method of *Plurality Voting* now puts $M(\beta)_x = 1_A(x)/|A|$, where 1_A is the indicator function of the set $A = \{x : p(x) \geq p(y) \text{ for all } y \in C\}$ of *plurality winners,* and $p(x) = |\{i \in E : \beta_i = x\}|$ is x's *plurality score*.

Note that for simplicity, in our version one cannot abstain under plurality voting, and hence every ballot is a bullet vote. Generally $|A| = 1$ except for ties, which means that this is a "deterministic" method in our terminology.

1.2.2.2 Approval Voting (AV). We formalize the ballot type of *Approval Ballot* as follows:

- $B(C) = \{0, 1\}^C$, i.e., each voter can either *approve* (by putting $b(x) = 1$) or *disapprove* (by putting $b(x) = 0$) each option $x \in C$ individually.
- $x\, P_b\, y$ iff $b(x) > b(y)$, i.e., iff $b(x) = 1$ and $b(y) = 0$, meaning that approving x is interpreted as stating a strict preference for x over all non-approved options.
- $b\, Q^C_x\, b'$ iff $b(x) > b'(x)$ and $b(y) = b'(y)$ for all $y \in C \setminus \{x\}$, i.e., "advancing" x means converting a non-approval of x into an approval of x.

Using this ballot type, the voting method of *Approval Voting* now puts $M(\beta)_x = 1_A(x)/|A|$, where 1_A is the indicator function of the set $A = \{x : a(x) \geq a(y) \text{ for all } y \in C\}$ of *approval winners* and $a(x) = \sum_{i \in E} \beta_i(x)$ is x's *approval score*.

There are two equivalent ways to "abstain" under approval voting: putting $b(x) \equiv 0$ for all $x \in C$ or putting $b(x) \equiv 1$ for all $x \in C$. Bullet voting for x means putting $b(x) = 1$ and $b(y) = 0$ for all other y.

1.2.2.3 Range Voting (RV). In our version of range voting, the ballot type of *Range Ballot* has:

- $B(C) = [0, 100]^C$, i.e., one can assign any real-valued *rating* $0 \leq b(x) \leq 100$ to each option $x \in C$ individually.[1]
- $x P_b y$ iff $b(x) > b(y)$ as before, meaning that a higher rating states a strict preference.
- $b Q_x^C b'$ iff $b(x) > b'(x)$ and $b(y) = b'(y)$ for all $y \in C \setminus \{x\}$, i.e., "advancing" x means raising its rating.

Using this ballot type, the voting method of *Range Voting* puts $M(\beta)_x = 1_R(x)/|R|$, where 1_R is the indicator function of the set $R = \{x : r(x) \geq r(y) \text{ for all } y \in C\}$ of *range winners* and $r(x) = \sum_{i \in E} \beta_i(x)$ is x's *range score*.

There are infinitely many equivalent ways to "abstain" under range voting: choose some $\alpha \in [0, 100]$ and put $b(x) \equiv \alpha$ for all $x \in C$. There are also infinitely many ways to "bullet vote" for x under range voting, and they are not (!) equivalent: choose $0 \leq \alpha < \gamma \leq 100$ and put $b(x) = \gamma$ and $b(y) = \alpha$ for all other $y \in C$.

In addition to the above three "scoring methods", we consider the following two more complicated "ranking methods".

1.2.2.4 Instant-Runoff Voting (IRV). In our version of instant-runoff voting, we use the following ballot type of *Truncated Ranking Ballot*:

- $B(C) = \left\{b \in (\mathbb{N} \cup \{\infty\})^C : b[b^{-1}[\mathbb{N}]] = \{1, \dots, |b^{-1}[\mathbb{N}]|\}\right\}$. In other words, one has to assign consecutive and distinct integer *ranks* $1, 2, \dots$ to any empty or nonempty subset of the options, leaving the other options unranked (here formally encoded by "rank" ∞).
- $x P_b y$ iff $b(x) < b(y)$ (!) since a smaller rank number indicates a "better place".
- $b Q_x^C b'$ iff $P_b|_{C\setminus\{x\}} = P_{b'}|_{C\setminus\{x\}}$,[2] $x P_b y$ whenever $x P_{b'} y$, $y P_{b'} x$ whenever $y P_b x$, but $P_b \neq P_{b'}$. In other words, advancing x means changing the ranks so that the resulting ordering P doesn't change except that some options ranked better than x before are now either ranked lower than x or not ranked at all, and/or x was not ranked before and is now ranked better than at least one option.

Using this ballot type, our simple version of the voting method of *Instant-Runoff Voting* (aka Single Transferable Vote, or Alternative Vote) runs like this: Initialize $D = C$ and repeat the following as long as $|D| > 1$: For each option $x \in D$, calculate the score $s(x) = \Big|\Big\{i \in E : \beta_i(x) <$

[1] Voter i's choice of $b(x)$ is what was denoted r_{ix} in the main text.

[2] If $R \subseteq X \times X$ is a binary relation on some set X and $S \subseteq X$ is a subset of X, then $R|_S = R \cap (S \times S)$ is the restriction of R to S.

$\beta_i(y)$ for all $y \in D \setminus \{x\}\}\Big|$. From the set of worst-scored options, $W = \{x \in D : s(x) \leq s(y)$ for all $y \in D\}$, draw a random member and remove it from D. The remaining member of D wins.

Abstention means not ranking any option (i.e., putting $b_x = \infty$ for all x). A bullet vote ranks exactly one option (i.e., puts $b_x = 1$ for some x and $b_y = \infty$ for all other y).

1.2.2.5 Simple Condorcet (SC). In our version of the Simple Condorcet method, we use the following ballot type of *Weak Ranking Ballot:*

- $B(C) = (\mathbb{N} \cup \{\infty\})^C$. In other words, one can assign arbitrary integer *ranks* $1, 2, \ldots$ to any empty or nonempty subset of the options, leaving the other options unranked (again encoded by ∞).
- $x\,P_b\,y$ iff $b(x) < b(y)$ as in IRV.
- $b\,Q^C_x\,b'$ iff $P_b|_{C\setminus\{x\}} = P_{b'}|_{C\setminus\{x\}}$, $x\,P_b\,y$ whenever $x\,P_{b'}\,y$, $y\,P_{b'}\,x$ whenever $y\,P_b\,x$, but $P_b \neq P_{b'}$. In other words, advancing x means changing the ranks so that the resulting ordering P doesn't change except that some options ranked better than x before are now either ranked equal to or lower than x or not ranked at all, and/or some options ranked equal to x before are now ranked lower than x or not ranked at all,

Note that we allow both that some options are ranked equal and some rank numbers are skipped since the only information our version of the *Simple Condorcet* method (aka the Minimax Condorcet or Simpson–Kramer method with pairwise opposition as score) is the ordering P_b. We put $M(\beta)_x = 1_O(x)/|O|$, where 1_O is the indicator function of the set $O = \{x : o(x) \leq o(y)$ for all $y \in C\}$ of *weak condorcet winners*, $o(x) = \max_{y\in C} o(x, y)$ is x's *worst opposition value*, and $o(x, y) = |\{i \in E : y\,P_{\beta_i}\,x\}|$ for all $x, y \in C$.

There are infinitely many equivalent ways to "abstain" under the Simple Condorcet method: choose some $\alpha \in \mathbb{N} \cup \{\infty\}$ and put $b(x) \equiv \alpha$ for all $x \in C$. The most straightforward is to not rank any option at all. There are also infinitely many equivalent ways to "bullet vote" for x under the Simple Condorcet method: choose $\infty \geq \alpha > \gamma \in \mathbb{N}$ and put $b(x) = \gamma$ and $b(y) = \alpha$ for all other $y \in C$.

After these five "deterministic" methods, we now turn to five "non-deterministic" methods, beginning with three from the literature.

1.2.2.6 Random Ballot (RB). The *Random Ballot* (aka Random Dictator) method uses Plurality Ballots but puts $M(\beta)_x = p(x)/N$. The interpretation is that one ballot is drawn uniformly at random to decide.

1.2.2.7 Full Consensus / Random Ballot (FC). An *FC Ballot* is basically a combination of two Plurality Ballots:

- $B(C) = C \times C$, i.e., each voter specifies one "proposed consensus" $b^1 \in C$ and one "fall-back" option $b^2 \in C$.
- $x\,P_b\,y$ iff $b^2 = x \neq y$, i.e., only the fall-back part of the ballot is interpreted as stating a strict preference for b^2 over all other options, while the consensus part is interpreted as inherently strategic.
- $b\,Q^C_x\,b'$ iff $b \neq b'$ and ($b^1 = x \neq b'^1$ or $b^1 = b'^1$) and ($b^2 = x \neq b'^2$ or $b^2 = b'^2$), i.e., "advancing" x means advancing it in at least one of the two ballot parts.

The method of *Full Consensus / Random Ballot (FC)* is now defined as in [1] (there called "Voting method 1"): $M(\beta)_x = 1$ if $\beta^1_i = x$ for all $i \in E$; otherwise $M(\beta)_x = p_2(x)/N$ for all $x \in C$ ("fall-back lottery"), where $p_2(x) = |\{i \in E : \beta^2_i = x\}|$ is x's *fall-back score*. The interpretation is that if all voters propose the same consensus, that option wins, otherwise the fall-back lottery applies.

A bullet vote is a bullet vote on both ballot parts. There is no way to abstain.

1.2.2.8 Full Consensus / Random Ballot / Ratings (RFC). Similarly, an *RFC Ballot* is a combination of two Plurality Ballots and a Range Ballot:

- $B(C) = C \times C \times [0, 100]^C$, i.e., each voter specifies one proposed consensus $b^1 \in C$, one fall-back option $b^2 \in C$, and a vector of ratings $b^3(x) \in [0, 100]$ for all $x \in C$.
- $x\,P_b\,y$ iff $b^3(x) > b^3(y)$, i.e., only the ratings part of the ballot is interpreted as stating preferences, while the other two parts are interpreted as inherently strategic.
- $b\,Q^C_x\,b'$ iff $b \neq b'$ and ($b^1 = x \neq b'^1$ or $b^1 = b'^1$) and ($b^2 = x \neq b'^2$ or $b^2 = b'^2$) and $b^3(x) \geq b^3(y)$, i.e., "advancing" x means advancing it in at least one of the three ballot parts.

The method of *Full Consensus / Random Ballot / Ratings (RFC)* is also defined as in [1] (there called "Voting method 2"). For all $j \in E$, let $r_j = \sum_{y \in C} p_2(y)\beta^3_j(y)/N$ be the rating of the fall-back lottery by voter j. Then put $M(\beta)_x = \frac{|A_x|}{N} + (1 - \frac{|A|}{N})\frac{p_2(x)}{N}$, where A_x is the set of all $i \in E$ for which $\beta^1_i = x$ and $\beta^3_j(x) \geq r_j$ for all $j \in E$ (i.e., whose proposed consensus is x and is preferred to the fall-back lottery by all voters according to their ratings), and $A = \bigcup_{y \in C} A_y$. The interpretation is that a voter i is drawn uniformly at random, and if i's proposed consensus β^1_i is unanimously preferred to the fall-back lottery $p_2(x)/N$, it wins, otherwise the fall-back lottery is applied.

A bullet vote is a bullet vote on all three ballot parts. There is no way to abstain.

1.2.2.9 Nash Lottery (NL).

The *Nash Lottery* method uses Range Ballots. Given β, $i \in E$, and $\ell \in L(C)$, let $r_i(\ell) = \sum_{x \in C} \ell(x)\beta_i(x)$ and $S(\ell) = -\sum_{i \in E} \log r_i(\ell)$. If there is a unique $\ell \in L(C)$ with $S(\ell) > S(\ell')$ for all $\ell' \in L(C) \setminus \{\ell\}$, we put $M(\beta) = \ell$. In the rare cases where $\arg\max_{\ell \in L(C)} S(\ell)$ is not a singleton, we use that ℓ which our numerical optimizer (the `minimize` function from the `scipy.optimize` Python package with method 'SLSQP') returns.

For formal theoretical analyses, one can use the following tie-breaker instead. Put $r_i^k = \sum_x \ell(x) \sqrt[k]{\beta_i(x)}$ for all $k = 1, 2, 3, \ldots$, and $S^k(\ell) = \sum_i \log r_i^k(\ell) \in [-\infty, \infty)$. Note that all S^k are continuous, continuously differentiable, and weakly concave functions of ℓ. Hence S^1 has a global maximum that is attained on a non-empty compact convex set $T^1 \subseteq L(C)$, and for all $k \geq 2$, S^k restricted to T^{k-1} has a global maximum that is attained on a non-empty compact convex set $T^k \subseteq T^{k-1}$. Then also $T = \bigcap_{k=1}^{\infty} T^k$ is non-empty compact convex, hence Lebesgue-measurable, and hence has a well-defined unique centre of mass $\ell = \int_T \ell d\ell / \int_T d\ell$ with $\ell \in T$ because of the convexity. We now put $M(\beta) = \ell$. The rationale for using concave functions of ratings to break ties is that in this way lotteries with lower entropy are preferred. The rationale for using the k-th square roots for this task is that in this way the tie-breaking is complete except in the case of clones (see below).

A bullet vote is to rate one option at > 0 and all others at 0, abstention means rating all options at the same value > 0.

In the context of "dichotomous preferences", a similar method based on Approval Ballots was studied in [2] under the name "Nash Max Product". The same idea is also common in the literature on fair division [3].

1.2.2.10 Maximal Partial Consensus (MaxParC, MPC).

For our simulations, we use this version of *MaxParC ballots:*

- $B(C) = [0, 100]^C$, i.e., one can assign any real-valued *willingness to approve* $0 \leq b(x) \leq 100$ to each option $x \in C$ individually.
- $x P_b y$ iff $b(x) = 100 > b(y)$ or $b(x) > 0 = b(y)$, i.e., we only interpret the special values 100 and 0 as "stated preferences" and treat all intermediate values as inherently strategic since their interpretation relates to other voters' willingnesses.
- $b\, Q_x^C\, b'$ iff $b(x) > b'(x)$ and $b(y) = b'(y)$ for all $y \in C \setminus \{x\}$, i.e., "advancing" x means raising the willingness to approve it.

The *Maximal Partial Consensus (MaxParC, MPC)* method now works as follows. A voter approves an option if enough other voters do so as well; a non-abstaining voter i is drawn uniformly at random; then from the highest-scoring options approved by i, one is drawn uniformly at random. Once it is decided who approves what, the procedure corresponds to what is described in [4], page 4 (last paragraph of section 3), and analysed in [2] under the name 'Conditional Utilitarian Rule'.

To define this formally, we will introduce the following mathematical objects. The set $A(x)$ will be the set of voters i that turn out to approve x since their willingness to approve x, $\beta_i(x)$, is properly larger than $100 \times (1 - |A(x)|/N)$. The quantity $a'(x)$ will be x's approval score plus a fractional part used for tiebreaking. The set A_i will be the set of options approved by i, and A'_i the set of highest-scoring options in A_i. Finally, $A(\emptyset)$ will be those voters who don't approve any option and thus "effectively abstain".

Formally, their definition is this: For all $x \in C$, let $A(x)$ be the largest subset $A \subseteq E$ such that $|A|/N + \beta_i(x)/100 > 1$ for all $i \in A$. Let $a'(x) = |A(x)| + \sum_{i \in E} \beta_i(x)/100N$. For all $i \in E$, put $A_i = \{x \in C : i \in A(x)\}$ and $A'_i = \arg\max_{y \in A_i} a'(y)$. Finally, put $A(\emptyset) = E - \bigcup_{x \in E} A(x)$. Then $M(\beta)_x = \sum\{1/|A'_i| : i \in E \text{ with } x \in A'_i\}/(N - |A(\emptyset)|)$.

A bullet vote is to rate one option at 100 and all others at 0, while abstention means rating all options at 0.

Note: since $A(x)$ can be found in $N \log N$ time, the total tallying complexity is $O(kN \log N)$.

1.3 Agent-based simulations experiments

1.3.1 Modeling individual preferences

We simulate voter preferences by using one of several different models to generate a profile of individual utility functions over options, $u_i : C \to \mathbb{R}$, and then derive individual utility functions over lotteries depending on each voters risk attitude type. Our *utility models* are the following.

1.3.1.1 Uniform model (Unif) In this simplest non-spatial utility model, each value $u_i(x)$ is drawn uniformly at random from the unit interval $[0, 1]$. The resulting preference orderings form what is usually called the *impartial culture* model [5].

1.3.1.2 Block model (BM) In this non-spatial utility model, there are $r \geq 1$ *voter blocks* whose expected relative sizes $s_1, \dots, s_r$ are drawn independently from a log-normal distribution such that $\ln s_j \sim N(0, h)$, where $h \geq 0$ is a *block size heterogeneity* parameter. In particular, if $h = 0$, all blocks are of similar size, while larger values of h will lead to ever smaller minorities.

For each voter i independently, the probability to belong to block j is then $s_j / \sum_{j'=1}^{r} s_{j'}$. Let $J(i)$ be i's block. Then the utility $u_i(x)$ that voter i would get from option x is now a sum of a block-dependent component and an individual component,

$$u_i(x) = U_{J(i)}(x) + \iota \varepsilon_i(x), \tag{1}$$

where all $U_{J(i)}(x)$ and $\varepsilon_i(x)$ are independent standard normal variables and $\iota > 0$ is an *individuality* parameter.

1.3.1.3 Spatial preference models In the spatial theory of voting [6] (also called "spatial cultures" in [5]), voters i and options x are represented by *ideal points* (or bliss points) η_i and *positions* ξ_x in a low-dimensional *policy space* $\mathbb{R}^d$, $d \geq 1$, and the utility $u_i(x)$ that voter i would get from option x depends in a monotonically decreasing fashion on the distance between η_i and ξ_x. We distinguish the following spatial voting models:

Linear homogeneous (LH) model. Utilities are decreasing linearly with distance,

$$u_i(x) = -\|\eta_i - \xi_x\|_1 \tag{2}$$

where the ξ_x are distributed independently and uniformly on the cube $[-1, 1]^d$, and the η_i are distributed independently and uniformly on the cube $[-\omega, \omega]^d$, where $\omega > 0$ is a *voter heterogeneity* parameter.

Quadratic homogeneous (QH) model. Utilities decrease quadratically with distance,

$$u_i(x) = -\|\eta_i - \xi_x\|_2^2, \tag{3}$$

the ξ_x are distributed independently according to the multivariate standard normal distribution, and the η_i according to the symmetric multivariate normal distribution with zero mean and standard deviation ω.

Gaussian homogeneous (GH) model. As in the quadratic homogeneous model, but with Gaussian utilities

$$u_i(x) = e^{-\|\eta_i - \xi_x\|_2^2/2\sigma^2} \tag{4}$$

for some $\sigma > 0$.

In addition to the above, rather classical spatial models, we also use the following three variants, which introduce some idea borrowed from what [5] calls "distributive cultures":

Gaussian allotment (GA) model. As in the Gaussian homogeneous model, but with each option having a different standard deviation $\sigma_x > 0$, so that

$$u_i(x) = e^{-\|\eta_i - \xi_x\|_2^2/2\sigma_x^2}/(\sqrt{2\pi}\sigma_x)^d. \tag{5}$$

The interpretation is that each option x allots a unit amount of total utility to all potential ideal points of voters using a symmetric multivariate normal distribution whose standard deviation σ_x represents the *broadness* of option x's "platform". Because of the normalization factor σ_x^{-d}, if two options have very close positions but different broadness, voters close to their position will prefer the "narrower" option and voters farther away will prefer the "broader" option.

Quadratic allotment (QA) model. As in the Gaussian allotment model, but with log-transformed utilities, resulting in a quadratic functional form:

$$u_i(x) = -\|\eta_i - \xi_x\|_2^2/2\sigma_x^2 - d\ln(\sqrt{2\pi}\sigma_x). \tag{6}$$

The interpretation is that here option x allots a unit amount of total wealth instead of a unit amount of total utility, and voters' utility is logarithmic in wealth.

Linear allotment (LA) model. As in the quadratic allotment model, but with ξ_x and η_i distributed on cubes as in the linear homogeneous model, and with a linearly decreasing utility:

$$u_i(x) = -\|\eta_i - \xi_x\|_1/\sigma_x - d\ln(2\sigma_x). \tag{7}$$

The interpretation is that each option x allots a unit amount of total wealth to all potential ideal points of voters using a symmetric multivariate exponential distribution with density $e^{-\|\eta_i-\xi_x\|_1/\sigma_x}/(2\sigma)^d$, and that utility is logarithmic in wealth.

Distribution of options' broadnesses. In the three allotment models, we draw the options' broadnesses σ_x independently from a log-normal distribution such that $\ln\sigma_x \sim N(\ln\sigma_0, \varrho)$, where $\sigma_0 > 0$ is the *median broadness* and $\varrho \geq 0$ is a *broadness heterogeneity* parameter. The three homogeneous models are then equivalent to the case $\varrho = 0$ of the allotment models.

1.3.1.4 Utility of uncertain prospects Regarding their preferences over uncertain prospects, represented as proper lotteries $\ell \in L(C)$ over options, we assume each voter is of one of three *risk attitude types* that determine how they derive utility functions over lotteries from their utility functions over options.

