Optimal Bounds for the No-Show Paradox via SAT Solving

Felix Brandt
TU München
Munich, Germany
brandtf@in.tum.de

Christian Geist
TU München
Munich, Germany
geist@in.tum.de

Dominik Peters
University of Oxford
Oxford, UK
dominik.peters@cs.ox.ac.uk

ABSTRACT

One of the most important desirable properties in social choice theory is *Condorcet-consistency*, which requires that a voting rule should return an alternative that is preferred to any other alternative by some majority of voters. Another desirable property is *participation*, which requires that no voter should be worse off by joining an electorate. A seminal result by Moulin [29] has shown that Condorcet-consistency and participation are incompatible whenever there are at least 4 alternatives and 25 voters. We leverage SAT solving to obtain an elegant human-readable proof of Moulin's result that requires only 12 voters. Moreover, the SAT solver is able to construct a Condorcet-consistent voting rule that satisfies participation as well as a number of other desirable properties for up to 11 voters, proving the optimality of the above bound. We also obtain tight results for set-valued and probabilistic voting rules, which complement and significantly improve existing theorems.

1. INTRODUCTION

In social choice theory, voting rules are usually compared using desirable properties (so-called axioms) that they may or may not satisfy. There are a number of well-known impossibility theorems—among which Arrow's impossibility is arguably the most famous—which state that certain axioms are incompatible with each other, and thus show the non-existence of voting rules that satisfy the given set of axioms. These results are important because they clearly define the boundary of what can be achieved at all. In stating such impossibility results, certain axioms are often not made explicit but are used as modeling assumptions. This is commonly the case for preference domain assumptions, and for bounds on the numbers of voters and alternatives. For instance, if there are only two alternatives, Arrow's theorem does not apply: the majority rule, among many others, satisfies the conditions used in Arrow's theorem. One impossibility whose proof requires unusually high bounds on the number of voters and alternatives is Moulin's *no-show paradox* [29], which states that the axioms of Condorcet-consistency and participation are incompatible whenever there are at least 4 alternatives and 25 voters. Now, Moulin proves that the bound on the number of alternatives is tight by showing that the maximin voting rule (with lexicographic tie-breaking) satisfies the desired properties when there are at most 3 alternatives. However, it was left open whether the more restrictive condition on the number of voters is tight as well.

The goal of this paper is to give tight bounds on the number of voters required for Moulin's theorem and related theorems that appear in the literature. For example, we show that Condorcet-consistency and participation are incompatible for 4 alternatives and 12 voters, but that the same result does not hold with 11 voters. To achieve results like this, we encode our problems as formulas in propositional logic and then use SAT solvers to decide their satisfiability and extract minimal unsatisfiable sets (MUSes) in the case of unsatisfiability. This approach is based on previous work by Tang and Lin [35], Geist and Endriss [20], Brandt and Geist [7], and Brandl et al. [3]. However, it turned out that a straightforward application of this methodology is insufficient to deal with the magnitude of the problems we considered. Several novel techniques were necessary to achieve our results. In particular, we extracted knowledge from computer-generated proofs of weaker statements and then used this information to guide the search for proofs of more general statements.

As mentioned above, Moulin's theorem shows a conflict between the axioms of Condorcet-consistency and of participation. Condorcet-consistency refers to one of the most influential notions in social choice theory, namely that of a *Condorcet winner*. A Condorcet winner is an alternative that is preferred to any other alternative by a majority of voters. Condorcet, after whom this concept is named, argued that, whenever a Condorcet winner exists, it should be elected [14]. A voting rule satisfying this condition is called *Condorcet-consistent*. While the desirability of Condorcet-consistency—as that of any other axiom—has been subject to criticism, many scholars agree that it is very appealing and so a large part of the social choice literature deals exclusively with Condorcet-consistent voting rules (e.g., [17, 26, 8]). *Participation* was first considered by Fishburn and Brams [18] and requires that no voter should be worse off by joining an electorate, or—equivalently—that no voter should benefit by abstaining from an election. The desirability of this axiom in any context with voluntary participation is evident. All the more surprisingly, Fishburn and Brams have shown that single transferable vote (STV), a common voting rule, violates participation and referred to this phenomenon as the *no-show paradox*. Moulin [29], perhaps even more startlingly, proved that no Condorcet-consistent voting rule satisfies participation when there are at least 25 voters.