Expected utility theory. We assume voters of *expected utility theory (EUT)* type evaluate the utility of a lottery of options $\ell \in L(C)$ by taking the expected value of the individual options' utilities, $u_i(\ell) = \sum_{x\in C}\ell(x)u_i(x)$. To see a major qualitative difference between the above linear, quadratic, and Gaussian models, consider the one-dimensional case of three options placed symmetrically at $\xi_A = -1$, $\xi_B = 1$, $\xi_C = 0$ with $\sigma_x \equiv 1$, and compare the utility a non-central voter at $\eta_i \geq 2$ will assign to the "compromise" option C and to the "polar" lottery $\ell = (A+B)/2$ (tossing a coin to decide between A and B). In the LH model, $u_i(C) = u_i(\ell)$, i.e., the voter is indifferent between the compromise option and the polar lottery. In the QH model, $u_i(C) > u_i(\ell)$, i.e., the voter prefers the compromise. In the GH model, $u_i(C) < u_i(\ell)$, i.e., the voter prefers the polar lottery. More generally, this means that the quadratic/Gaussian models tend to have a larger/smaller number of potential consensus options than the linear models, respectively.

Cumulative prospect theory. Motivated by recent empirical evidence [7], we assume that only about 20 percent of voters are of EUT type, while the remaining 80 percent evaluate ℓ instead as follows.

40 percent are of *low-expectations cumulative prospect theory (LCP)* type. Such a voter i treats all options she prefers to her least-desired one as "gains", hence sorts the options by descending utility, $u_i(x_1) \geq u_i(x_2)\ldots \geq u_i(x_k)$, calculates the cumulative probabilities $c_j =$

$\sum_{j'=1}^{j} \ell(x_{j'})$, so that $c_0 = 0$ and $c_k = 1$, and then evaluates ℓ as

$$u_i(\ell) = \sum_{j=1}^{k} w_j u_i(x_j), \tag{8}$$

where $w_j = W(c_j) - W(c_{j-1})$, $W(c) = \delta c^\gamma/(\delta c^\gamma + (1-c)^\gamma)$ is the *probability weighting function* with $W(0) = 0$ and $W(1) = 1$, and we choose $\delta = 0.926$ and $\gamma = 0.377$ following the pooled group estimates for gains from [7].

The remaining 40 percent are of *high-expectations cumulative prospect theory (HCP)* type. Such a voter i treats all options except her favourite one as "losses", hence sorts the options by *ascending* utility, $u_i(x_1) \leq u_i(x_2) \ldots \leq u_i(x_k)$, calculates the cumulative probabilities $c_j = \sum_{j'=1}^{j} \ell(x_{j'})$, so that $c_0 = 0$ and $c_k = 1$, and then evaluates ℓ as

$$u_i(\ell) = \sum_{j=1}^{k} w_j u_i(x_j), \tag{9}$$

where w, W are as above, but now with $\delta = 0.991$ and $\gamma = 0.397$ following the pooled group estimates for losses from [7].

Generic utilities. Note that due to the involvement of independent continuously distributed utility components, in all our utility models the resulting utility functions u_i will be generic with probability one in the following sense. Different rational-valued lotteries $\ell \neq \ell' \in L(C) \cap \mathbb{Q}^C$ will have different utilities $u_i(\ell) \neq u_i(\ell')$, so that each voter will have strict preferences over all pairs of rational-valued lotteries. In particular, each voter will have a unique *favourite* option $f_i = \arg\max_{x \in C} u_i(x) \in C$.

1.3.2 Voting behaviour

1.3.2.1 Polling and final voting rounds. In our simulation model, the actual voting round is preceded by a number $R > 0$ of polling rounds. In each polling round, voters are asked to name their favourite and all options they "approve" of, and the total *favourite polling scores* $f^p(x)$ and *approval polling scores* $a^p(x)$ are published so that voters can base their voting behaviour in later polling rounds and the final voting round on this information. The actual voting round is then assumed to consist of an initial ballot that can then be changed for some time in an interactive phase as a response to the current ballots' tallying results, so that our setup allows for the simulation of the emergence of strategic equilibria.

We assume that each voter i is of either of five *behavioural types* $\tau(i)$: sincere, lazy, heuristic, trial-and-error, or factional. Heuristic and factional voters together form the set of "strategic" voters, which we assume to make up about half the electorate.[3] Trial-and-error and factional voters together form the set of "interactive" voters who will potentially change their ballots during the interactive phase. For lack of empirical data on interactive voting systems, we assume that these also make up about half of the electorate, while the rest will stick to their ballots during the interactive phase. Lazy voters, filling in their ballots in the simplest possible and non-strategic way, are assumed to make up about one sixth of the electorate.[4] We assume that about another sixth is sincere and fills in their ballots non-strategically to best represent their actual preferences. Together with the above assumptions, this implies that also about one sixth is heuristic, using a simple form of strategic reasoning, about one sixth is of trial-and-error type, starting sincerely but testing simple modifications during the interactive phase, and about one third is of factional type, starting heuristically and following best-response strategies proposed by their faction leaders during the interactive phase. Although some studies also suggest that some voters do not sufficiently understand elections in order to vote sincerely or at least properly lazily, but will rather vote more or less erratically, we do not include an erratic type here since it would only increase the noise in the data.

In addition to this "middle"scenario, we test two further scenarios, one "strategic" and one "lazy",[5] with behavioural types distributed according to the probabilities listed here:

	behavioural types scenario							
type	**lazy**	**middle**	**strategic**	**all-L**	**all-S**	**all-T**	**all-H**	**all-F**
L (lazy)	1/3	1/6	1/20	1	0	0	0	0
S (sincere)	1/3	1/6	1/20	0	1	0	0	0
T (trial-and-error)	1/9	1/6	1/5	0	0	1	0	0
H (heuristic)	1/9	1/6	1/5	0	0	0	1	0
F (factional)	1/9	1/3	1/2	0	0	0	0	1

[3]One of the few countries in which the election outcome can be used for a rough assessment of the percentage of voters who take into account strategic reasoning is Germany because of the strategic incentive to "split vote" by voting for different parties with your first and second votes. Several studies show that in recent parliamentary elections about half of the voters who had an incentive to "split vote" because their favoured party had no chance of winning a direct mandate actually did split their vote ([8], p.17). Using an elaborate methodology, [9] classified voters in Germany's 2013 parliamentary elections into several behavioural types and found that about 15.9 per cent had an incentive to split and did split, while 11.9 per cent had an incentive to split and didn't split ([9] Table 1), i.e., 57 per cent of those with an incentive to vote strategically did so.

[4]One of the few systems in which the election outcome can be used for a rough assessment of the percentage of voters who supply less information on their ballot than would be advisable is the Single Transferable Vote system used in Ireland's parliamentary elections because voters may keep their submited ranking so short that during the iterative tallying process their vote gets "exhausted" and thus essentially wasted. Election outcomes suggest that between 10 and 25 per cent of voters are "lazy" in this sense (the number of exhausted votes can be calculated easily from public data, e.g., on `https://en.wikipedia.org/wiki/Dublin_Central_%28D%C3%A1il_%C3%89ireann_constituency%29`, by comparing the elected candidates total vote turnout with the number of cast ballots)

[5]In many US elections, voters may use a so-called "straight ticket" which might be interpreted as indicating a certain level of lazyness. As there are often up to or even more than half of all voters using straight ticket voting, we assume that in the "lazy" scenario one third of the voters are lazy.

We now specify our behavioural assumptions for the five types.

1.3.2.2 Sincere voters ($\tau(i) = S$). A sincere voter fills in her ballot b in a certain way that represents her "true" preferences, in particular so that the stated preferences P_b are compatible with her true preferences in the sense that $x \, P_b \, y$ implies $u_i(x) > u_i(y)$. Since many ballot types allow for more than one way of sincere voting, we make the following explicit assumptions for the different voting methods. Several of these make use of the *benchmark lottery* $\ell \in L(C)$ whose winning probabilities are proportional to the latest favourite polling scores, $\ell(x) = f^p(x)/N$, and on its expected utility $u_i(\ell) = \sum_{x \in C} \ell(x) u_i(x)$. The MaxParC sincere strategy also makes use of an estimate $\alpha \in [0, 1]$ of the proportion of lazy voters in the electorate.

- In the first polling round, she names her true favourite and approves all x with above-average utility, using equal weights for all options.
- In later polling rounds, she names her true favourite and approves all x with above-average utility, using weights based on the latest favourite polling scores.
- In the actual voting round, her initial and final ballot is determined like this:
  - In Plurality and Random Ballot, she marks her true favourite: $b = f_i$.
  - In Approval Voting, she marks all x with at-least-average utility, where the average is weighted with the latest favourite polling scores so that "approval" is with respect to the benchmark of the currently most relevant seeming options: $b(x) = 1$ iff $u_i(x) \geq u_i(\ell)$.
  - In Range Voting and the Nash Lottery, she assigns ratings from 0 to 100 proportional to utility: $b(x) = 100 \frac{u_i(x) - \min_{y \in C} u_i(y)}{\max_{y \in C} u_i(y) - \min_{y \in C} u_i(y)}$.
  - In IRV and Simple Condorcet, she ranks all x with at-least-average utility as in Approval Voting, in correct order of preference: $b(x) = |\{y \in C : u_i(y) \geq u_i(x)\}|$ iff $u_i(x) \geq u_i(\ell)$, else $b(x) = \infty$.
  - In FC, she marks her true favourite as "favourite" and marks that option as "consensus" which has the highest approval polling score among those options she herself approves: $b = (\arg\max_{x \in C} u_i(x), \arg\max_{x \in C,\, u_i(x) \geq u_i(\ell)} a^p(x))$.[6]
  - In RFC, she combines a sincere FC ballot with a sincere Range Voting ballot.
  - In MaxParC, she assigns a willingness of 0 to all non-approved options, and willingness values from 100α to 100 scaling linearly with utility for all other options: $b(x) = 0$ if $u_i(x) < u_i(\ell)$, else $b(x) = 100\left(\alpha + (1 - \alpha) \frac{u_i(x) - u_i(\ell)}{\max_{y \in C} u_i(y) - u_i(\ell)}\right)$.[7]

[6] In the rare cases where several $a^p(x)$ are equal, we use $f^p(x)$ as a first-order tie-breaker and $u_i(x)$ as a second-order tie-breaker.

[7] This is the simplest sincere voting heuristic for MaxParC that (i) guarantees that my share of winning prob-

1.3.2.3 Lazy voters ($\tau(i) = L$). A lazy voter marks or ranks (only) her true favourite in Plurality, Random Ballot, Approval Voting, IRV, and Simple Condorcet, marks the same option as consensus in FC and RFC, and gives it a rating/willingness of 100 and all others a rating/willingness of zero in Range Voting, the Nash Lottery, RFC, and MaxParC ("bullet voting").

1.3.2.4 Heuristic voters ($\tau(i) = H$). Heuristic voters try to adjust their voting behaviour to that of the other voters in order to increase the chances of preferred options and avoid "wasting their vote". But since their information is restricted to polling scores, they can only act boundedly rational. In addition, we assume they do not employ full optimization given that data but rather use more or less simple or moderately complex "heuristic" strategies [10] mainly based on the idea of "exaggerating" their stated preferences regarding the two options between which a nip-and-tuck race seems most likely [11, 12], and possibly taking into account next-most likely nip-and-tuck races as well. For the more complex voting methods, we do however allow for heuristics that require basic computational tasks such as forming sums, products and ratios and following simple decision trees.

- In polling rounds, she acts as in Plurality and Approval Voting, while in the actual voting round, her initial and final ballot is determined as follows.
- In Plurality, she marks her preferred option among the two best-placed in the latest favourite polling scores: $b = y$ if $u_i(y) > u_i(z)$, else $b = z$, where $y = \arg\max_{x\in C} f^p(x)$ and $z = \arg\max_{x\in C\setminus\{y\}} f^p(x)$.[8]
- In Approval Voting, she marks all x she prefers to the option y leading the latest approval polling scores, and marks y iff she prefers y to the runner-up z in the latest approval polling scores: $b(x) = 1$ iff ($u_i(x) > u_i(y)$ or $x = y$ and $u_i(y) > u_i(z)$), where $y = \arg\max_{x\in C} a^p(x)$ and $z = \arg\max_{x\in C\setminus\{y\}} a^p(x)$.[9]
- In Range Voting, she applies the same strategy as in Approval Voting to find her "approved" options, then assigns a rating of 100 to approved options and a rating of 0 to the other options.[10]

ability goes to an option which I prefer to the benchmark lottery, that (ii) leads to full consensus if applied by all and if a potential full consensus exists, and that (iii) otherwise leads to partial consensus with a high probability if the utility-by-distance curves are rather concave (as in the LH, QH, LA and QA models) than convex (as in the GH and GA models). See Section 2.3.1 for a more detailed discussion of this and for alternative heuristic formulas for sincere voting under MaxParC.

[8]The rationale is that your vote is most relevant in a nip-and-tuck race, and the most likely nip-and-tuck race is between the favourite poll's leader and runner-up, so that you should vote for your preferred one among those two.

[9]This is called the "leader rule" in [13], see also [14, 15]. Since the most likely nip-and-tuck race is between the approval poll's leader and runner-up, you should approve only your preferred among those two. Since the next likely nip-and-tuck race is between the leader and some other option, you should also approve all options you prefer to the leader.

[10]Following the same rationale.

- In IRV, she denotes the options in descending order of their latest approval polling scores as $x_0, x_1, \dots, x_{k-1}$ and then constructs her ranking as follows: In rank 1 she puts either x_0 or x_1 depending on which she prefers, and labels the other option as y. Then, for each rank $r = 2, 3, \dots, k-1$, she puts either y or x_r in rank r, depending on which she prefers, and labels the other option as the new y.[11]

- In Simple Condorcet, she finds y, z as in Approval Voting, assigns a tied rank of one to her preferred option among y, z and all options she prefers to both, assigns sincere ranks to those other options she prefers to at least one of y, z, and doesn't rank the less preferred option among y, z and all she considers even less desirable.

- In Random Ballot, FC and RFC, she acts like a sincere voter.

- In the Nash Lottery, she first uses the latest favourite polling scores to compute the utility υ of the benchmark lottery with probabilities $f^p(x)/N$. She then computes her rating for any x based on x's apparent chances as estimated by $f^p(x)/N$ and on the difference between $u_i(x)$ and υ as follows. If she is of EU type, she has $\upsilon = \sum_{x \in C} u_1(x) f^p(x)/N$ and uses
$$b(x) = 1 + \frac{f^p(x)(u_i(x) - \upsilon)}{\max_{y \in C} f^p(y)(\upsilon - u_i(y))}, \tag{10}$$
where the denominator is chosen so that the smallest resulting rating is exactly zero.[12] If she is LCP or HCP type, she similarly uses
$$b(x_j) = 1 + \frac{w_j(u_i(x_j) - \upsilon)}{\max_{j'=1}^{k} w_{j'}(\upsilon - u_i(x_{j'}))}, \tag{11}$$
with x_j, w_j as described in the LCP and HCP models above.

- In MaxParC, she applies the same strategy as in Approval Voting to find her "approved" options, then assigns to an approved option x a willingness that is at least as large as her sincere MaxParC willingness for x and large enough to make sure she is counted as approving x should x's approval score be as predicted by $a^p(x)$. More precisely, she puts
$$b(x) = \max\{b^s(x), 101 - 100a^p(x)/N\}, \tag{12}$$

[11] The rationale here is that the most likely nip-and-tuck race is between x_0 and x_1, in which case her vote must go to the better of those from beginning on. Her 2nd ranked option only becomes relevant when the 1st ranked gets eliminated during the tally, in which case the most likely race is between the other (y) and x_2, so she should rank the better of those two 2nd. Her r-th ranked option only becomes relevant when all higher ranked get eliminated, in which case the most likely race is between the one option among $x_0, \dots, x_{r-1}$ not yet ranked, which is the current y, and x_r, so she should rank the better of those two next. See also [16], Fig. 3, and [17].

[12] The rationale is that if some options appear to have higher chances than others, exaggerating one's preferences regarding these options will increase one's influence (see Section 2.3.2 for a formal derivation of this heuristic strategy). In the special case where all $f^p(x)$ are equal, the strategy reduces to voting sincerely.

where $b^s(x) = \max\{0, 100\frac{u_i(x)-u_i(\ell)}{\max_{y\in C} u_i(y)-u_i(\ell)}\}$. To a non-approved option x, she assigns a willingness at most her sincere willingness and small enough so that she is counted as not approving x:

$$b(x) = \min\{b^s(x), 99 - 100a^p(x)/N\}. \tag{13}$$

1.3.2.5 Interactive voting. We assume that in the final voting round voters have to submit a ballot but can then still change their ballots continuously over some time interval during which they can observe the resulting tally statistics in real-time, thus introducing the possibility to interactively test voting strategies and react on others' strategies. We assume there are two additional types of voters who will change their votes during this interval: "trial-and-error" voters and "factional" voters, while all other voters submit an initial ballot as described above and don't change it afterwards. This interactive phase is simulated long enough so that in typical situations a strategic equilibrium can emerge. Before the interactive phase starts, these voters behave like heuristic voters. In the interactive phase, they vote as follows:

1.3.2.6 Trial-and-error voters ($\tau(i) = T$). The interactive phase consists of a large number of consecutive time points, at each of which some percentage of the trial-and-error voters will update their ballots. When a trial-and-error voter i updates her ballot, she picks a random modification out of a set of elementary modifications that depend on the ballot type (see below), submits the modified ballot, observes the resulting change in utility u_i due to all simulateneous modifications, and either sticks to or undoes the modification. She undoes the modification if either u_i has decreased or if u_i has stayed constant and the modification was towards a strictly less "sincere" ballot (see below).

We assume these elementary modifications:

- On a Plurality, FC, or RFC Ballot, either the favourite or the consensus option may be replaced by any other option.
- On an Approval Ballot, one can add or remove approval for a single option.
- On a Range, MaxParC, or RFC Ballot, one can replace the rating or willingness value for a single option by any value in $[0, 100]$.
- A Truncated or Weak Ranking Ballot b may be replaced by another such ballot b' if for some option x, $P_b|_{C\setminus\{x\}} = P_{b'}|_{C\setminus\{x\}}$, i.e., a single option x may be moved to an arbitrary new position in the ranking, making place for it by shifting the other ranks if necessary.

In Random Ballot, trial-and-error voters will always vote sincerely since that is a dominant strategy.

A modified ballot b' is strictly less sincere than b if:

- In Plurality and FC: $u_i(b') < u_i(b)$ (resp. $u_i(b'_1) < u_i(b_1)$ for FC).

- In Approval Voting, IRV, Simple Condorcet, and MaxParC: $e(b') > e(b)$ with $e(b) = |\{(x,y) \in C^2 : u_i(x) > u_i(y) \text{ but } y\, P_b\, x\}|$ (number of wrongly stated binary preferences).
- In Range Voting and the Nash Lottery: $||b' - b^s|| > ||b - b^s||$, where b^s is the sincere ballot described above.
- In RFC: $u_i(b'_1) \leq u_i(b_1)$ and $||b'_3 - b^s|| \geq ||b_3 - b^s||$, but not both equal.

Trial-and-error voters behave as sincere voters during polling and also start the interactive phase with a sincere ballot.

1.3.2.7 Factional voters ($\tau(i) = F$). Since strategic voting can be much more effective when coordinating with other voters having similar preferences, we assume that voters of this type change their ballots as follows during the interactive phase, starting it with a heuristic ballot as described above, and after voting as heuristic voters in the polling rounds, too.

For each $x \in C$, we consider the "faction" F_x of all voters i with $\tau_i = F$ favouring x, $F_x = \{i \in E : \tau_i = F, f_i = x\}$. Each faction F_x is assumed to possess enough information and computing capabilities to calculate a *best unanimous response* to all other voters' current ballots, which is a voting behaviour where all $i \in F_x$ submit the same filled-in ballot and no other unanimous voting behaviour of all $i \in F_x$ would generate a strictly higher total utility $U = \sum_{i \in F_x} u_i$ given that all other voters $j \in E \setminus F_x$ submit the same ballots as before. The assumption that factions cannot coordinate their members to vote differently even if that might be better than all voting the same way can be interpreted as a form of bounded rationality.

During the interactive phase, each faction, whether small or large, has the same constant probability rate for updating their ballots, leading to a Poisson process of updates by randomly picked factions. When a faction F_x updates their ballots, they replace their current ballots by a best unanimous response to all other voters' current ballots as follows:

- In Plurality, they find the plurality scores $p(y)$ resulting from all other voters' ballots, find the set A of options less than $|F_x|$ many votes behind the leader, $A = \{y \in C : p(y) + |F_x| > \max_{z \in C} p(z)\}$, and vote for that $y \in A$ which maximizes U: $b = \arg\max_{y \in A} U(y)$ with $U(y) = \sum_{i \in F_x} u_i(y)$.
- In Approval Voting, they find y in the same way as in Plurality, only using approval scores instead of plurality scores, and then bullet vote for it: $b(y) = 1$, $b(z) = 0$ for all $z \neq y$.
- In Range Voting, they find y in the same way as in Plurality, only using Range Voting scores divided by 100 instead of plurality scores, and then bullet vote for it: $b(y) = 100$, $b(z) = 0$ for all $z \neq y$.
- In IRV, they find the best response truncated ranking ballot by constructing a set A of "candidate" truncated rankings $(x_1, x_2, \ldots, x_\ell)$ that cover all possible results they can effect by submitting identical ballots, and then select the member of A that gives the best

result. A is constructed iteratively by adding ever longer truncated rankings as follows. Given all other voters' ballots, they start by finding the set Y of options $y \in C$ for which y survives the elimination process during the tally strictly longer when they rank y 1st than when they submit a blank ballot. They put $A = \{(y) : y \in Y\}$. Then, for each ranking $(x_1, x_2, \dots, x_\ell) \in A$ with $\ell < k-1$, they find the set Y of options y for which y survives the elimination process during the tally strictly longer when they submit the longer ranking $(x_1, x_2, \dots, x_\ell, y)$ than when they submit the shorter ranking $(x_1, x_2, \dots, x_\ell)$. They add all those ballots $(x_1, x_2, \dots, x_\ell, y)$ to A and iterate until no further ballots are added. One can show that for each possible truncated ranking ballot, there is a member of A that has the same effect when used as the unanimous ballot of all faction members, and $|A| \leq 2^k$.[13]

- In Simple Condorcet, they find the binary opposition values $o(y, z)$ resulting from all other voters' ballots, and put $o(y) = \max_{z \in C} o(y, z)$, $o_0 = \min_{y \in C} o(y)$, and $A = \{y \in C : o(y) < o_0 + |F_x|\}$. For each $y \in A$, they put $A_y = \{z \in A : o(z) < o(y)\}$, and check whether there is a function $g : A_y \to C$ such that $\{(z, g(z)) : z \in A_y\}$ is acyclic and for all $z \in A_y$, $o(z, g(z)) + |F_x| > o(y)$. If this is the case, y can be made the winner by ranking the options in any way that ranks y first and ranks each $z \in A_y$ below its $g(z)$. Among these $y \in A$, they find the one with the largest $U(y)$ and submit any ranking which ranks y first, ranks each $z \in A_y$ below its $g(z)$, and doesn't rank any further options. If no such y exists, they submit a bullet vote for x.

- In Random Ballot, they mark their true favourite since that is a dominant strategy: $b = \arg\max_{x \in C} u_i(x)$.

- In FC, they mark their true favourite as "favourite" and find an optimal option for marking as "consensus" by computing the resulting U for all of the k many possible choices.

- In RFC, they submit sincere ratings and find an optimal combination of options for marking as "favourite" and "consensus" by computing the resulting U for all of the k^2 many possible combinations.

- In the Nash Lottery, they try to find a (globally) best response by starting with a common ballot derived by averaging the faction members' sincere Nash Lottery ballots (see above) and then following a simple steepest ascent optimization algorithm until reaching a (local) optimum of $U(x)$. Although this local optimum might not be a globally best response, we assume they use the resulting ballot anyway, which can be considered an additional form of bounded rationality.

- In MaxParC, they compare the results of all the 2^k many ballots b_A of the form $b_A(x) = 100$ if $x \in A$ and $b_A(x) = 0$ else, for some subset of options $A \subseteq C$. They identify that A which maximizes U given all others' ballots. Note that the corresponding b_A is a unanimous

[13]See Section 2.3.3 for a proof sketch.

best response since only the resulting approvals matter. For this A, they calculate the approval scores a_y that would result from using ballot b_A in MaxParC given all others' ballots. Then they define

$$w(y) = \sum_{i \in F_x} \max\{0, 100 \frac{u_i(y) - u_i(\ell)}{\max_{z \in C} u_i(z) - u_i(\ell)}\}/|F_x|, \tag{14}$$

which is the average sincere ballot of the faction members, and use the ballot with

$$b(y) = \max\{100(1 - a_y/N), w(y)\} \tag{15}$$

for $y \in A$ and

$$b(y) = \min\{99(1 - a_y/N), w(y)\} \tag{16}$$

for $y \notin A$, which leads to the same approvals as b_A and is thus also a best unanimous response. In other words, they use that best unanimous response which is closest to the average sincere ballot.

- In polling rounds, they act as in Plurality and Approval Voting.

Note that if a steady state emerges, it approximates a pure-strategy Nash equilibrium between the trial-and-error voters as individual players and the factions as aggregate players, which will in general however not be a strong or coalition-proof equilibrium since although we regard factions, we do not regard inter-factional coalitional strategies. Also, the process may also lead to cyclic or more complex attractors rather than a steady state.

1.3.3 Experiment design

We generated $M = 1,293,906$ many independent group decision problems, drawing their parameters independently from the following probability distributions (where parameter names in code are set in `this font`):

Number of voters N. We drew odd numbers between 9 and 999 such that $\log_{10} N$ was approximately uniformly distributed in the interval $[1, 3]$.

Number of options k. Uniformly in $\{3, \dots, 9\}$.

Preference models. Uniformly in {Unif, BM, GA, QA, LA}.

BM parameters. For the block model: number of voter blocks `Bmr` $= r \sim \text{Unif}\{2, 5, 9\}$, block size heterogeneity `Bmh` $= h \sim \text{Unif}\{0, 1\}$, individuality `Bmiota` $= \iota \sim \text{Unif}\{0.1, 0.5\}$.

Spatial model parameters. For GA, QA, and LA: policy space dimension `dim` $= d \sim \text{Unif}\{1, 2, 3\}$, voter heterogeneity `omega` $= \omega \sim \text{Unif}\{1, 2, 3, 5\}$, option broadness heterogeneity `rho` $= \rho \sim \text{Unif}\{0, 1/3, 2/3, 1\}$, where $\rho = 0$ corresponds to the homogenous cases GH, QH, LH.

Risk attitude scenarios. Uniformly in {all-EUT, all-LCP, all-HCP, mixed}, where in 'mixed' 20% of the voters are EUT, 40% LCP, and 40% HCP.

Number of polling rounds R. Uniformly in $\{1, 2, 3, 5, 7, 10\}$.

Behavioural type scenarios. Uniformly in {lazy, middle, strat, all-L, all-S, all-T, all-H, all-F}

The following parameters were not varied:

Length of interactive phase. 100 time points.

Trial-and-error frequency. At each time point, 50% of the trial-and-error voters updated their ballots.

Factional update probability. At each time point, each faction had a 10% probability to update their ballots.

For each group decision problem, we constructed a second problem in which a randomly chosen option was replaced by a compromise option y that was constructed from set C_0 of the remaining $k - 1$ options, to analyse the effect that a specifically designed compromise option would have. Depending on the preference model, voters' preferences about y were constructed as follows: In Unif and BM, the compromise got the average utility of the other options, $u_i(y) = \sum_{x\in C_0} u_i(x)/(k-1)$; in GH, QH, and LH, the compromise's position ξ_y was chosen to be a weighted average of the other options' positions ξ_x, with weights w_x proportional to first-preference support and inversely proportional to options' platforms' broadness σ_x:

$$\xi_y = \sum_{x\in C_0} w_x \xi_x / \sum_{x\in C_0} w_x, \tag{17}$$

$$w_x = |\{i \in E : x = \arg\max_{z\in C} u_i(z)\}|/\sigma_x. \tag{18}$$

For each of these $2M$ decision problems, we simulated R rounds of polling. Finally, for each of the ten voting methods independently, we simulated an initial voting round and an interactive voting phase[14] based on the same polling results, and determined all options' resulting winning probabilities ℓ both after the initial voting round and after the interactive phase.

[14]For the basically deterministic IRV, we did not calculate tie probabilities since this would have been too costly due to the iterative nature of the method; instead, we resolved ties in IRV randomly, so that the resulting lottery always appeared to be a sure-thing lottery instead of the true tying lottery. For the NL method, the optimization problem $\max S(\ell)$ was solved using Sequential Least Squares Programming (SLSQP); to avoid a convergence failure due to singular Jacobian matrices because of zero ratings, we added 10^{-5} to all ratings (the maximal rating always being 100). In the interactive phase of NL, a faction's best response ratings optimization problem was solved using Constrained Optimization By Linear Approximation (COBYLA) since that converged better than SLSQP.

1.3.4 Social welfare metrics

To measure the welfare effects of the tested voting methods, we use a set of metrics which are based on three different social welfare measures (utilitarian, Gini–Sen, and egalitarian welfare), taken either on an absulute or a relative scale, and either taken before or after the interactive phase of the simulations, giving a total of twelfe different metrics per problem and method.

All these measures are aggregating the voters' individual utility $u_i(\ell)$ they get from the resulting lottery ℓ, as modelled by the various utility models discussed above. In applications where there is only a single decision taken, these measures must hence be interpreted as measuring the 'ex ante' efficiency of the method, as opposed to the 'ex post' efficiency that would be based on the utilities of the actual options chosen by the resulting lottery. In applications where we imagine a sequence of decisions, our efficiency metrics can be interpreted as measuring the long-run efficiency of the method over the whole sequence of decisions.

1.3.4.1 Utilitarian welfare The simplest and most popular measure is the one proposed by average utilitarianism, $W_{\text{util.}}(\ell) = \sum_{i\in E} u_i(\ell)/|E|$.

Since in all our utility models, lottery utility $u_i(\ell)$ is a linear combination of option utilities $u_i(x)$, the lottery that maximizes $W_{\text{util.}}(\ell)$ is a sure-thing lottery.

1.3.4.2 Gini–Sen welfare As $W_{\text{util.}}(\ell)$ is insensitive to redistribution of utility across voters, and hence to inequality between voters' utilities, we also use two inequality-averse metrics, the first of which is the Gini–Sen welfare function $W_{\text{Gini}}(\ell) = \sum_{i\in E} \sum_{j\in E\setminus\{i\}} \min\{u_i(\ell), u_j(\ell)\}/|E|(|E|-1)$.

As a motivating story one can imagine voters meet in a large sequence of bilateral meetings and each time evaluate the welfare status of society welfare in terms of the smaller of their two utilities, and overall social welfare is then measured by the average of all these individual pairwise evaluations. Another motivation for the same metric is that it can be seen as an "inequality-adjusted" version of $W_{\text{util.}}$, since we have $W_{\text{Gini}}(\ell) = W_{\text{util.}}(\ell)(1 - I_{\text{Gini}}(\ell))$, where $I_{\text{Gini}}(\ell)$ is the well-known Gini coefficient of inequality in utilities $u_i(\ell)$ [18].

Note that the lottery ℓ^* which maximizes $W_{\text{Gini}}(\ell)$ can be expected to be a proper lottery rather than a sure-thing lottery. This is because randomization tends to reduce inequality more than it reduces average utility.

1.3.4.3 Egalitarian welfare As the most extremely inequality-averse welfare metric, we also consider the egalitarian one, $W_{\text{egal.}}(\ell) = \min_{i\in E} u_i(\ell)$. As in the case of Gini–Sen welfare, maximization of ex-ante egalitarian welfare usually requires randomization.

1.3.4.4 Absolute and relative welfare metrics All three welfare metrics measure welfare on the same scale as individual utility, hence are hard to compare directly across different utility models since these use quite different scales. Also, for some models their distribution is quite skewed, having a long lower tail. In addition to the above *absolute* welfare metrics, we therefore also compare the *relative* metrics

$$relW_{\text{util./Gini/egal.}}(\ell) = \frac{W_{\text{util./Gini/egal.}}(\ell) - \min_{x\in C} W_{\text{util./Gini/egal.}}(x)}{\max_{x\in C} W_{\text{util./Gini/egal.}}(x) - \min_{x\in C} W_{\text{util./Gini/egal.}}(x)} \in [0, \infty] \tag{19}$$

which rescale the welfare so that the sure-thing (!) lotteries giving the lowest and highest welfare get scores 0 and 1, respectively. This design allows us to interpret values larger than 1 as welfare gains from randomization.

Still, as it turned out, the relative versions of Gini–Sen and egalitarian welfare often take very large values for nondeterministic methods and thus now have a very skewed distribution with a long upper tail. For this reason, we also study an *alternative relative* version of all three metrics, defined as

$$altrelW_{\text{util./Gini/egal.}}(\ell) = 2\frac{W_{\text{util./Gini/egal.}}(\ell) - \min_{x\in C} W_{\text{util./Gini/egal.}}(x)}{W_{\text{util./Gini/egal.}}(\ell) + \max_{x\in C} W_{\text{util./Gini/egal.}}(x) - 2\min_{x\in C} W_{\text{util./Gini/egal.}}(x)} \tag{20}$$

$$= \frac{2relW_{\text{util./Gini/egal.}}(\ell)}{1 + relW_{\text{util./Gini/egal.}}(\ell)} \in [0, 2]. \tag{21}$$

These are now restricted to the interval $[0, 2]$, again taking a value of 0 and 1 for the sure-thing (!) lotteries giving the lowest and highest welfare.

1.3.4.5 "Cost of fairness" As an alternative to the above relative welfare metrics, one can also compare welfare differences between methods with utility differences within the electorate to assess the influence of method choice on welfare. In analogy to the notion of a "price of anarchy" [19], we therefore define a "relative cost of fairness",

$$CF = \frac{W^{RV}_{\text{util}} - W^{MPC}_{\text{util}}}{W^{RV}_{\text{util}} - W^{RV}_{\text{egal}}}, \tag{22}$$

where the numerator is the absolute difference in average voter utility between the best deterministic method RV's result and the best proportional method MaxParC's result (which could be termed the "absolute cost of fairness"), and the denominator is the difference between the average and minimum voter utility under RV (which could be termed the "absolute egalitarian inequality").

1.3.5 Randomization metrics

To measure the degree of *randomization* a voting method actually applies, we computed two established entropy measures, Shannon entropy and Rényi entropy of degree two, and the maximal probability $\max_{x\in C} \ell_x$, again applied to the results before and after the interactive phase. This gives six randomization metrics in total per problem and method.

1.3.6 Voter satisfaction metrics

As another type of performance indicators, we computed each voter's "satisfaction level"

$$\frac{u_i(\ell) - \min_{x \in C} u_i(x)}{\max_{x \in C} u_i(x) - \min_{x \in C} u_i(x)} \in [0, 1], \tag{23}$$

which would be zero if i's least preferred option won for sure, and unity if i's favourite won for sure. Based on these, we report average satisfaction levels in the whole electorate and, to assess possible advantages of strategic behaviour, by behavioural type.

1.3.7 Consequentialist preferences over methods

Finally, to get an idea of which methods voters would chose if that choice was itself performed by majority voting, we counted for each decision problem how many voters would prefer the lottery resulting from some method A to that resulting from some method B.

2 Detailed Results

2.1 Properties of voting methods

2.1.1 Basic consistency properties

2.1.1.1 Anonymity. A voting method is *anonymous* iff it treats all voters alike, i.e., iff its result is invariant under permutations of voters. All considered methods have this property.

2.1.1.2 Neutrality. A voting method is *neutral* iff it treats all options alike, i.e., iff the resulting winning probabilities of any two options x, y are swapped when x, y are swapped on all ballots. All considered methods have this property.

2.1.1.3 Pareto-efficiency w.r.t. stated preferences. An option y is *Pareto-dominated* w.r.t. stated preferences iff there is another option x with $x \, P_{\beta_i} \, y$ for all $i \in E$. A voting method is *Pareto-efficient* w.r.t. stated preferences iff all Pareto-dominated options get zero winning probability. All considered methods except FC and RFC fulfill this. Since in FC, only the fall-back option is interpreted as stating a preference, a Pareto-dominated y might still be everyone's proposed consensus and win. Similarly, since in RFC only the ratings are interpreted as preferences, y still might be named by someone as fall-back option and thus have positive winning probability.

It is more difficult to check whether also an option which is Pareto-dominated w.r.t. *true* preferences will have zero winning probability, since this depends on whether and how voters behave strategically. Our numerical simulations at least suggests that all of the considered methods, including FC and RFC, fulfill this criterion under normal circumstances.

2.1.2 Monotonicity properties

Although there are a number of variants of the 'monotonicity' criterion, we here focus on two variants of Woodall's 'mono-raise' monotonicity [20], which differ only really for non-deterministic methods, and one properly weaker property related to Woodall's 'mono-add-plump' monotonicity.

2.1.2.1 Strong mono-raise monotonicity. A voting method is *strongly mono-raise monotonic* iff the winning probability of an option y cannot increase if a different option x is advanced on one ballot: $M(\beta)_y \leq M(\beta')_y$ whenever $x \neq y$, $\beta_i \, Q^C_x \, \beta'_i$ for some $i \in E$, and $\beta_j = \beta'_j$ for all $j \in E \setminus \{i\}$.

2.1.2.2 Weak mono-raise monotonicity. A voting method is *weakly mono-raise monotonic* iff the winning probability of an option x cannot decrease if x is advanced on one ballot: $M(\beta)_x \geq M(\beta')_x$ whenever $\beta_i \, Q^C_x \, \beta'_i$ for some $i \in E$ and $\beta_j = \beta'_j$ for all $j \in E \setminus \{i\}$.

2.1.2.3 Weak mono-raise-abstention monotonicity. We call a voting method *weakly mono-raise-abstention monotonic* iff the winning probability of an option x cannot decrease if x is advanced on an abstention ballot: $M(\beta)_x \geq M(\beta')_x$ whenever $\beta_i \, Q^C_x \, \beta'_i$ for some $i \in E$, $\beta_j = \beta'_j$ for all $j \in E \setminus \{i\}$, and β'_i is an abstention ballot.

Obviously, strong mono-raise monotonicity implies weak mono-raise monotonicity, which in turn implies weak mono-raise-abstention monotonicity.

For PV, AV, RV, SC, and RB it is straightforward to prove all three forms of mono-raise monotonicity (exercise left to the reader). IRV is known to violate both strong and weak mono-raise monotonicity [20] but is easily seen to fulfill weak mono-raise-abstention monotonicity.

FC fulfills weak but not strong mono-raise monotonicity since if z is everyone's proposed consensus, advancing x on some consensus ballot destroys the consensus so that someone's fall-back option $y \neq x$ can get gets positive winning probability.

RFC violates both weak and strong mono-raise monotonicity. Consider the case of three options x, y, z and two voters who both name z as consensus and rate (x, y, z) at $(0, 3, 1)$. If one names x and the other y as fall-back, x and y both get winning probability $1/2$, but if both name x as fall-back, z wins for sure.

NL violates strong mono-raise monotonicity. Again consider three options x, y, z and two voters $1, 2$ who rate them as $\beta_1(x, y, z) = (1/2, 0, 1/6)$ and $\beta_2(x, y, z) = (0, 1, 3/4)$. Then $M(\beta)_y = 0$. But if we increase $\beta_1(x)$ to 1, we get $M(\beta)_y = 1/2 > 0$. Numerical simulations suggest that NL fulfills weak mono-raise monotonicity, which we conjecture but were not able to prove yet unfortunately.

Regarding weak mono-raise-abstention monotonicity, it was shown in [2] that NL (there called "Nash Max Product") fulfills a roughly equivalent condition of "Strict Participation" when all ballots are "dichotomous" in the sense that all ratings are either zero or 100 (or some other common, fixed, positive number). Using the Envelope Theorem, we can give an alternative proof of weak mono-raise-abstention monotonicity for arbitrary ratings. *Proof.* Let E, C be fixed, consider some $i \in E$ and $x \in C$, assume $\beta : E \to B(C)$ is fixed except for its entry $\beta_i(x)$, and assume $\beta_i(y) = \epsilon$ for all $y \in C \setminus \{x\}$. We will study the change of the NL probabilities $p^*(\alpha) = M(\beta)$ as a function of the parameter $\alpha = \beta_i(x)$ and show that $\alpha' > \alpha$ implies $p^*(\alpha') \geq p^*(\alpha)$, which will suffice to prove the claim. $p^*(\alpha)$ is the solution of the maximization of the continuously differentiable function $f(p, \alpha) = \sum_{j \in E} \log h_j(p, \alpha)$ with $h_j(p, \alpha) = \sum_{y \in C} \beta_j(y) p_y$ and $\beta_i(x) = \alpha$ under the constraint $g(p, \alpha) = \sum_{y \in C} p_y = 1$. Let $V(\alpha) = \max_{p, g(p,\alpha)=1} f(p, \alpha)$ be the corresponding maximum. Since the constraint is independent of the parameter α, the envelope theorem implies that

$$V'(\alpha) = \left. \frac{\partial f(p, \alpha)}{\partial \alpha} \right|_{p=p^*(\alpha)} = \frac{p^*(\alpha)_x}{h_i(p^*(\alpha), \alpha)}. \tag{24}$$

Now assume an infinitesimal increase in α from $\alpha = \alpha_0$ to $\alpha = \alpha_1 = \alpha_0 + d\alpha$ with $d\alpha > 0$. Then the above implies

$$V(\alpha_1) = V(\alpha_0) + V'(\alpha_0) d\alpha = f(p^*(\alpha_0), \alpha_0) + \frac{p^*(\alpha_0)_x}{h_i(p^*(\alpha_0), \alpha_0)} d\alpha, \tag{25}$$

but also

$$V(\alpha_1) = f(p^*(\alpha_1), \alpha_1) \tag{26}$$

$$= f(p^*(\alpha_1), \alpha_0) + \left. \frac{\partial f(p, \alpha)}{\partial \alpha} \right|_{p=p^*(\alpha_1), \alpha=\alpha_0} d\alpha \tag{27}$$

$$= f(p^*(\alpha_1), \alpha_0) + \frac{p^*(\alpha_1)_x}{h_i(p^*(\alpha_1), \alpha_0)} d\alpha. \tag{28}$$

Since optimization means that $f(p^*(\alpha_0), \alpha_0) \geq f(p^*(\alpha_1), \alpha_0)$, this implies

$$0 \leq \frac{f(p^*(\alpha_0), \alpha_0) - f(p^*(\alpha_1), \alpha_0)}{d\alpha} \tag{29}$$

$$= \frac{p^*(\alpha_1)_x}{h_i(p^*(\alpha_1), \alpha_0)} - \frac{p^*(\alpha_0)_x}{h_i(p^*(\alpha_0), \alpha_0)} \tag{30}$$

$$= \frac{p^*(\alpha_1)_x}{(\alpha_0 - \epsilon) p^*(\alpha_1)_x + \epsilon} - \frac{p^*(\alpha_0)_x}{(\alpha_0 - \epsilon) p^*(\alpha_0)_x + \epsilon}, \tag{31}$$

Since $\epsilon > 0$, this implies $p^*(\alpha_1)_x \geq p^*(\alpha_0)_x$. Since this holds for all values α_0 of α, we have shown that $p^*(\alpha)_x$ is a weakly increasing function of α as claimed. *Q.E.D.*

MaxParC fulfills all three forms. *Proof.* It suffices to show that if (i) x is advanced by one voter i from $\beta_i(x) = r$ to $\beta_i(x) = r' > r$, (ii) some voter j is drawn at random, and (iii) some option

$y \neq x$ is in the set A'_j after the change, then y must have been in A'_j before the change and A'_j can only have grown due to the change. It is easy to see that $A(x)$ can only have grown and that no other $A(z)$ has changed, hence $a'(x)$ has properly grown but no other $a'(z)$ has changed. So if y is in $A'_j = \arg\max_{z \in A_j} a'(z)$, the value of $\max_{z \in A_j} a'(z)$ has not changed, hence y must have been in this set before, and the only change in A'_j can be that now x is also in A'_j. This means A'_j can only have grown and thus y's winning probability decreased. *Q.E.D.*

2.1.3 Further consistency properties

2.1.3.1 Independence from Pareto-dominated alternatives. The idea of this criterion is that the "removal" of a Pareto-dominated option y from all ballots should have no effect on the winning probabilities. It was first introduced by Steve Eppley on the election-methods emailing list.[15] Since for some methods it is not obvious how a filled ballot will change when an option is removed (e.g., what should one assume about how a filled Plurality Ballot will change if the marked option is removed?), we do not study this in a formal way here but rather discuss it verbally.