We leverage SAT solving to obtain an elegant human-readable proof of Moulin's result that requires only 12 voters. While computer-aided solving techniques allow us to tackle difficult combinatorial problems, they usually do not provide additional insight into these problems. To our surprise, the computer-aided proofs we found possess a certain kind of symmetry that has not been exploited in previous proofs. Moreover, the SAT solver is able to construct a Condorcet-consistent voting rule that satisfies participation as well as a number of other desirable properties for up to 11 voters, proving the optimality of the above bound. In 99.8% of all cases, this computer-generated voting rule selects alternatives that are returned by the maximin voting rule and, in contrast to maximin, only selects alternatives from the top cycle. As a practical consequence of our theorem, strategic abstention need not be a concern for Condorcet-consistent voting rules when there are at most 4 alternatives and 11 voters, for instance when voting in a small committee. We also use our techniques to provide optimal bounds for related results about set-valued and probabilistic voting rules [23, 36]. In particular, we give a tight bound of 17 voters for the optimistic preference extension, 14 voters for the pessimistic extension, and 12 voters for the stochastic dominance preference extension. These results are substantial improvements of previous results. For example, the previous statement for the pessimistic extension requires an additional axiom, at least 5 alternatives, and at least 971 voters [23]. Our results are summarized in Table 1.

2. RELATED WORK

The no-show paradox was first observed by Fishburn and Brams [18] for the STV voting rule. Ray [31] and Lepelley and Merlin [27] investigate how frequently this phenomenon occurs in practice. The main theorem addressed in this paper is due to Moulin [29] and requires at least 25 voters. This bound was recently brought down to 21 voters by Kardel [24]. Simplified proofs of Moulin's theorem are given by Schulze [33] and Smith [34]. Holzman [22] and Sanver and Zwicker [32] strengthen Moulin's theorem by weakening Condorcet-consistency and participation, respectively. Duddy [15] shows the incompatibility of Condorcet-consistency and weaker notions of participation when allowing weak preferences. Pérez [30] considers these notions in the context of set-valued voting rules and shows that all common Condorcet extensions except the maximin rule and Young's rule violate these properties. Jimeno et al. [23] prove variants of Moulin's theorem for set-valued voting rules based on the optimistic and the pessimistic preference extension. Determining optimal bounds on the number of voters for these paradoxes has been recognized as an open problem. For example, Pérez notes that "a practical question, which has not been dealt with here, refers to the number of candidates and voters that are necessary to invoke the studied paradoxes" ([30], p. 614) and Duddy [15] concludes that "we do not know what upper bound is imposed on the number of potential voters by the conjunction of Condorcet consistency and [...] the participation principle in the case of linear orderings. And these upper bounds may fall as the number of candidates rises. These are important open problems since voting is often conducted by small groups of individuals." The influence of the number of voters and alternatives has recently also been studied in other contexts of social choice theory (see, e.g., [10, 11]).

When assuming that voters have incomplete preferences over sets or lotteries, participation and Condorcet-consistency can be satisfied simultaneously and positive results for common Condorcet-consistent voting rules (such as the top cycle) have been obtained by Brandt [6] and Brandl et al. [3, 4]. A particularly positive result was recently obtained for maximal lotteries, a probabilistic Condorcet extension due to Fishburn [5]. Maximal lotteries satisfy a notion of participation that lies in between strong SD-participation and weak SD-participation.