For methods using a ballot type that lets voters rate or rank all options independently (AV, RV, SC, NL and MaxParC), let us assume "removal" means leaving the other options' ratings unchanged. For IRV, let us assume "removal" implies decreasing the ranks of the later-ranked options by one. Then those six methods all fulfill this criterion, and so do Plurality and Random Ballot whenever no voter has named y as favourite (which rational voters wouldn't).

As this criterion implies Pareto-efficiency, FC and RFC do not fulfill it.

We note that in particular many Condorcet-type methods, which elect a winner of all pairwise comparisons for sure if such an option exists, including the 'Ranked Pairs' method by Nicolaus Tideman [21] and the 'Beatpath' method by Markus Schulze [22], fail this criterion.

2.1.3.2 Independence from losing options. This criterion demands that the removal of any option y receiving zero winning probability must have no effect on the winning probabilities. This is a variation of the famous 'Independence of Irrelevant Alternatives' criterion, and is stronger than Independence from Pareto-dominated alternatives if Pareto-efficiency is given.

It is easy to see that again AV, RV, RB, and NL fulfill this and FC and RFC do not. For PV, some voters may have voted for y and now vote for the current runner-up and make it win.

IRV and SC also do not, as can be seen from the example of three options x, y, z and three factions $F_{1,2,3}$ of sizes $4, 3, 2$ and rankings $F_1 : x > y > z$, $F_2 : y > z > x$, $F_3 : z > x > y$. Both methods elect x for sure but elect z if y is removed.

MaxParC also fulfills this criterion. *Proof.* Removal of y does not change who approves which other options. If y has zero winning probability, every voter who approves y also approves some higher-scoring option. Hence for no voter the set of highest-scoring approved options changes. Thus all other options' winning probabilities are unaffected. *Q.E.D.*

[15] http://lists.electorama.com/pipermail/election-methods-electorama.com//2003-March/107700.html

An even stronger variant of 'Independence of Irrelevant Alternatives' that can also be interpreted as a form of 'monotonicity' goes as follows: removing any option y from E must not decrease any remaining option x's winning probability. NL probably violates this while MaxParC clearly fulfills it.

2.1.3.3 Independence from cloned options ("clone-proofness"). Another type of criterion deals with the addition of an option y, called a 'clone', that is very "similar" to some existing option x. Since "similarity" can be defined in different ways depending on the ballot type, we restrict our interest here to the special case where y is a unique 'exact clone' of x, meaning all voters are truly indifferent between x and y but not between these two and any further option z. We demand that in that case and under plausible assumptions on voter's voting behaviour, the addition of y shall not change the winning probability of any other option $z \notin \{x, y\}$.

Let us assume that after the addition of y, voters will assign y the exact same approval, rating, or ranking (if tied rankings are allowed, otherwise an adjacent ranking) as x, and will name any option $z \notin \{x, y\}$ as favourite or proposed consensus iff they named the same option before the addition, only possibly switching from naming x to now naming y. It is then easy to see that AV, RV, IRV, SC, RB, RFC, NL and MaxParC all fulfill this form of 'exact clone independence', while PV and FC do not.

Note that there are other, stronger, forms of clone-independence, including the one discussed in [21], that some variants of IRV, many Condorcet-type methods, NL, and MaxParC might not fulfill.

2.1.3.4 Revelation of preferences. Some voting methods have the property that, sometimes depending on the level of strategic behaviour, voters' filled-in ballots reveal all or part of their preferences. Under RB, for example, whenever a voter has an option she strictly prefers to all other options (a unique favourite), it is a weakly dominant strategy to specify that option. This form of "strategy-proofness" can be interpreted as implying that RB "reveals unique favourites" (but nothing else about voters' preferences).

As another example, it was shown in [1] that under RFC, whenever a voter has preferences conforming to expected utility theory (see below) with some utility function u, she has no strategic incentive to specify different ratings than a properly rescaled version of u, and hence RFC can be said to "reveal von-Neumann–Morgenstern utility functions" (but no preferences that do not conform to expected utility theory). Still, RFC is not strategy-proof in the sense that there always exist weakly dominant strategies, since in its other two ballot components, a rational voter may want to name a proposed consensus option that depends on others' preferences, and may have incentives to name a different option as "fall-back" than her favourite. This shows that full preference revelation is related to but neither implied by nor stronger than strategy-proofness.

FC reveals favourites but its consensus ballot component is strategic. NL and MaxParC also reveal favourites in the sense that a voter has no incentive to not rank her favourite first or to rate it below 100, but usually has an incentive to rate all other options strictly below 100. Neither of them however reveal much more of a voter's preferences.

AV and RV don't reveal favourites since typically a rational voter has an incentive to approve (or rate at 100) some additional options. Still, in the case where voters have no information about others' preferences, AV and RV can be said to reveal something about a voter's preferences, because in that case a rational expected utility theory voter would approve (or rate at 100) all options she prefers to drawing an option uniformly at random (and would rate all other options at 0), so that one can infer that she strictly prefers each approved to each disapproved option.

Similarly, under NL, expected utility theory voters who use the zero-information heuristic derived in the end of 2.3.2 also reveal their full preferences, but this heuristic might not be a weakly dominant strategy under zero information, so rational voters might not use it. Under MaxParC, the linear heuristic derived in 2.3.1 reveals the above-average part of a voter's utility function but is also typically not weakly dominant under zero information.

IRV and SC also do not reveal favourites. For IRV, consider three options and six voters and assume voter 1 has preferences $1 : A > B > C$ and the others vote $A > B$, A, $B > C$, $C > A$, C. Then if 1 votes sincerely, B is removed and a coin toss between A and C results. But if 1 votes $B > A > C$, A still gets probability 1/2 but now B gets 1/3 and C gets 1/6, which 1 strictly prefers. For SC, consider three options and three voters and assume voter 1 has preferences $1 : A > B > C$ and the others vote $B = C > A$, $C > A > B$. Then 1 would want to vote $A = B > C$ or $B > A > C$ to ensure a coin toss between B and C rather than voting sincerely $A > B > C$ and getting C for sure.

2.1.4 Proportional allocation of effective power

This criterion requires that in every situation (C, E) and for every option $x \in C$ and group of voters $G \subseteq E$, there must be a way of voting $\beta_G \in B(C)^G$ for G so that for all ways of voting $\beta_{-G} \in B(C)^{E\setminus G}$ of the other voters, the winning probability of x is at least as large as G's relative size: $M(\beta_G, \beta_{-G})_x \geq |G|/|E|$. A related criterion was discussed for the special case of 'dichotomous preferences' under the name 'Core Fair Share' in [2].

Since all considered methods are neutral and anonymous, one can summarize the power distribution by drawing the maximal winning probability $h(s)$ a group of size $|G|/|E| = s$ can guarantee any option x of their choice under the various methods, as is done in Fig. 1 of the main text. For all considered deterministic methods, this "effective decision power" is basically a step function with the value zero for $s < 1/2$ and one for $s > 1/2$ (blue line). Only for $s = 1/2$ the value depends on the method's detailed treatment of ties, which we do not discuss here. Note that $h(s) + h(1 - s) \leq 1$, so the integral of $h(s)$ from 0 to 1 is at most 1/2 but can be properly smaller, e.g. for the Borda Score method, where h is essentially a step function switching from 0 to 1 at $s = 2/3$. By contrast, for all considered non-deterministic methods, it can easily be seen that any group G can guarantee x a probability at least $|G|/|E|$ by simply bullet-voting for x, so effective decision power is simply equal to s (green line), which we call *proportional allocation of effective power.*

Note that of course there are also non-deterministic neutral anonymous methods with different allocations of effective power. E.g., one could draw a sequence of plurality ballots at random until one option was named twice ("first to get two"), giving a smooth but S-shaped nonlinear power curve $h(s) = s^2(3 - 2s)$.

For NL, it is the specific use of the logarithm that gives a linear power curve. Indeed, consider a method that puts $M(\beta) = \arg\max_\ell S(\ell)$ for $S(\ell) = \sum_{i\in E} f(r_i(\ell))$, some weakly increasing and continuously differentiable function f, and $r_i(\ell) = \sum_{x\in C} \ell_x \beta_i(x)$. If the power curve is linear, then whenever a group of voters $G \notin \{\emptyset, E\}$ bullet-votes for x and the other voters $E \setminus G$ bullet-vote for y, we must have $p := M(\beta)_x = |G|/|E| =: s$ and $q := M(\beta)_y = 1 - |G|/|E| = 1 - s$, hence the first-order condition

$$0 = (\partial_p - \partial_q)S\,|_{p=s,q=1-s} = sf'(100s) + (1 - s)f'(100(1 - s)) \tag{32}$$

implies $f'(100s) \propto 1/s$ for all rational numbers $s \in (0, 1)$ and thus $f(r) \propto \log r$ for all real numbers $r \in (0, 100)$.

2.1.5 Consensus supporting properties

In this section, we show that our two focus methods NL and MPC support both full and partial consensus even with strategic voters. To do so, we show that the respective potential consensus options result both from sincere voting (see 2.3.1 for a discussion of sincere voting in MaxParC) and in several forms of strategic equilibrium in archetypial decision situations.

2.1.5.1 Nash Lottery supports full consensus

Assumptions. We assume two equal-sized factions F_1, F_2 of m many voters each, and three options A, B, D. Voters in F_1 have von-Neumann–Morgenstern utility function $u_1(A, B, D) = (1, 0, u)$ and submit ratings $r_1(A, B, D) = (1, 0, r)$, those in F_2 have $u_2(A, B, D) = (0, 1, v)$ and submit $r_2(A, B, D) = (0, 1, s)$ with $r, s \in (0, 1)$ and $1/2 < u, v < 1$, so that both factions prefer D to a coin toss between A and B.

Resulting lottery and expected utilities. If A, B, D get probabilities $p, q, 1 - p - q$, the resulting expected utilities are

$$U_1 = p + (1 - p - q)u, \tag{33}$$
$$U_2 = q + (1 - p - q)v, \tag{34}$$

and the Nash sum is

$$f = m \log(p + (1 - p - q)r) + m \log(q + (1 - p - q)s).$$

Because f is concave in both p and q, the unique pair p, q maximizing f can be found as follows. Given $q \in [0, 1]$, f is maximized by that $p \in [0, 1 - q]$ which is closest to the point

$p_0(q)$ of zero slope,

$$0 = \partial_p f = \frac{1-r}{p+(1-p-q)r} + \frac{-s}{q+(1-p-q)s}, \tag{35}$$

$$p_0(q) = \frac{(1-q)(1-2r)s+(1-r)q}{2(1-r)s}. \tag{36}$$

Similarly, given p, f is maximized by that $q \in [0, 1-p]$ closest to

$$q_0(p) = \frac{(1-p)(1-2s)r+(1-s)p}{2(1-s)r}. \tag{37}$$

If we introduce the notation $[x]_0^y = \max(0, \min(x, y))$, the maximum of f is thus attained where

$$p = \left[\frac{(1-q)(1-2r)s+(1-r)q}{2(1-r)s}\right]_0^{1-q}, \tag{38}$$

$$q = \left[\frac{(1-p)(1-2s)r+(1-s)p}{2(1-s)r}\right]_0^{1-p}. \tag{39}$$

Depending on r, s, the solution (p, q) found by the Nash Lottery method and resulting utilities (U_1, U_2) are the following:

$$(p, q, U_1, U_2) = \begin{cases} (\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \frac{1}{2}) & r+s<1, \\ \frac{(1-2r, 0, 1-2r+u, v)}{2-2r} & s \le \frac{1}{2}, r \ge 1-s, \\ \frac{(0, 1-2s, u, 1-2s+v)}{2-2s} & r \le \frac{1}{2}, s \ge 1-r, \\ (0, 0, u, v) & s, r \ge \frac{1}{2}. \end{cases} \tag{40}$$

Outcome with sincere voters. Sincere voters put $r = u > 1/2$ and $s = v < 1/2$ and thus get $p = q = 0$, i.e., the consensus option D wins for sure.

Strategic equilibria between factions. To analyse strategic incentives for the two factions, we treat F_1, F_2 as the players of a two-player game in which they simultaneously choose $r, s \in (0, 1)$, and study the Nash equilibria (NE) of that game. Given some s, F_1's best responses are the following: If $s < \min(\frac{1}{2}, 1-u)$, each $r < 1-s$ is a best response. If $1-u \le s < \frac{1}{2}$, each $r \ge 1-s$ is a best response. If $s > \frac{1}{2}$, only $r = 1-s$ is a best response. If $s = \frac{1}{2} \ge u$, each $r \le \frac{1}{2}$ is a best response. Finally, if $s = \frac{1}{2} < u$, only $r = \frac{1}{2}$ is a best response.

Since we assume $u, v > \frac{1}{2}$, this results in the following sets of NE. Any combination $r, s < \frac{1}{2}$ is a NE giving only $U_1 = U_2 = \frac{1}{2}$. Any combination $(r, 1-r)$ with $r \in (1-v, \frac{1}{2})$ is a NE giving $(U_1, U_2) = \frac{(1+u-2r, v)}{2-2r}$. Any combination $(1-s, s)$ with $s \in (1-u, \frac{1}{2})$ is a NE giving $(U_1, U_2) = \frac{(u, 1+v-2s)}{2-2s}$. Finally, $s = r = \frac{1}{2}$ is the "focal" NE giving $(U_1, U_2) = (u, v)$ and the largest utility sum $U_1 + U_2$ of all NE.

Summary. The above analysis shows that in this scenario, Nash Lottery supports full consensus with both sincere voters (who would put $r = u$ and $s = v$) and strategic voters (who would put $s = r = \frac{1}{2}$). We conjecture that similar calculations will show that the same holds with more and unequally sized factions.

2.1.5.2 Nash Lottery supports partial consensus Assume that to the above we now add a third faction F_3 of size $n - 2m$ and a fourth option C, and utilities $u_3(A, B, D, C) = (0, 0, 0, 1)$, $u_1(C) = u_2(C) = 0$.

Strategic equilibria between factions. F_3 has a dominant strategy to bullet-vote for D, i.e., put $r_3(A, B, D, C) = (0, 0, 0, 1)$, and F_1, F_2 have no reason not to put $r_1(C) = r_2(C) = 0$. If we parameterize the probabilities of (A, B, D, C) as $(p(1-w), q(1-w), (1-p-q)(1-w), w)$, the Nash sum becomes

$$f = m\log(p + (1-p-q)r) + m\log(q + (1-p-q)s) + 2m\log(1-w) + (n-2m)\log(w),$$

which is maximized by $w = 1 - 2m/n$ and the same values of (p, q) as above. Since F_1, F_2's utilities are proportional to the case above,

$$U_1 = (1-w)(p + (1-p-q)u), \tag{41}$$
$$U_2 = (1-w)(q + (1-p-q)v), \tag{42}$$

the strategic analysis is the same as before, so putting $r = s = \frac{1}{2}$ is again the utility-maximizing and focal equilibrium.

Outcome with sincere voters. Sincere voters in F_1 and F_2 still put $r = u > 1/2$, $s = v < 1/2$, $\beta_i(C) = 0$, and those in F_3 bullet-vote for C. They still get $p = q = 0$ and in addition $w = 1 - 2m/n$, i.e., the partial consensus option D now gets probability $1 - w = 2m/n = (|F_1| + |F_2|)/|E|$, as required.

This shows that in this scenario, Nash Lottery also supports partial consensus. We conjecture that the same holds with more and unequally sized factions, and with several partial consensusses between different sets of factions.

2.1.5.3 MaxParC supports full consensus Note that under MaxParC, one has never an incentive to approve a worst-liked option or to disapprove one's favourite. Since the following is not restricted to voters with expected utility preferences, we don't use von-Neumann–Morgenstern utility functions u_i here but rather state a voter's preferences over lotteries of options ℓ, ℓ' by means of the binary relations $\ell\, P_i\, \ell'$ (strict preference for ℓ over ℓ'), $\ell\, R_i\, \ell'$ (weak preference), and $\ell\, E_i\, \ell'$ (indifference), only assuming that R_i is a quasi-ordering (not necessarily complete) and that P_i and E_i are its antisymmetric and symmetric parts.

Assumptions. Assume $m \geq 2$ factions F_j with sizes $N_j \geq 1$, $N = \sum_{j=1}^m N_j$, with distinct true favourites x_j, and assume voters are indifferent about pairs of other factions' favourites: $x_{j'} E_i x_{j''}$ for all $i \in F_j$ if $j' \neq j \neq j''$. Let ℓ_b be the benchmark lottery of drawing a random voter's favourite: $\ell_b(x_j) = N_j/N$. Assume there is just one more option y, and this is a potential full consensus option: $y\, P_i\, \ell_b$ for all i.

Assume the MaxParC ballot profile β has $\beta_i(x_j) = 100$, $\beta_i(x_{j'}) = 0$, and $0 < \beta_i(y) \leq 100/N$ for all $j \neq j'$ and $i \in F_j$. Then each $i \in F_j$ approves x_j and y, hence all vote for y and y is the sure winner.

Outcome with sincere voters. As discussed in 2.3.1, there is no unique way to vote "sincerely" in MaxParC, hence we rather discuss the results of voters applying one of the heuristics discussed there.

A voter applying the conservative satisficing heuristic rates their favourites x_j at 100 and the compromise y at $100(1 - u_i(\ell_b)/u_i(y)) > 0$. Also with the informed satisficing heuristic, the linear heuristic, and the hyperbolic heuristic, voters rate y at > 0. So if all voters apply one of these heuristics, y wins for sure.

Nash equilibrium. Since no $i \in F_j$ can make any $i' \notin F_j$ vote for x_j, the only way i could only improve the result would be by making the vote of some $i' \in F_j$ go to x_j instead of to y. But this is only possible by lowering $\beta_i(y)$ to zero (either certainly or with some positive probability), which will make everyone disapprove y and vote for their favourites, resulting in ℓ_b. Since no mixture of ℓ_b with the sure-thing lottery ℓ_y is an improvement for i, β is a *Nash equilibrium in pure strategies.*

Likewise, any group of voters $G \subseteq F_j$ from the same faction could only improve the result for each of them by making the vote of some $i' \in F_j$ go to x_j. As above, this is only possible if at least one $i \in G$ lowers $\beta_i(y)$ to zero, again resulting in ℓ_b, which is no improvement. Hence β remains a Nash equilibrium when some group of voters from the same faction is considered to act as one player; in particular, if each faction is considered one player (this could be called a "factional Nash equilibrium").

Still, as with other voting methods, it is easy to see that there are many other Nash equilibria (e.g., the less efficient one where everyone bullet-votes, resulting in ℓ_b), so the criterion of being a Nash equilibrium is not sufficiently discriminatory and stronger game-theoretic solution concepts are called for.

Strong Nash equilibrium. Assume a proper subgroup $G \subset E$ intersecting at least two factions, let's say it intersects the factions $F_1, \dots, F_r$. Assume the voters in G change their ballots in some way that improves the result for them all. Assume some $i \in G$ stops approving y. Then the votes of all $i' \notin F_1 + \dots + F_r$ will go to their favourites, hence, for $j = 1 \dots r$, at least $N_j + 1$ votes must go to x_j for this to be an improvement for all in G, which is impossible since there are only $N_1 + \dots + N_r$ votes left to distribute. Hence no $i \in G$ stops approving y, and no x_j gets approval by all voters, so y is still the sure winner, and there is no improvement for G after all.

Finally, assume the whole electorate could improve the result for all. Then there would be $\ell' \neq \ell_b$ with $\ell' \, P_i \, y$ for all i. Hence there would be $\ell'' \neq \ell_b$ with $\ell'' \, P_i \, y$ for all i and $\ell''(y) = 0$. But $\ell'' \neq \ell_b$ implies there is j with $\ell''(x_j) < \ell_b(x_j)$, so all $i \in F_j$ would have $y \, P_i \, \ell_b \, P_i \, \ell''$, a contradiction.

This shows that β is a *strong Nash equilibrium,* i.e., no group whether small or large, unanimous or cross-faction, has an incentive to deviate from β.

Other methods. Under majoritarian methods, in particular PV, AV, RV, IRV, and SC, there is usually no Nash equilibrium between the factions that would give y positive winning probability, simply because whenever $N_j > N/2$ for some j, faction F_j will enforce that x_j wins. Also, RB does not support full consensus since for all F_j it is strictly dominant to vote for x_j. FC and RFC however do support full consensus, as shown in [1].

With sincere voters, only AV and RV also support full consensus, while PV, IRV and SC would still elect x_j whenever $N_j > N/2$.

2.1.5.4 MaxParC favours full over partial consensus

Assumptions. As a generalization of the above, assume now that there is an additional option z, considered a potential partial consensus by the union $H = F_1 + \cdots + F_h \subset E$ of some of the factions, and considered equally bad by all others, so that $\ell_{b/z} \, P_i \, \ell_b$ for all $i \in H$ and $z \, E_i \, x_j$ for all $i \notin H + F_j$, where $\ell_{b/z}$ is the result of all $i \in H$ voting for z and all others voting for their favourites: $\ell_{b/z}(z) = |H|/N$ and $\ell_{b/z}(x_j) = |F_j|/N$ for all $j > h$.

Consider an extension of the above ballot profile β with $0 < \beta_i(z) \leq 100(1 + N - |H|)/N$ and $\beta_{i'}(z) = 0$ for all $i \in H$, $i' \notin H$. Note that then all $i \in F_j$ approve x_j, all $i \in H$ approve z, and all approve y, hence again y is the sure winner.

Outcome with sincere voters. A voter applying one of the heuristic in 2.3.1 will rate y at > 0 but will rate z at 0 if she is not a member of H. Hence if all voters apply some of these heuristics, y will be strictly more approved than z and still win for sure.

Strong Nash equilibrium. If all $i \in H$ consider full consensus still better than their potential partial consensus, i.e., $y \, P_i \, \ell_{b/z}$ for all $i \in H$, then we will show that β is again a strong Nash equilibrium, at least when all $i \in H$ have von Neumann–Morgenstern expected utility functions $u_i(\ell)$ over lotteries. Assume some group G can improve the result to some lottery ℓ' by modifying their ballots. If no $i \in G$ stopped approving y, y would remain the sure winner, hence some $i \in G$ stops approving y and the votes of all $i' \in H - G$ go to z, while those of $i \in F_j - G$ for $j > h$ go to x_j. As above, for all j with $F_j \cap G \neq \emptyset$, at least $N_j + 1$ votes must then go to x_j for this to be an improvement for all $i \in F_j \cap G$, hence less than

$|H|$ votes are left that could go to either z or some of $x_1, \dots, x_h$. Those $i \in F_j \cap G$, $j \leq h$, have $Nu_i(\ell') = v_z u_i(z) + v_j u_i(x_j)$, where v_z, v_j are the votes going to z or x_j, respectively, and $v_z + \sum_{j=1}^h v_j < |H|$. Since $Nu_i(y) > Nu_i(\ell_{b/z}) = |H|u_i(z)$ and $Nu_i(y) > Nu_i(\ell_b) = N_j u_i(x_j)$, we have $u_i(y) < u_i(\ell') = [v_z u_i(z) + v_j u_i(x_j)]/N < [v_z/|H| + v_j/N_j]u_i(y)$, i.e., $v_j > N_j(1 - v_z/|H|)$ for all $j \leq h$, thus $|H| - v_z > \sum_{j=1}^h v_j > |H|(1 - v_z/|H|) = |H| - v_z$, a contradiction. Note that the same kind of argument can be made if there are several partial potential consensus options z,z',$\dots$, if each pair of corresponding supporting groups H,H' is either disjoint or one contains the other (so that they form a hierarchy).

In other words, no group has an incentive to deviate from electing a good enough full consensus, even if a whole hierarchy of narrower and broader partial consensus options is available.

2.1.5.5 MaxParC supports a single partial consensus

Assumptions. Assume the same situation as in 2.1.5.4, but without the full consensus option y, so that only the partial consensus option z remains besides the favourites x_j. Then the same ballot profile β, just with y removed, leads to the partial consensus result $\ell(z) = |H|/N$ and $\ell(x_j) = |F_j|/N$ for $j > h$.

Outcome with sincere voters. If the voters from H apply the hyperbolic heuristic, they rate z at $100(1 - \ell(x_j)/u_i(z))$ which is by assumption larger than $100(1 - |H|/N)$, so they all end up voting for z.

With the linear heuristic, however, they may rate z too low for getting their votes since $100(u_i(z) - \ell(x_j))/(1 - \ell(x_j))$ might be smaller than $100(1 - |H|/N)$.

Strong Nash equilibrium. We can show this β is again a strong Nash equilibrium when voters have von Neumann–Morgenstern utilities. Assume some $i \in G \cap H$ stops approving z. Then, for each j with $G \cap F_j \neq \emptyset$, at least $N_j + 1$ votes must go to x_j for this to be an improvement for all in G, but for each j with $G \cap F_j = \emptyset$, all N_j votes go to x_j, a contradiction as above. So no $i \in G \cap H$ stops approving z. If $G \cap H \neq \emptyset$, some $i \in G - H$ must vote for z for those voters to profit from the deviation. But then not enough votes are left in $G - H$ to make all $i \in G - H$ profit as well. So $G \cap H = \emptyset$, but since $E - H$ has no potential for even partial consensus, they cannot improve over the benchmark lottery either. This completes the proof.

Other methods. Again, under majoritarian methods, in particular PV, AV, RV, IRV, and SC, there is no Nash equilibrium between the factions that would give z positive winning probability if one of the factions is in a majority. Also, RB does not support partial consensus since for all F_j it is strictly dominant to vote for x_j. FC and RFC also fail to support partial consensus: If z wins because the fallback was not invoked, the voters in $E \setminus H$ can cause the fallback to be invoked and have strict incentive to do so; if the fallback is invoked, no voter in any $F_j \subset H$ will vote for z since they have then a strict incentive to vote for x_j instead.

2.1.5.6 MaxParC supports disjoint partial consensuses If several disjoint groups of factions exist each of which has a potential partial consensus, the situation can get a little trickier, and the canonical ballot profile might not be a strong Nash equilibrium but only a coalition-proof equilibrium. We treat a simple special case first to demonstrate this.

Example. Assume $N = 6$, four factions of sizes $N_1 = N_4 = 1$ and $N_2 = N_3 = 2$ with favourites $x_1 \dots x_4$, and two potential partial consensus options z, z', with utilities as in the following table:

	faction	F_1	F_2	F_3	F_4
	size	1	2	2	1
utility	100	x_1	x_2	x_3	x_4
	75	z	z	z'	z'
	0	rest	rest	rest	rest

A canonical ballot profile β that realizes both partial consensuses is given by this table:

	faction	F_1	F_2	F_3	F_4
willingness $\beta_i(x)$	100	x_1	x_2	x_3	x_4
	51	z	z	z'	z'
	0	rest	rest	rest	rest

Although this is a Nash equilibrium between the invididual voters and a Nash equilibrium between the four factions, it is not a strong Nash equilibrium since the two middle factions can profit from approving their mutual favourites, i.e., deviating as follows:

	faction	F_2	F_3
willingness $\beta'_i(x)$	100	x_2	x_3
	51	z	z'
	35	x_3	x_2
	0	rest	rest

This will result in F_2, F_3 approving both x_2, x_3 so that these options get a higher approval (4) than z, z' (having 3), the votes of F_2, F_3 now go to x_2, x_3 in equal shares (due to the tiebreaker), and those of F_1, F_4 still go to z, z'.

Still, the above deviation by F_2, F_3 is not *coalition-proof* since each of these two factions has an incentive to betray the other by not performing the agreed deviation after all, i.e., by deviating from the planned deviation. E.g., if F_2 defects in this way, we have the profile

	faction	F_1	F_2	F_3	F_4
willingness $\beta''_i(x)$	100	x_1	x_2	x_3	x_4
	51	z	z	z'	z'
	35			x_2	
	0	rest	rest	rest	rest

which now makes F_2's and F_3's votes both go to x_2, profiting F_2 even more and leaving F_3 with strictly less than under β. Because of this risk of being betrayed by F_2, F_3 has few incentives to agree with F_2 to perform the original deviation β'.

Conjecture. More generally, we conjecture that under quite general conditions, there will be at least a certain type of coalition-proof equilibrium (similar to [23]) which results in the election of a broad consensus.

More specifically, consider the following type of situation: There are $M \geq 2$ disjoint blocks $B_1, \dots, B_M$ of voters, each block B_k having size $N_k = |B_k|$ and consisting of $m_k \geq 2$ disjoint factions $F_{k1}, \dots, F_{km_k}$, and we assume their sizes $N_{kj} = |F_{kj}|$ are all at least $2M$. Each faction F_{kj} has a distinct favourite option x_{kj}, each block B_k a potential partial consensus option y_k. No other options exist. Let $\ell_b(x_{kj}) = N_{kj}/N$ define the benchmark lottery and $\ell_c(y_k) = N_k/N$ define the partial consensus lottery. Each voter $i \in F_{kj}$ has a von Neumann–Morgenstern utility function with $u_i(x_{kj}) = 1 > u_i(y_k) > N_{kj}/N_k$ and $u_i(z) = 0$ for all other options. The expected utility for $i \in F_{kj}$ resulting from some ballot profile β' under MaxParC is thus $u_i(\beta') = M(\beta')(x_{kj}) + u_i(y_k)M(\beta')(y_k)$, where $M(\beta')$ is the resulting lottery.

Now consider the following "canonical" ballot profile β: For $i \in F_{kj}, \beta_i(x_{kj}) = 100$, $100(N - N_k)/N < \beta_i y_k \leq 100(1+N-N_k)/N$, and $\beta_i(z) = 0$ for all other options. Note that under MaxParC with this ballot profile, each $i \in F_{kj}$ approves x_{kj} and y_k, hence ends up voting for y_k, so that the resulting lottery is ℓ_c as desired.

Also assume that any group G of voters can secretly plan to deviate from β, leading to a modified profile β' with $\beta'_i = \beta_i$ for all $i \notin G$, but that no member of G can be sure that the others will actually perform the deviation; rather, any subgroup $H \subset G$ can secretly plan a further deviation β'' from β', with $\beta''_i = \beta'_i$ for all $i \notin H$.

Then we conjecture that with the above strategy profile β, if there is a group G with a deviation β' from β that strictly profits all members (i.e., $u_i(\beta') > u_i(\beta)$ for all $i \in G$), there is a subgroup $H \subset G$ with a further deviation β'' from β' that strictly profits all its members (i.e., $u_i(\beta'') > u_i(\beta')$ for all $i \in H$) and is strictly worse than β for at least one member of G (i.e., there is $i' \in G$ with $u_{i'}(\beta'') < u_{i'}(\beta)$).

2.1.5.7 Summary We have shown in this section that NL and MaxParC both support full and partial consensus in a number of archetypical decision situations, whereas all other eight studied methods and all majoritarian methods do not.

2.2 Results of simulation experiments

We simulated $2M$ = 2,587,812 decision problems in total and stored the resulting welfare, randomization, and satisfaction metrics of all ten methods and the preferences over all method pairs.

To analyse voter behaviour during the interactive phase, we also stored (i) the share of factional updates that led to a change in the faction's ballots (metric `moverate`), (ii) the share of trial-and-error updates that did not lead to a change in the voter's ballot (metric `keeprate`), and (ii) the share of problems in which the ballots after the interactive phase differed from before that phase (metric `interactivechanged`).

Table S1 gives an overview of all metrics' mean values.

In addition to univariate and bivariate statistics for all metrics, we also fitted an OLS generalized linear regression model for each metric Y, separately for each preference model U, using the following parameters as explanatory variables: dummy variables for the voting method (using RV as reference method); log-transformed numbers of voters (`nvoters`), options (`noptions`), and polling rounds (`npolls`); shares of LCP (`rshare_LCP`), HCP (`rshare_HCP`), strategic (`sshare_S`), trial-and-error (`sshare_T`), heuristic (`sshare_H`), and factional (`sshare_F`) voters; a dummy indicating whether the first option was a constructed compromise option (`with_compromise`); and the parameters of the preference model (`log(Bmr), Bmh, Bmiota` or `dim, log(omega), rho`). The regression analysis shows that the case number was large enough to distinguish the influences of all explanatory variables since almost all estimated coefficients were significantly different from zero.

2.2.1 Social welfare

2.2.1.1 Absolute welfare metrics All six absolute welfare metrics (`Wutil_initial`, `Wutil_final`, `Wgini_initial`, `Wgini_final`, `Wegal_initial`, `Wegal_final`) had considerably left-skewed distributions across problems. When distinguishing by preference model (`umodel`), one can see that this is due to the spatial preference models, and that their location depends strongly on the preference model (Fig. S1).

In the block (`BM`) and uniform (`unif`) preference models, the majoritarian methods (`PV, AV, RV, IRV, SC`) generated slightly larger utilitarian and slightly smaller egalitarian absolute welfare than the proportional methods (`RB, FC, RFC, NL, MPC`), being roughly equivalent on the intermediate Gini–Sen welfare metric. In the QA and LA models, the majoritarian methods also outperformed the proportional ones in the Gini–Sen and egalitarian absolute welfare metrics, most significantly in the QA model, less so in the LA model. In the GA model, the differences between methods were still statistically significant (e.g., Tbl. S2) but negligible in comparison to the overall dispersion of welfare across problems.

Throughout the regression models, more voter blocks (larger `BMr`), larger policy-space dimension (`dim`), and larger spatial voter heterogeneity (`omega`) decreased welfare, and so did a larger number of voters except in the `Wgini/GA` case. More options and larger block size heterogeneity (`BMh`) increased welfare. Larger individuality (`Bmiota`) increased utilitarian but decreased Gini–Sen and egalitarian welfare. For the spatial option broadness parameter (`rho`) and shares of non-EUT voters, there was no clear pattern. More pre-voting polling rounds had a statistically significant but very small negative influence.

Larger shares of sincere, trial-and-error, and factional voters and lower shares of lazy voters tended to increase welfare, the share of heuristic voters had no clear influence. Surprisingly, adding a constructed compromise option only increased welfare in the spatial models though

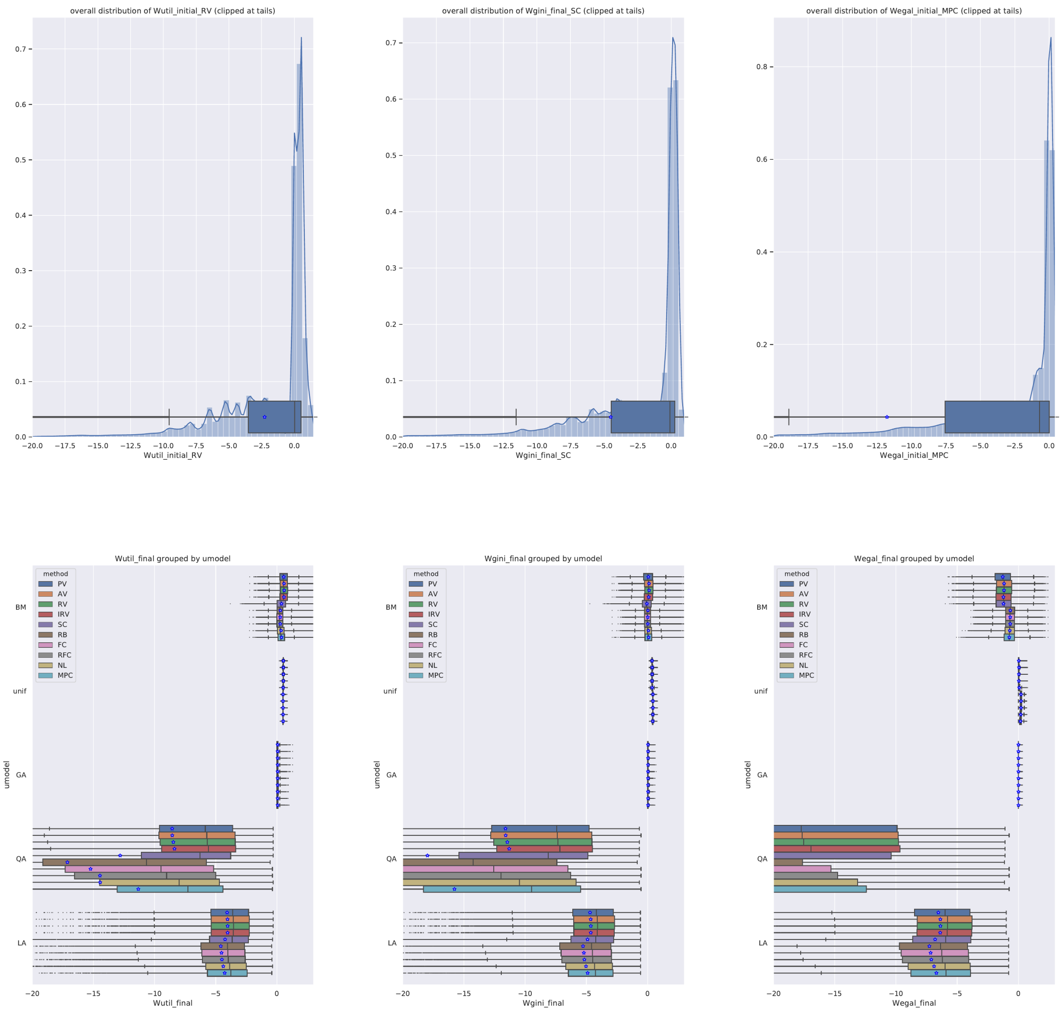

Figure S1: Distribution of absolute social welfare across decision problems. Top: histograms, kernel density estimators, and boxplots with means for three example metrics/methods. Bottom: distribution of final absolute welfare by preference model and method.

not in the block and uniform models.

2.2.1.2 Relative welfare metrics Because of the skewed distributions of absolute welfare, we also analysed the two versions of relative welfare metrics, Since the first of these was skewed in the other direction in case of the nondeterministic methods and inequality-averse metrics, we focus on the second, alternative version of relative welfare metrics here, which were much more balanced (Fig. S2).

Interestingly, in the uniform preference model, `alt_relWgini` and `alt_relWegal` had a particular trimodal distribution for the proportional methods, which performed similarly to the deterministic methods in most cases, but much better in a somewhat smaller cluster of cases and much worse in a still smaller cluster of cases.

Looking at the fairly inequality-averse Gini–Sen welfare metric in its "middle" version `alt_relWgini` more closely in a regression analysis, we see that this metric typically increased with the no. of options (except for the uniform preference model); it decreased with the no. of polling rounds, the share of lazy voters; the addition of a constructed compromise option, the risk attitude scenario, and the spatial broadness parameter had no consistent influence across preference models. In contrast to their influence on the absolute welfare metric `Wgini`, a larger no. of voter blocks `BMr`, individuality `BMiota`, policy space dimension `dim` increased `alt_relWgini`.

In addition, the grouped boxplots in Fig. S3 show some specific influences. A larger no. of voters increased `alt_relWgini` considerably for the proportional, but not for the majoritarian methods. The behavioural type scenario influenced `alt_relWgini` under the Simple Condorcet method much more than the other methods: interestingly, the all-sincere scenario produced much less welfare than the all-lazy one. Larger spatial option broadness `rho` had opposite effects for the majoritarian methods (increasing `alt_relWgini`) and the proportional methods (decreasing `alt_relWgini`): when options had narrow appeal (or candidates a narrow platform), proportional methods performed considerably better, with broadly appealing options (or platforms) it was the other way around. Finally, `alt_relWutil` decreased with increasing spatial voter heterogeneity `omega` more strongly for majoritarian methods, their clearest advantage mostly restricted to narrow spatial voter distributions.

2.2.1.3 Frequency of best-performing methods If we focus on the more qualitative question of which methods perform "best" how often, we can study the share of decision problems in which the largest welfare was (i) only provided by one or more majoritarian methods (red in Fig. S4), (ii) only provided by one or more proportional methods (green), or (iii) provided by at least one method from both groups (yellow). In Fig. S4, this is shown for all three final welfare metrics (`Wutil, Wgini, Wegal`), grouped by parameters that made a significant difference.

One can see that across all three welfare metrics, proportional methods were performing best according to this statistic in the Gaussian allotment and uniform preference models, and with fewer lazy voters.

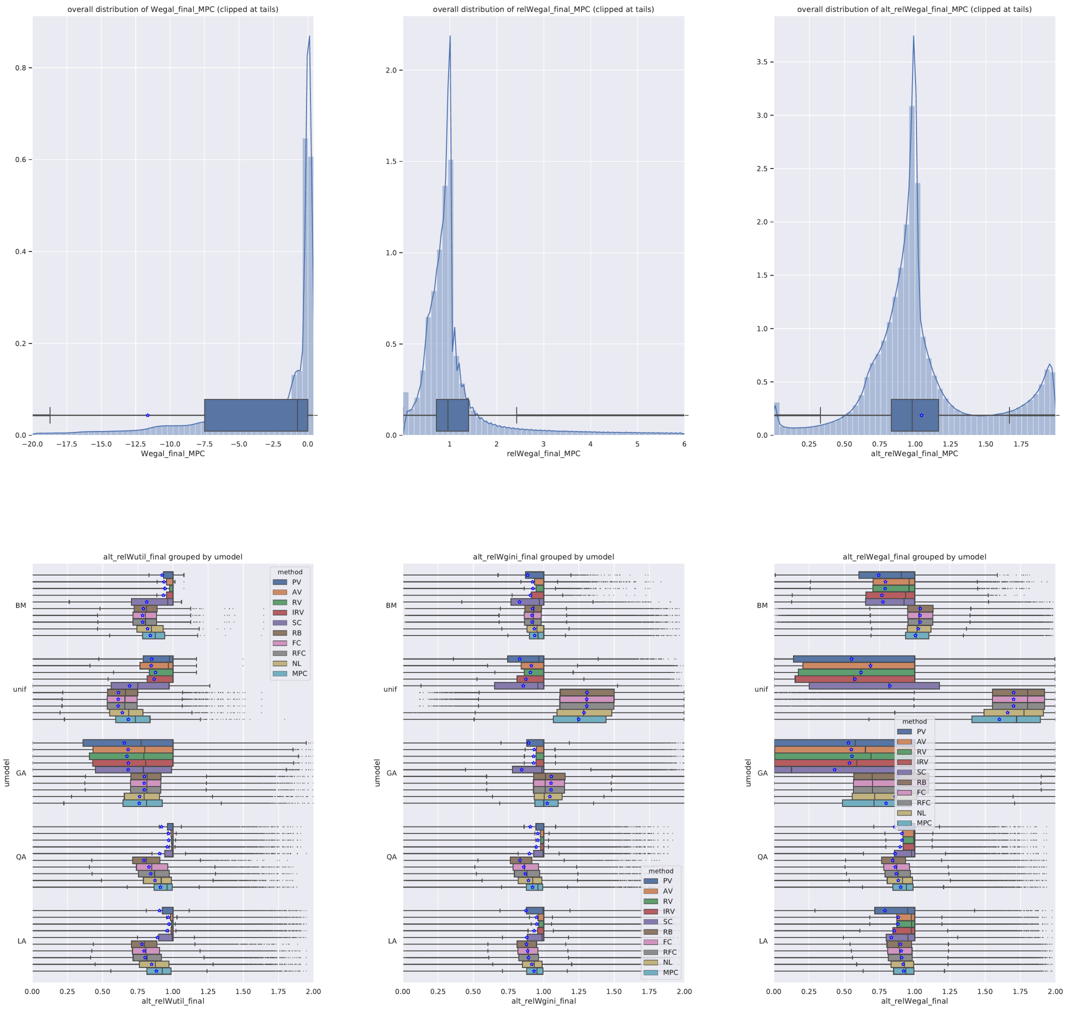

Figure S2: Distribution of relative social welfare across decision problems. Top: comparison of the two versions of relative welfare with absolute welfare for one combination of metric and method. Bottom: distribution of final alternative relative welfare metrics by preference model and method.

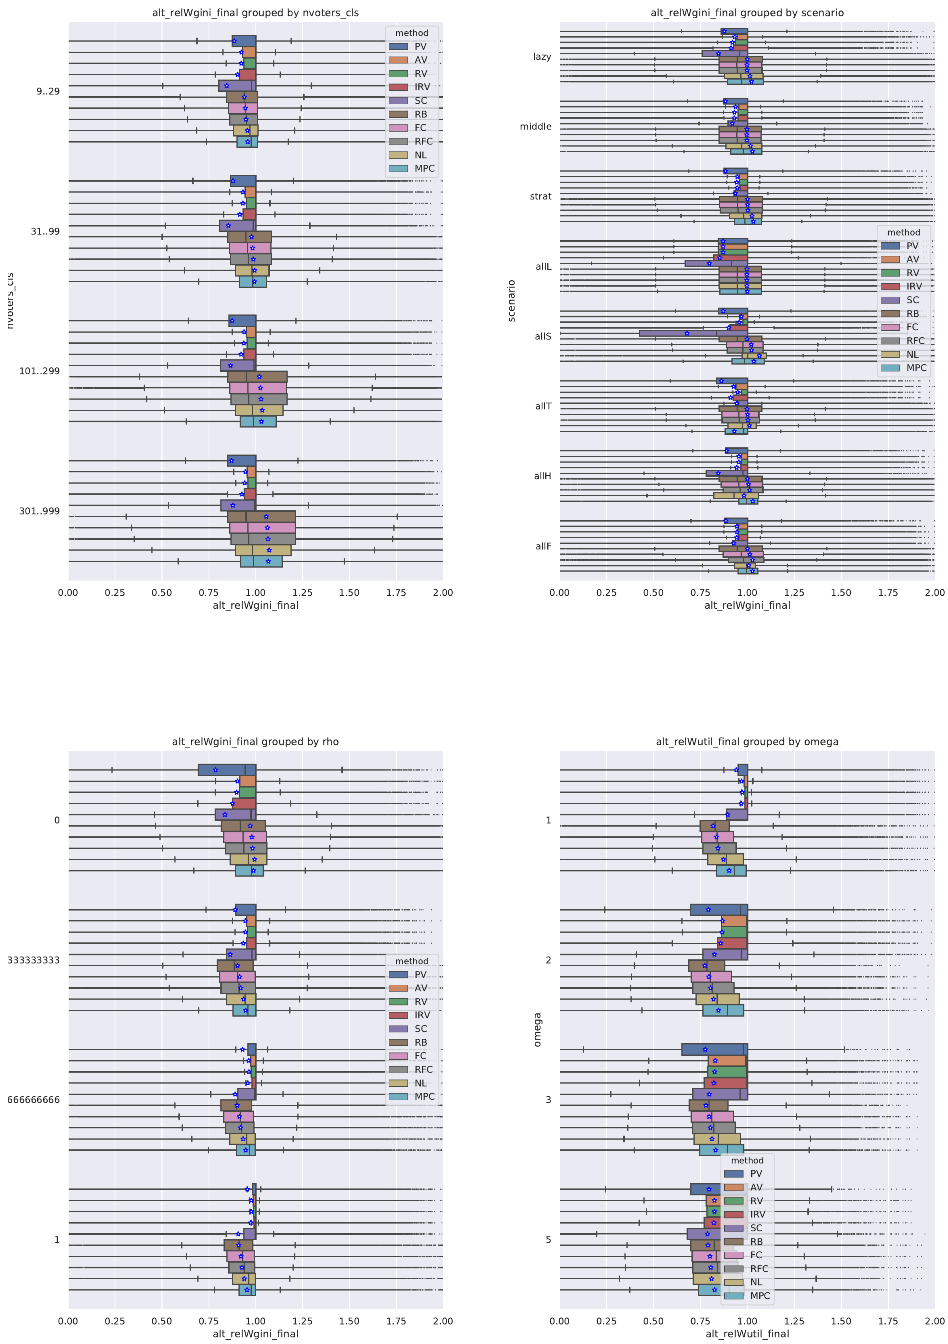

Figure S3: Distribution of relative social welfare across decision problems. Specific influences of no. of voters, behavioural type scenario, spatial option broadness, and spatial voter heterogeneity.

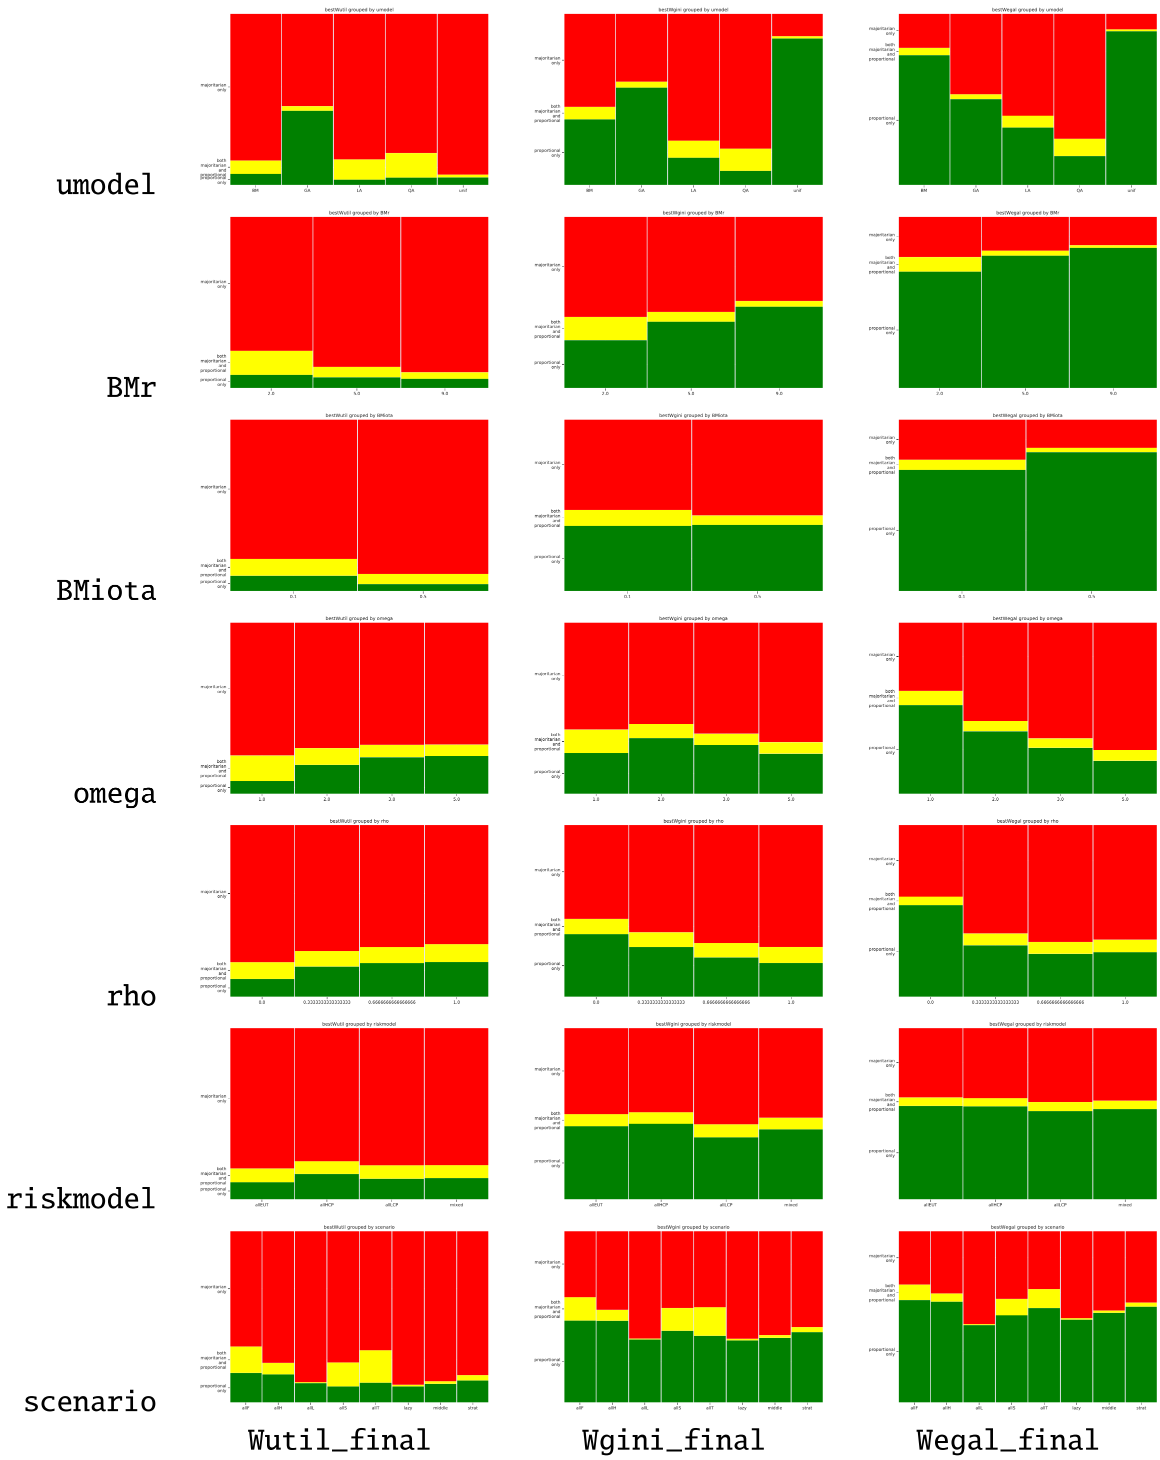

Figure S4: Frequency of majoritarian and proportional methods being best according to final utilitarian (left column), Gini–Sen (middle), and egalitarian (right) social welfare, by various parameters (rows).

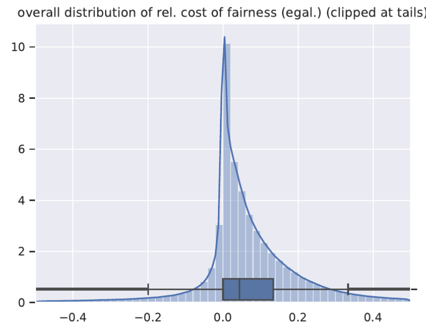

Figure S5: Distribution of the "relative cost of fairness" across decision problems.

W.r.t. Gini–Sen and egalitarian welfare, they performed best with larger no. of voter blocks, more individuality, larger spatial voter heterogeneity, and lower spatial option broadness; w.r.t. utilitarian welfare, this was the other way around.

2.2.1.4 "Cost of fairness" As a final indicator of the social welfare effects of using proportional instead of majoritarian methods, we show the distribution of the above-defined cost of fairness measure across all simulated decision problems in Fig. S5. It shows that, typically, the decrease (if any) in average voter utility one gets by switching from the best majoritarian method (Range Voting) to the best proportional (and thus "fair") method in our study (MaxParC), was about an order of magnitude smaller than the difference between the utilitarian and egalitarian absolute social welfare resulting from using Range Voting (which can be seen as a natural measure of absolute inequality in voters' utility). The 2%-trimmed mean of this "relative cost of fairness" was 0.08.

2.2.1.5 Summary Overall, one can conclude that proportional methods can well compete with majoritarian ones regarding social welfare effects and that the welfare assessment depends strongly on the choice of welfare metric used: the more inequality-averse the welfare metric, the more it favours proportional methods. In addition, welfare effects of method choice depend strongly on the distribution of voters' preferences, which suggests that proportional methods may particularly well perform in situations with heterogeneous voters and when many options have a comparatively narrow appeal.

Since proportional methods achieve this by randomization, and hence only over sequences of several decisions (which is why we used 'ex-ante'/long-run welfare metrics here), we need to look next at the amount of randomization actually used.

2.2.2 Randomization

In "deterministic" methods, randomization is only used to resolve the odd tie, hence Shannon entropy (and similarly Rényi entropy) is mostly zero, sometimes $\log 2$, and rarely larger, and maximal option probability is mostly one, sometimes $1/2$, and rarely smaller. Under RB, entropy is distributed in a left-skewed distribution with a peak at $\log$ `noptions`, leading to a mixture distribution with mean around 1.3. Under FC, this distribution is further mixed with a peak at 0 representing the cases where the full compromise was found. RFC was similar to FC regarding randomization.

Under NL, this probability of finding a full compromise was almost twice as large; while it still shows a mixture of left-skewed distributions with peaks at $\log k$ for some integer k, these k are now generally smaller than for RB since the optimization of the Nash sum typically leaves some options with zero probability. The mean entropy for NL is thus smaller, 1.05. For MPC, this is even more pronounced, with two clear peaks at 0 and $\log 2$ and mean entropy initially 0.9. For MPC, the interactive phase changed entropy the most, bringing it down to 0.8 on average.

For RB and FC, the 25% quantile [and mean] of the largest option probability were only at around 0.3 [0.43], for NL is was at around 0.35 [0.53], and for MPC initially around 0.45 [0.6] and finally (after the interactive phase) around 0.5 [0.65].

Regression analysis reveals that Shannon entropy increases (and max. probability decreases) with the no. of options, no. of voters (slightly), no. of voter blocks, and share of lazy voters (considerably). Adding a constructed compromise option decreased entropy (and increased max. probability) even if it had no clear welfare effects. Surprisingly, the share of non-expected-utility, risk- (and hence randomization-) averse voters had no clear effect on the level of randomization, and neither had the no. of pre-voting polls.

Fig. S7 shows some further parameter influences on individual methods' level of randomization. With more options and more trial-and-error voters, MPC's advantage becomes more pronounced, while NL randomizes less if all voters are sincere or are factionally strategic, the latter reflecting the fact that the factional strategy in NL was implemented as an optimization problem in our simulations which however might be hard to solve in reality.

In the next subsection we will see whether strategizing pays off and whether it gives a relative advantage.

2.2.3 Satisfaction

In order to see whether and when different behavioural types of voters have advantages, we study the distribution of the average satisfaction of all voters of a certain type across our simulations. Fig. S8 shows that typical shape of these distributions depends very much on the method, but is almost identical for the two polar behavioural types of lazy voters and factionally strategic voters (and also for the other three types). In other words, voting strategically does not so much give a comparative advantage to the strategic voters over the lazy ones, but rather increases overall welfare (as seen above).

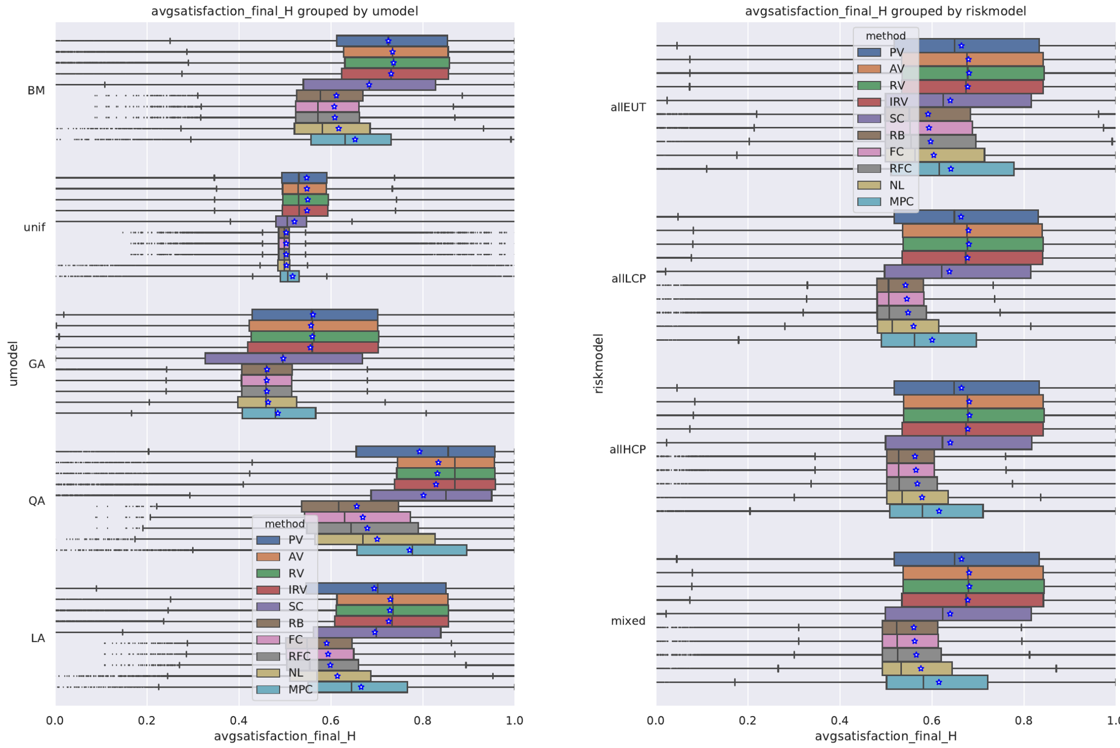

Figure S6: Distribution of the final average satisfaction of heuristic voters across decision problems, by method.

As can be seen in Fig. S6, the majoritarian methods fare somewhat better regarding average voter satisfaction with mean values around 0.66 compared to MPC's mean of around 0.61, again quite much depending on the preference model. As can be expected, for the proportional methods the risk-averse non-EUT voters were less satisfied.

2.2.4 Preferences over methods

If voters were to decide between the ten voting methods and would use for this "meta-decision" a pairwise comparison method such as Simple Condorcet or any other majoritarian method, which voting method would win?

If one assumes that voters are purely consequentialist and judge a method only by its generated utility, the surprising answer seems to be that they would then end up with Instant-Runoff Voting. As table S1 shows, IRV is the Condorcet Winner of this meta-decision since it would win a pairwise decision against all other nine methods.[16]

However, in an actual decision about a future voting method, people would however probably also have non-consequentialist preference components related to fairness, consistency, and other criteria such as those discussed in section 2.1.

2.2.5 Further conclusions

Our experiments indicate that assessments via agent-based simulations involving individual preferences will usually depend very much on the particular assumptions about voters' preference distributions, whether from a spatial or other model (or "culture"), and on the particular functional forms (e.g., linear, quadratic, or Gaussian) and parameter values used in these models.

They also seem to suggest that using an interactive phase only rarely has any considerable effect on the most important metrics, with the decrease in randomization under MaxParC being an exception. This might however be due to our very restricted assumptions on what agents can do during the interactive phase. Future work, whether empirical or numerical, should therefore consider the possibilities of information gathering, communication and other forms of social dynamics during the interactive phase.

To this end, we are currently developing a social app that offers an interactive version of MaxParC for making everyday group decisions, which we plan to use in empirical studies to assess the real-world performance and social dynamics of nondeterministic proportional consensus decision making methods.

2.3 Derivation of heuristics and strategies

2.3.1 Sincere voting under MaxParC

Since the MaxParC ballot explicitly asks for a quantity (the level "willingness" to approve an option) whose meaning ultimately depends on ones beliefs about the other voters' preferences, voting under MaxParC always incorporates some form of "strategic" thinking in some sense, so it is in a way pointless to ask what "the" sincere way of filling in a MaxParC ballot is. Rather, one may apply any of a number of different heuristics that all lead to a sincere ballot in the sense that more-preferred options are assigned higher willingness values.

2.3.1.1 Conservative satisficing heuristic. This sincere voting heuristic is based on the idea to assign to any option y a willingness $b(y)$ just small enough to guarantee that if I end up approving y and all who don't approve y approve their favourite only, the resulting lottery will not be worse than the benchmark lottery ℓ that would result if all approve their favourite only.

[16]Which method would win if the meta-decision was made using a proportional method, we can only speculate here since our current results do not provide us a way to predict this.

To find this willingness value $b(y)$, a EU-type voter would proceed as follows. Assume $\ell(x) > 0$ for all $x \in C$ (otherwise ignore those x for which $\ell(x) = 0$ in the following). Sort the options into an ordering $x_1, x_2, \dots$ by descending utility, so that $u_i(x_1) > u_i(x_2) > \dots > u_i(x_k)$ and $x_1 = f_i$ is the favourite of i. For all $a = 0 \dots k$, let $F_a = \sum_{c=a+1}^{k} \ell(x_c)$ and $U_a = \sum_{c=a+1}^{k} \ell(x_c) u_i(x_c)$, noting that $1 - \ell(f_i) = F_1 > \dots > F_k = 0$, $U_0 = u_i(\ell)$ and $U_k = 0$. For each $\beta \in [0, 1 - \ell(f_i)]$, let $a(\beta)$ the smallest a with $F_a \leq \beta$, noting that $a(0) = k$ and $a(1 - \ell(f_i)) = 1$. For each $y \in C \setminus \{f_i\}$ with $u_i(y) \geq u_i(\ell)$, let

$$V_y(\beta) = (1 - \beta) u_i(y) + U_{a(\beta)} + (\beta - F_{a(\beta)}) u_i(x_{a(\beta)}), \tag{43}$$

noting that $V_y(0) = u_i(y) \geq u_i(\ell)$ and $V_y(1 - \ell(f_i)) = u_i(\ell) - \ell(f_i)(u_i(f_i) - u_i(y)) < u_i(\ell)$. $V_y(\beta)$ is the evaluation of the lottery that results if the βN voters whose favourites I like least (this is the "conservative" aspect of the heuristic) approve their favourite only while the rest (including me) approve also of y so that y will win with probability $1 - \beta$ while the rest of the winning probability goes to options I rather don't like. The "satisficing" aspect of the heuristic is to be satisfied if this evaluation is not worse than that of the benchmark lottery. Hence one lets $b(y) = 100\beta$ for the largest β with $V_y(\beta) \geq u_i(\ell)$, so that $0 \leq b(y) < 100(1 - \ell(f_i))$, and completes the ballot by putting $b(f_i) = 100$ and $b(y) = 0$ for all $y \in C$ with $u_i(y) < u_i(\ell)$.

LCP-type voters may use the same formula based on the probability weights w_j instead of the actual probabilities $\ell(x_j)$.

Also for HCP-type voters, one can easily derive a similar formula.

Note that this heuristic indeed produces a sincere ballot since $u_i(y') > u_i(y)$ will imply $b(y) > b(y')$. If one has to expect that a fraction $\alpha \in [0, 1]$ of all voters is lazy, one would adjust $b(y)$ to $b(y) = 100(\alpha + (1 - \alpha)\beta)$ if $\beta > 0$.

2.3.1.2 Informed satisficing heuristic. If more information about the other voters' preferences is available or can be estimated, one may rather want to apply this heuristic in which i assumes that if a fraction of voters j will eventually approve some option y, it will be those j with the largest relative utility $\rho_j(y) := \frac{u_j(y) - u_j(\ell)}{u_j(f_j) - u_j(\ell)}$. Hence let us assume i has beliefs regarding the distribution of $\rho_j(y)$ inside each faction F_x and hence can sort the voters into an ordering $j_1, \dots, j_N$ by descending $\rho_j(y)$, so that $\rho_{j_1}(y) > \dots > \rho_{j_N}(y)$ and so that she knows f_{j_a} for all a. If the first $n \leq N$ voters in this ordering rather assign their winning probability to y than to f_{j_a}, i's utility becomes

$$V_y(n) = (n u_i(y) + \sum_{a=n+1}^{N} u_i(f_{j_a}))/N. \tag{44}$$

Now if i is satisfied if this is no smaller than $u_i(\ell)$, she would seek the smallest n with $\rho_{j_n}(y) \leq \rho_i(y)$ and $V_y(n') \geq u_i(\ell)$ for all $n' > n$ and put $b(y) = 100(1 - n/N)$, or, if there is no such n, put $b(y) = 0$. Note that also this heuristic has $b(y) > 0$ iff $u_i(y) \geq u_i(\ell)$ (since then $n < N$).

However, this heuristic may produce insincere ballots in which $b(y') < b(y)$ despite $u_i(y') > u_i(y)$ since the voter ordering used for $b(y')$ may be completely different than the one used for

$b(y)$. Still, one can argue that in many situations, the ballot will be approximately sincere. This is because often (i) $\rho_{j_n}(y) \leq \rho_i(y)$ will imply $V_y(n') \geq u_i(\ell)$ for all $n' > n$, and hence $b(y) \approx 100(1 - |\{j : \rho_j(y) \geq \rho_i(y)\}|)$, and (ii) the distribution of $\rho_j(y)$ in the electorate will be similar for all relevant options y, so that $b(y)$ is approximately monotonic in $\rho_i(y)$. In a spatial model with concave utilities $u_i(y) = f(\|\eta_i - \xi_y\|)$ (such as the LH and QH models) and smoothly and widely distributed voter and option positions, i will indeed prefer y to a lottery of favourites of those voters j with $\rho_j(y) \geq \rho$ for any ρ since those voters are distributed approximately uniformly and symmetrically around η_y, so the average distance from η_i to their favourites is at least $\|\eta_i - \xi_y\|$, translating into an expected utility from the lottery that is below $u_i(y)$ since f is concave. More particularly, both in the 1-dimensional *LH* and the 2-dimensional *QH* model, the number of voters j with $\rho_j(y) \geq \rho_i(y)$ scales roughly linearly with $1 - \rho_i(y)$, hence $b(y)$ will scale roughly linearly with $\rho_i(y)$, whereas in a higher-dimensional model, $b(y)$ will become a concave function of $\rho_i(y)$.

Our next two heuristics mimic this linear or concave behaviour to some extent with much simpler formulae.

2.3.1.3 Linear heuristic. A much simpler heuristic is the one we assume in our simulations, where $b(y) = 100\left(\alpha + (1-\alpha)\frac{u_i(y)-u_i(\ell)}{\max_{x \in C} u_i(x)-u_i(\ell)}\right)$ for all $y \in C$ with $u_i(y) \geq u_i(\ell)$, i.e., one assigns a willingness of 0 to options worse than the benchmark, 100 to one's favourite, and interpolates linearly between 100α and 100 based on the options' utilities, where $\alpha \in [0, 1]$ is the expected share of lazy voters in the electorate.

One motivation for this heuristic is that under certain assumptions, it can be interpreted as an approximation of the conservative satisficing heuristic. Assume the number of options is large, their utilities $u_i(y)$ for i are distributed uniformly, say (without loss of generality) between 0 and $u_i(f_i) = 1$, and their benchmark winning probabilities $\ell(y)$ are not correlated to i's evaluations $u_i(y)$. Then $u_i(\ell) \approx 1/2$, F_a and $a(\beta)$ decrease approximately linearly in a or β, respectively, $U_{a(\beta)} \approx \beta^2/2$, and $\beta - F_{a(\beta)}$ is small. Hence $V_y(\beta) \approx (1-\beta)u_i(y) + \beta^2/2$, which equals $u_i(\ell)$ for $\beta \approx 2u_i(y) - 1$, hence $b(y) \approx 100\left(\alpha + (1-\alpha)\frac{u_i(y)-u_i(\ell)}{\max_{x \in C} u_i(x)-u_i(\ell)}\right)$ for all y with $u_i(y) \geq u_i(\ell)$. The same derivation can be made under the weaker assumption that only those options y with $u_i(y) \geq u_i(\ell)$ are numerous, have uniformly distributed $u_i(y)$ and have $\ell(y)$ uncorrelated to $u_i(y)$. These assumptions are, e.g., approximately fulfilled if k is large and utility follows the *LH* model. If, instead, utility depends more concavely on distance, as in the *QH* model, the linear heuristic will tend to produce larger willingness values than the conservative satisficing heuristic, hence will produce more compromise outcomes which however may sometimes be worse than the benchmark lottery for some voters. On the contrary, if utility depends more convexly on distance, as in the tails of the *GH* model, the linear heuristic will tend to produce smaller willingness values than the conservative satisficing heuristic, hence may sometimes not produce a partial consensus when there is a potential one.

2.3.1.4 Hyperbolic heuristic.

A little less simple is the heuristic that puts $b(f_i) = 100$, $b(y) = 0$ for all y with $u_i(y) < u_i(\ell)$, and $b(y) = 100(1 - \frac{u_i(\ell)-\min_j u_i(f_j)}{u_i(y)-\min_j u_i(f_j)})$ for all other y, which has a hyperbolical rather than a linear dependency on $u_i(y)$, growing fast for $u_i(y)$ slightly above $u_i(\ell)$ and much slower for $u_i(y)$ approaching $u_i(f_i)$.

Also this formula can be derived as an approximation of the conservative satisficing heuristic, under different assumptions on the distribution of utility. Assume that i considers all other options than f_i that occur as favourites of any voter as approximately equally bad, so that we can assume $u_i(f_i) = 1$ and $u_i(f_j) \approx 0$ for all j with $f_j \neq f_i$. Then $u_i(\ell) \approx \ell(f_i)$, $V_y(\beta) \approx (1-\beta)u_i(y)$, hence $b(y) \approx 100(1 - \frac{u_i(\ell)}{u_i(y)}) \approx 100(1 - \frac{u_i(\ell)-\min_j u_i(f_j)}{u_i(y)-\min_j u_i(f_j)})$. Since these assumptions on utility are even more extremely "convex" than in the GA model, the hyperbolic heuristic may be a better choice in Gaussian utility situations than the linear heuristic.

2.3.2 Heuristic Nash Lottery strategy

Assume $N \gg 1$ and $C = \{1, \dots, k\}$, put $m = k-1$, $e = (1, \dots, 1) \in \mathbb{R}^m$, $p = (\ell_1, \dots, \ell_m)$, $v = u_{1k}$, $w = (u_{11} - v, \dots, u_{1m} - v)$. We focus on voter 1's choice of ratings $r_{1*} = \beta_i$ and consider $s = r_{1k} \geq 0$, $t = (r_{11} - s, \dots, r_{1m} - s)$ with $t_x \geq -s$ the control variables, all vectors being column vectors. Then the Nash sum (= log of Nash lottery target function) is

$$f(p|s,t) = g(p|s,t) + h(p) \tag{45}$$

with

$$g(p|s,t) = \log(s + p^\top t), h(p) = \sum_i \log(r_{ik} + \sum_x p_x r_{ix}), \tag{46}$$

where summation over i means $i = 2 \dots N$ (likewise for j) and summation over x means $x = 1 \dots m$ (likewise for y, z).

Since $N \gg 1$, we can approximate

$$f(p|s,t) = g(q|s,t) + d^\top G(s,t) + h(q) + d^\top H d/2, \tag{47}$$

where

$$q = \arg\max_p h(p), \tag{48}$$

$$d = p - q, \tag{49}$$

$$G(s,t)_x = \partial_{p_x} g(p|s,t)|_{p=q} = \frac{t}{s + q^\top t} = \gamma(s,t)t, \tag{50}$$

$$H_{xy} = \partial_{p_x}\partial_{p_y} h(p)|_{p=q} = \sum_i \partial_{p_x} \frac{r_{iy}}{r_{ik} + \sum_z p_z r_{iz}} = -\sum_i \frac{r_{ix} r_{iy}}{(r_{ik} + \sum_z p_z r_{iz})^2} \tag{51}$$

with

$$\begin{aligned}
\gamma(s,t) &= \frac{1}{s+q^\top t} > 0, && (52)\\
\partial_s\gamma(s,t) &= -\frac{1}{(s+q^\top t)^2} = -\gamma(s,t)^2 < 0, && (53)\\
\nabla_t\gamma(s,t) &= -\frac{q}{(s+q^\top t)^2} = -\gamma(s,t)^2 q. && (54)
\end{aligned}$$

Assume that H is nonsingular with

$$I = H^{-1}, \tag{55}$$

and note that H, I are symmetric and negative semidefinite. Assume that $q_x > 0$ for all x and $t \neq 0$, so that $s + q^\top t > 0$ and $\gamma(s,t) < \infty$.

The Nash lottery $p^*(s,t)$ is that $p \in [0,1]^m$ with $e^\top p \leq 1$ which maximizes $f(p|s,t)$. Assume this is an interior solution (e.g., since there is at least one bullet voter for each option), then the first-order condition is

$$0 = \nabla_p f(p|s,t) = \nabla_d(d^\top G(s,t) + d^\top Hd/2) = G(s,t) + Hd, \tag{56}$$

hence

$$\begin{aligned}
d^*(s,t) &= -IG(s,t) = -\gamma(s,t)It, && (57)\\
\partial_s d^*(s,t) &= \gamma(s,t)^2 It, && (58)\\
\nabla_t d^*(s,t)^\top &= \gamma(s,t)^2 qt^\top I - \gamma(s,t)I. && (59)
\end{aligned}$$

Voter 1's expected utility is then

$$U(s,t) = v + p^*(s,t)^\top w = v + q^\top w + d^*(s,t)^\top w = v + q^\top w - \gamma(s,t)t^\top Iw. \tag{60}$$

If she considers abstaining (which is equivalent to putting $t \equiv 0$ and an arbitrary $s > 0$, w.l.o.g. $s = 1$) and wonders what small change in ratings Δr would improve her utility most, she would calculate

$$\begin{aligned}
\partial_{r_{1x}} U(s,t)|_{t\equiv 0} &= -\partial_{t_x}(\gamma(s,t)t^\top Iw)|_{t\equiv 0} = -(\partial_{t_x} t^\top|_{t\equiv 0})Iw = -(Iw)_x, && (61)\\
\partial_{r_{1k}} U(s,t)|_{t\equiv 0} &= -(\partial_s - \sum_x \partial_{t_x})(\gamma(s,t)t^\top Iw)|_{t\equiv 0} = \sum_x (Iw)_x. && (62)
\end{aligned}$$

Not knowing H and hence I, voter 1 might use the following heuristic to estimate an approximate H from the latest favourite polling data, simply assuming every other voter $j > 1$ is lazy and submits a bullet vote $r_{jx} = 1$ for their favourite option x, putting all others to $r_{jy} = 0$. In that case, assuming $f^p(x) > 0$ for all x,

$$H_{xy} \approx N^2\left(1/f^p(k) + \delta_{xy}/f^p(x)\right) \tag{63}$$

where $\delta_{xy} = 1$ iff $x = y$, else $\delta_{xy} = 0$. Hence H is a matrix filled with equal positive entries plus some positive diagonal. Its inverse then has

$$I_{xx} \approx \zeta f^p(x)\,(N - f^p(x)) \tag{64}$$

$$I_{xy} \approx -\zeta f^p(x) f^p(y) \tag{65}$$

for some $\zeta > 0$ and all $x \neq y$. This would imply that voter 1's utility grows fastest in the direction Δr_1 with

$$\Delta r_{1x} = -(Iw)_x \approx \zeta f^p(x)\left(N w_x - \sum_y f^p(y) w_y\right) = \zeta N f^p(x)(u_{1x} - \upsilon), \tag{66}$$

$$\Delta r_{1k} = \sum_x (Iw)_x \approx \zeta N \sum_x f^p(x)\,(\upsilon - u_{1x}) = \zeta N f^p(k)(u_{1k} - \upsilon), \tag{67}$$

where $\upsilon = \sum_{x=1}^k u_{1x} f^p(x)/N$ is voter 1's expected utility of the benchmark lottery based on the latest favourite polling data. A natural heuristic is then that voter 1 moves her ratings from $r_{1x} \equiv r_{1k} = 1$ as much in the above direction as is possible without any rating getting negative, i.e., putting

$$r_{1x} = 1 + \rho f^p(x)(u_{1x} - \upsilon) \geq 0 \tag{68}$$

for all $x = 1 \dots k$, where

$$\rho = 1/f^p(y)(\upsilon - u_{1y}) > 0, \tag{69}$$

$$y = \arg\min_{x=1}^{k} f^p(x)(u_{1x} - \upsilon), \tag{70}$$

so that $r_{1y} = 0$. In the special case where $f^p(x) \equiv N/k$ (e.g. before the first poll), we get $y = \arg\min_{x=1}^k u_{1x}$, $\rho = k/N(\upsilon - u_{1y})$, and $r_{1x} = \frac{u_{1x} - \min_{z=1}^k u_{1z}}{\upsilon - \min_{z=1}^k u_{1z}} \propto u_{1x} - \min_{z=1}^k u_{1z}$, i.e., voter 1 would then vote sincerely. If, however, some options appear to have much higher chances than others, she would exaggerate her stated preferences regarding those options that appear to have higher chances (high $f^p(x)$) while playing down her stated preferences regarding those options that appear to have lower chances, which can result in rating some promising well-liked compromise option higher than her favourite if the latter has low chances, or rating some lurking less-liked compromise option lower than a very improbable least-liked option.

2.3.3 Factional unanimous best response in IRV

We show that w.l.o.g., one can restrict the analysis on the described set A of ballots. First, assume some ballot ranks some option y which however gets eliminated before all higher-ranked options are eliminated. Then submitting a shorter ballot with y left out instead leads to the exact same tally process. Second, assume y is ranked but gets eliminated at the same point as when submitting the shorter ballot with y left out. Then submitting the shorter ballot also leads to the exact same tally process. Hence we can restrict our focus on ballots ranking only options that survive the elimination process strictly longer than when not ranked, and don't get eliminated before any higher-ranked option. For any ballot $b = (x_1, x_2, \ldots, x_\ell) \in A$, let $a(b) \in \{0, 1, \ldots, k-1\}$ be the number of options eliminated strictly before x_ℓ when submitting b, and assume that also $b' = (x_1, x_2, \ldots, x_\ell, y) \in A$. Note that if submitting b, y is eliminated after at least $a(b) + 1$ many options, but is not the winner (otherwise $b' \notin A$), hence there are at most $k - a(b) - 2$ many different $y \in C$ such that $(x_1, x_2, \ldots, x_\ell, y) \in A$. Thus the number $a'(b, y)$ of options eliminated strictly before y when submitting b (not b'!) is one of the numbers in $\{a(b)+1, \ldots, k-2\}$ and is different for all y for which $(x_1, x_2, \ldots, x_\ell, y) \in A$. If submitting b', y must survive longer than when submitting b, hence

$$a(b') > a'(b, y) \geq a(b) + 1 > a'((x_1, x_2, \ldots, x_{\ell-1}), x_\ell) + 1, \tag{71}$$

i.e., $a(b') \geq a(b) + 2$. This implies that any ballot $b = (x_1, x_2, \ldots, x_\ell) \in A$ can be uniquely encoded via a sequence of integers $(a'(\emptyset, x_1),\ a'((x_1), x_2),\ a'((x_1, x_2), x_3),\ \ldots,\ a'((x_1, x_2, \ldots, x_{\ell-1}), x_\ell))$ that fulfils

$$a'((x_1, \ldots), x_i) + 1 < a'((x_1, \ldots), x_{i+1} \tag{72}$$

for all i. There are less than 2^k such sequences in $0, \ldots, k-1$, hence $|A| \leq 2^k$.

References

[1] J. Heitzig, F. W. Simmons, *Social Choice and Welfare* **38**, 43 (2012).

[2] H. Aziz, A. Bogomolnaia, H. Moulin, *ACM EC 2019 – Proceedings of the 2019 ACM Conference on Economics and Computation* pp. 753–781 (2019).

[3] H. Moulin, *Fair division and collective welfare* (MIT Press, 2004).

[4] C. Duddy, *Mathematical Social Sciences* **73**, 1 (2015).

[5] J.-F. Laslier, *Handbook on Approval Voting* (2010), pp. 311–335.

[6] R. Carroll, J. B. Lewis, J. Lo, K. T. Poole, H. Rosenthal, *American Journal of Political Science* **57**, 1008 (2013).

[7] A. Bruhin, H. Fehr-Duda, T. Epper, *Econometrica* **78**, 1375 (2010).

[8] J. Behnke, S. Hergert, F. Bader, Stimmensplitting – Kalkuliertes Wahlverhalten unter den Bedingungen der Ignoranz (2004).

[9] J. Sommer, Wer wählt strategisch und warum? Eine Analyse strategischen Wahlverhaltens bei der Bundestagswahl 2013, Phd dissertation, Heinrich-Heine-Universität Düsseldorf (2015).

[10] G. Gigerenzer, W. Gaissmaier, *Annual review of psychology* **62**, 451 (2011).

[11] K. van der Straeten, J. F. Laslier, N. Sauger, A. Blais, *Social Choice and Welfare* **35**, 435 (2010).

[12] J. S. Bower-Bir, N. J. D'Amico, A Tool for All People, but Not All Occasions: How Voting Heuristics Interact with Political Knowledge and Environment (2013).

[13] J. Laslier, *Journal of Theoretical Politics* **21**, 113 (2009).

[14] R. Myerson, R. Weber, *American Political Science Review* **87**, 102 (1993).

[15] A. Dellis, *Handbook on Approval Voting*, J.-F. Laslier, M. R. Sanver, eds. (Springer, 2010), chap. 18, pp. 431–454.

[16] H. Moulin, *Econometrica* **47**, 1337 (1979).

[17] P. K. Bag, H. Sabourian, E. Winter, *Journal of Economic Theory* **144**, 1278 (2009).

[18] A. Sen, *Journal of Public Economics* **3**, 387 (1974).

[19] E. Koutsoupias, C. Papadimitriou, *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Springer Verlag, 1999), vol. 1563, pp. 404–413.

[20] D. R. Woodall, *Discrete Applied Mathematics* **77**, 81 (1997).

[21] T. N. Tideman, *Social Choice and Welfare* **4**, 185 (1987).

[22] M. Schulze, *Social Choice and Welfare* **36**, 1 (2011).

[23] B. D. Bernheim, B. Peleg, M. D. Whinston, *Journal of Economic Theory* **42**, 1 (1987).

	PV	AV	RV	IRV	SC	RB	FC	RFC	NL	MPC
anonymous	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
neutral	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Pareto-efficient w.r.t. stated preferences	yes	yes	yes	yes	yes	yes	no	no	yes	yes
strongly mono-raise monotonic	yes	yes	yes	no	yes	yes	no	no	no	yes
weakly mono-raise monotonic	yes	yes	yes	no	yes	yes	yes	no	yes ?	yes
weakly mono-raise-abstention mon.	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
independent from Pareto-dominated alternatives	partial	full	full	full	full	partial	no	no	full	full
independent from losing options	no	full	full	no	no	partial	no	no	full	full
independent from exact clones	no	yes	yes	yes	yes	yes	no	yes	yes	yes
stronger forms of clone-proofness	no	yes	yes	no	yes	yes	no	?	?	?
strategy-freeness	no	no	no	no	no	yes	no	no	no	no
reveals preferences	no	some	some	no	no	fav.	fav.	utility	fav.	fav.
allocates power proportionally	no	no	no	no	no	yes	yes	yes	yes	yes
supports full consensus with sincere voters	no	yes	yes	no	no	no	yes	yes	yes	yes
supports full consensus with strategic voters	no	no	no	no	no	no	yes	yes	yes	yes
supports partial consensus with strategic voters	no	no	no	no	no	no	no	no	yes	yes

	PV	AV	RV	IRV	SC	RB	FC	RFC	NL	MPC
`moverate`	0.15	0.178	0.193	0.0257	0.0768	0	0.00145	0.013	0.488	0.327
`keeprate`	0.289	0.315	0.248	0.293	0.356	0	0.208	0.211	0.198	0.282
`interactivechanged`	0.0697	0.192	0.113	0.0934	0.307	0	0.403	0.423	0.62	0.482
`Eshannon_initial`	0.0344	0.0703	0.0173	0	0.0795	1.32	1.25	1.23	1.06	0.914
`Eshannon_final`	0.0211	0.0434	0.0119	0	0.0515	1.32	1.28	1.26	1.05	0.809
`Erenyi2_initial`	0.0344	0.0703	0.0173	0	0.0795	1.21	1.14	1.12	0.958	0.792
`Erenyi2_final`	0.0211	0.0434	0.0119	0	0.0515	1.21	1.18	1.15	0.961	0.689
`maxprob_initial`	0.976	0.954	0.988	1	0.949	0.423	0.456	0.466	0.539	0.602
`maxprob_final`	0.985	0.971	0.992	1	0.966	0.423	0.448	0.46	0.532	0.649
`pcompromise_initial`	0.0883	0.308	0.236	0.171	0.309	0.0833	0.147	0.186	0.214	0.257
`pcompromise_final`	0.0921	0.27	0.234	0.192	0.431	0.0833	0.156	0.193	0.221	0.278
`Wutil_initial`	-2.3	-2.36	-2.27	-2.25	-4.06	-4.17	-3.8	-3.7	-4.04	-2.97
`Wutil_final`	-2.29	-2.28	-2.26	-2.24	-3.22	-4.17	-3.78	-3.62	-3.58	-2.93
`Wgini_initial`	-3.18	-3.26	-3.13	-3.09	-5.76	-5.86	-5.3	-5.15	-5.66	-4.07
`Wgini_final`	-3.17	-3.14	-3.11	-3.09	-4.51	-5.86	-5.27	-5.03	-4.98	-4.02
`Wegal_initial`	-8.43	-8.77	-8.32	-8.15	-16.8	-18.6	-16.5	-15.9	-17.1	-11.8
`Wegal_final`	-8.4	-8.33	-8.24	-8.13	-12.7	-18.6	-16.3	-15.4	-14.8	-11.6
`relWutil_initial`	0.813	0.857	0.87	0.857	0.681	0.637	0.653	0.657	0.694	0.723
`relWutil_final`	0.813	0.849	0.861	0.851	0.747	0.637	0.651	0.657	0.69	0.729
`relWgini_initial`	0.857	0.968	0.931	0.896	0.765	2.1e+20	2.1e+20	2.1e+20	2.1e+20	2.1e+20
`relWgini_final`	0.851	0.947	0.934	0.908	0.885	2.1e+20	2.1e+20	2.1e+20	2.1e+20	2.1e+20
`relWegal_initial`	1.93e+20	3.6e+26	1.91e+20	1.91e+20	1.91e+20	2.2e+67	2.2e+67	2.2e+67	2.2e+67	2.21e+67
`relWegal_final`	1.92e+20	1.35e+59	1.91e+20	1.91e+20	1.31e+21	2.2e+67	2.2e+67	2.2e+67	2.2e+67	2.21e+67
`alt_relWutil_initial`	0.85	0.886	0.893	0.884	0.739	0.754	0.764	0.767	0.792	0.813
`alt_relWutil_final`	0.849	0.879	0.886	0.879	0.796	0.754	0.762	0.767	0.788	0.814
`alt_relWgini_initial`	0.881	0.948	0.937	0.915	0.782	0.998	1.01	1.01	1.02	1.02
`alt_relWgini_final`	0.877	0.934	0.933	0.917	0.861	0.998	1	1.01	1.01	1.01
`alt_relWegal_initial`	0.701	0.781	0.749	0.717	0.68	1.07	1.08	1.08	1.08	1.07
`alt_relWegal_final`	0.694	0.764	0.75	0.726	0.744	1.07	1.08	1.08	1.07	1.04
`avgsatisfaction_initial_F`	0.664	0.684	0.684	0.68	0.616	0.564	0.568	0.57	0.588	0.614
`avgsatisfaction_final_F`	0.665	0.68	0.682	0.678	0.649	0.564	0.569	0.575	0.604	0.618
`avgsatisfaction_initial_H`	0.664	0.684	0.684	0.679	0.615	0.564	0.568	0.57	0.588	0.614
`avgsatisfaction_final_H`	0.664	0.68	0.681	0.677	0.639	0.564	0.566	0.57	0.579	0.618
`avgsatisfaction_initial_L`	0.657	0.672	0.672	0.667	0.581	0.564	0.564	0.565	0.6	0.597
`avgsatisfaction_final_L`	0.657	0.669	0.67	0.666	0.606	0.564	0.563	0.564	0.589	0.601
`avgsatisfaction_initial_S`	0.657	0.683	0.684	0.677	0.577	0.564	0.574	0.574	0.617	0.608
`avgsatisfaction_final_S`	0.658	0.68	0.681	0.676	0.602	0.564	0.572	0.574	0.608	0.613
`avgsatisfaction_initial_T`	0.658	0.68	0.684	0.678	0.578	0.564	0.573	0.574	0.619	0.609
`avgsatisfaction_final_T`	0.661	0.676	0.682	0.677	0.65	0.564	0.567	0.569	0.594	0.616
`pctprefer_PV_over`	—	18	15.4	13.3	34.3	65.4	63.4	63.1	61.2	58.7
`pctprefer_AV_over`	19.2	—	9.47	12.5	33.5	69.2	65.2	65.2	62.8	59.1
`pctprefer_RV_over`	17.3	10.4	—	10.8	33.3	69.1	65.3	65.2	63.2	59.4
`pctprefer_IRV_over`	15.8	14.1	11.6	—	32.8	68	65.1	64.9	63.1	60.1
`pctprefer_SC_over`	26.8	24.2	23.5	22.8	—	61.7	59.3	58.9	57	53.5
`pctprefer_RB_over`	34.4	30.6	30.7	31.9	38.1	—	20.6	20.4	37.4	26.9
`pctprefer_FC_over`	35.1	31.1	31.1	32.4	38.8	24.3	—	21.4	39.2	27.8
`pctprefer_RFC_over`	35.5	31.4	31.4	32.7	39.2	28.5	25.8	—	40	28.9
`pctprefer_NL_over`	36.9	34	33.4	34.2	41	62.4	59.5	58.7	—	42.3
`pctprefer_MPC_over`	38.4	34.1	34.4	35.5	42.4	63.6	59.5	58.7	56.4	—
	PV	AV	RV	IRV	SC	RB	FC	RFC	NL	MPC

Table S1: Level of compliance with voting method consistency criteria, and average performance metrics from agent-based simulations.

OLS Regression Results

Dep. Variable:	Wgini_final	R-squared:	0.661
Model:	OLS	Adj. R-squared:	0.661
Method:	Least Squares	F-statistic:	2.663e+05
Date:	Sat, 04 Apr 2020	Prob (F-statistic):	0.00
Time:	15:34:58	Log-Likelihood:	1.0244e+07
No. Observations:	5124153	AIC:	-2.049e+07
Df Residuals:	5124130	BIC:	-2.049e+07
Df Model:	22		
Covariance Type:	HC1		

	coef	std err	z	P>|z|	[0.025	0.975]
Intercept	0.1581	0.000	1223.248	0.000	0.158	0.158
PV	-0.0018	6.48e-05	-27.023	0.000	-0.002	-0.002
AV	5.993e-05	6.55e-05	0.916	0.360	-6.84e-05	0.000
IRV	-0.0002	6.53e-05	-3.193	0.001	-0.000	-8.05e-05
SC	-0.0030	6.58e-05	-44.931	0.000	-0.003	-0.003
RB	0.0019	6.47e-05	29.673	0.000	0.002	0.002
FC	0.0018	6.47e-05	28.243	0.000	0.002	0.002
RFC	0.0018	6.47e-05	28.405	0.000	0.002	0.002
NL	0.0017	6.5e-05	26.489	0.000	0.002	0.002
MPC	0.0014	6.52e-05	22.143	0.000	0.001	0.002
log(nvoters)	0.0003	1.06e-05	25.800	0.000	0.000	0.000
log(noptions)	0.0046	4.01e-05	115.265	0.000	0.005	0.005
with_compromise	0.0004	2.9e-05	12.443	0.000	0.000	0.000
rshare_LCP	0.0004	4.05e-05	10.453	0.000	0.000	0.001
rshare_HCP	0.0009	4.08e-05	22.671	0.000	0.001	0.001
log(npolls)	6.738e-05	1.87e-05	3.607	0.000	3.08e-05	0.000
sshare_S	0.0005	5.79e-05	9.036	0.000	0.000	0.001
sshare_T	0.0007	5.69e-05	13.150	0.000	0.001	0.001
sshare_H	0.0012	5.68e-05	20.744	0.000	0.001	0.001
sshare_F	0.0013	5.42e-05	23.297	0.000	0.001	0.001
dim	-0.0465	1.95e-05	-2384.367	0.000	-0.047	-0.046
log(omega)	-0.0447	3.11e-05	-1437.194	0.000	-0.045	-0.045
rho	-0.0018	3.89e-05	-46.644	0.000	-0.002	-0.002

Omnibus:	909575.095	Durbin-Watson:	1.060
Prob(Omnibus):	0.000	Jarque-Bera (JB):	1549242.812
Skew:	1.174	Prob(JB):	0.00
Kurtosis:	4.319	Cond. No.	63.3

Table S2: Generalized linear model for Gini–Sen absolute social welfare in the GA preference model.

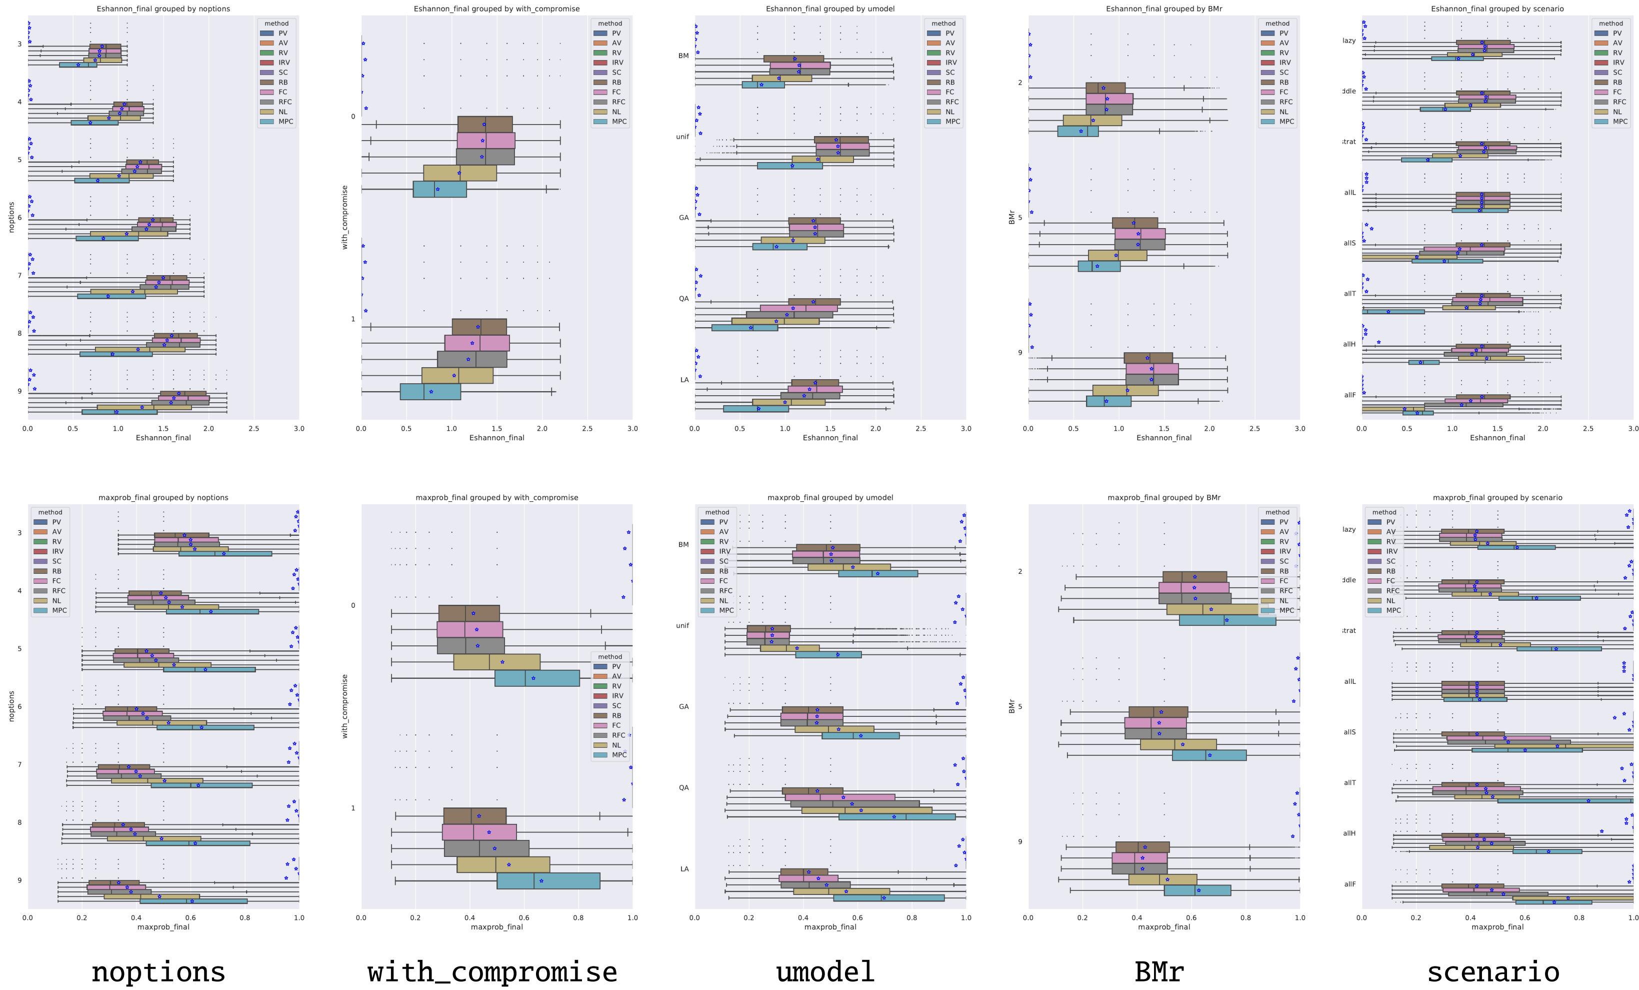

Figure S7: Statistics for final Shannon entropy (top) and maximal option probability (bottom) across decision problems for all methods, grouped by parameters with considerable influence.

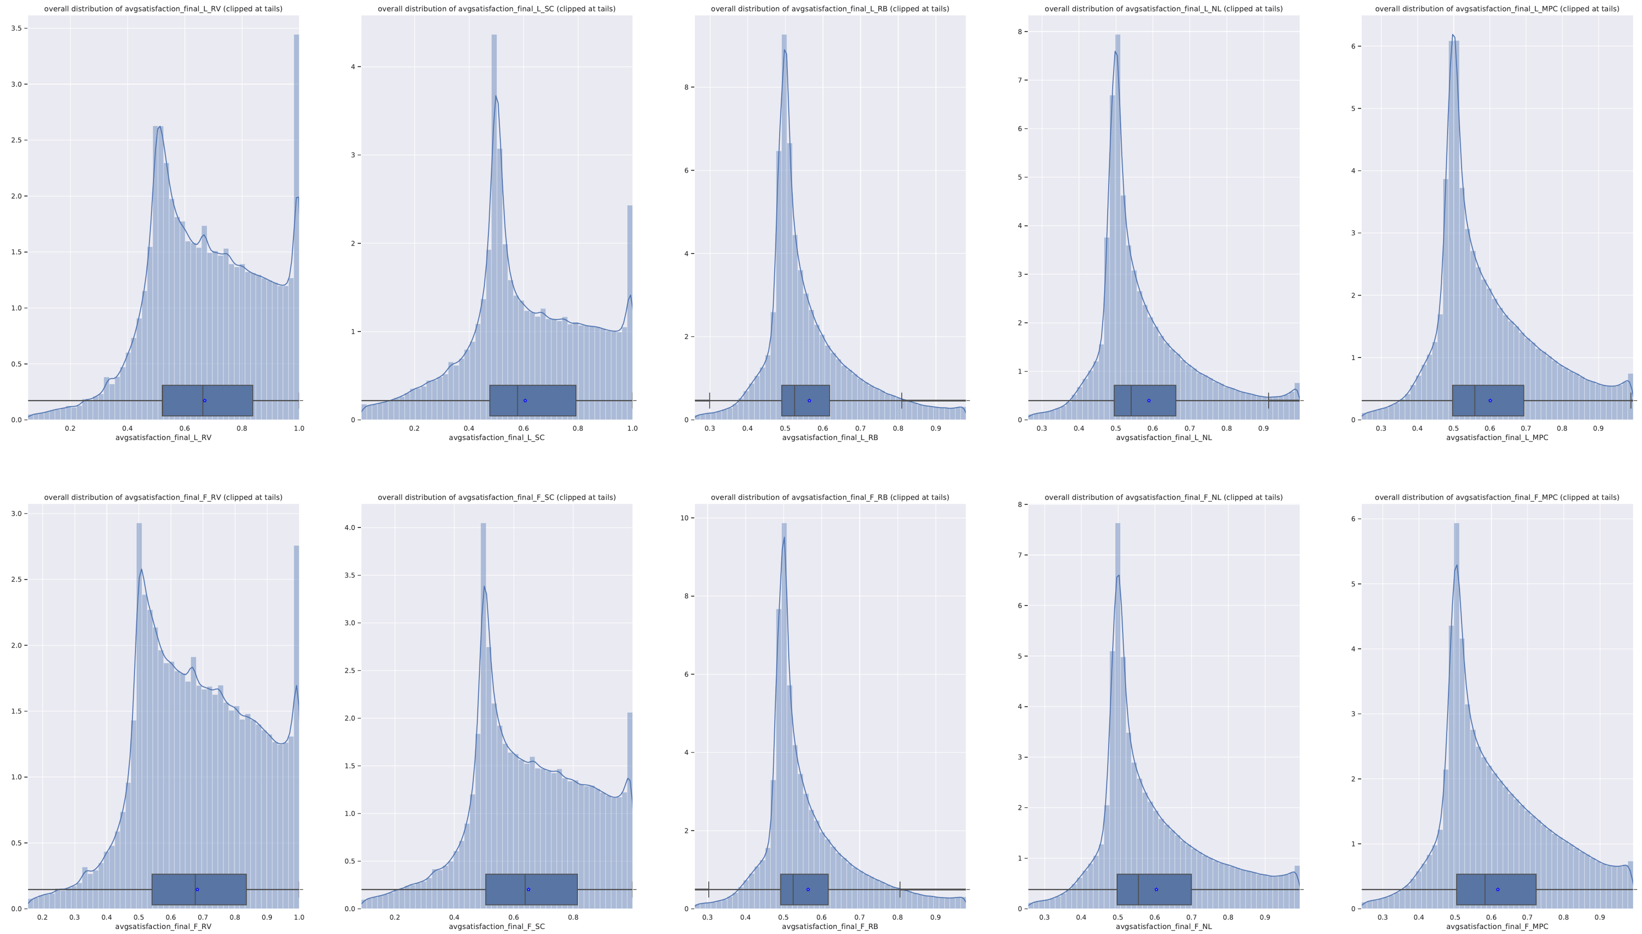

RV SC RB NL MPC

Figure S8: Distribution of final average satisfaction of lazy (top) and factionally strategic (bottom) voters across decision problems for selected methods.