The computer-aided techniques in this paper are inspired by Tang and Lin [35], who reduced well-known impossibility results from social choice theory—such as Arrow's theorem—to finite instances, which can then be checked by a SAT solver. This methodology has been extended and applied to new problems by Geist and Endriss [20], Brandt and Geist [7], and Brandl et al. [3]. The results obtained by computer-aided theorem proving have already found attention in the social choice community [12]. More generally, SAT solvers have also proven to be quite effective for other problems in economics. A prominent example is the ongoing work by Fréchette et al. [19] in which SAT solvers are used for the development and execution of the FCC's upcoming reverse spectrum auction. In some respects, our approach also bears some similarities to *automated mechanism design* (see, e.g., [13]), where desirable properties are encoded and mechanisms are computed to fit specific problem instances.

3. PRELIMINARIES

Let A be a set of m alternatives and $\mathcal{N}$ be a set of n voters. Whether no-show paradoxes occur depends on the exact values of m and n. By $\mathcal{E}(\mathcal{N}) := 2^{\mathcal{N}} \setminus \{\emptyset\}$ we denote the set of *electorates*, i.e., non-empty subsets of $\mathcal{N}$. For many of our results, we will take $A = \{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$, and we use the labels $\mathtt{x}, \mathtt{y}$ for arbitrary elements of A.

A *(strict) preference relation* is a complete, antisymmetric, and transitive binary relation on A. The preference relation of voter i is denoted by $\succcurlyeq_i$. The set of all preference relations over A is denoted by $\mathcal{R}$. For brevity, we denote by $\mathtt{abcd}$ the preference relation $\mathtt{a} \succcurlyeq_i \mathtt{b} \succcurlyeq_i \mathtt{c} \succcurlyeq_i \mathtt{d}$, eliding the identity of voter i, and similarly for other preferences.

A *preference profile* R is a function from an electorate $N \in \mathcal{E}(\mathcal{N})$ to the set of preference relations $\mathcal{R}$. The set of all preference profiles is thus given by $\mathcal{R}^{\mathcal{E}(\mathcal{N})}$. For the sake of adding and deleting voters, we define for any preference profile $R \in \mathcal{R}^N$ with $(i, \succcurlyeq_i) \in R$, and $j \in \mathcal{N} \setminus N$, $\succcurlyeq_j \in \mathcal{R}$

$$R - i := R \setminus \{(i, \succcurlyeq_i)\}, \quad R + (j, \succcurlyeq_j) := R \cup \{(j, \succcurlyeq_j)\}.$$

If the identity of the voter is clear or irrelevant we sometimes, in slight abuse of notation, refer to $R - i$ by $R - \succcurlyeq_i$, and write $R + \succcurlyeq_j$ instead of $R + (j, \succcurlyeq_j)$. If k voters with the same preferences $\succcurlyeq_i$ are to be added or removed, we write $R + k \cdot \succcurlyeq_i$ and $R - k \cdot \succcurlyeq_i$, respectively.

The *majority margin* of R is the map $g_R \colon A \times A \to \mathbb{Z}$ with $g_R(\mathtt{x}, \mathtt{y}) = |\{i \in N \mid \mathtt{x} \succcurlyeq_i \mathtt{y}\}| - |\{i \in N \mid \mathtt{y} \succcurlyeq_i \mathtt{x}\}|$. The majority margin can be viewed as the adjacency matrix of a *weighted tournament* T_R. We say that a preference profile R *induces* the weighted tournament T_R.

An alternative $\mathtt{x}$ is called the *Condorcet winner* of R if it wins against any other alternative in a majority contest, i.e., if $g_R(\mathtt{x}, \mathtt{y}) > 0$ for all $\mathtt{y} \in A \setminus \{\mathtt{x}\}$. Conversely, an alternative $\mathtt{x}$ is the *Condorcet loser* if $g_R(\mathtt{x}, \mathtt{y}) < 0$ for all $\mathtt{y} \in A \setminus \{\mathtt{x}\}$.

Our central object of study are *voting rules*, i.e., functions that assign every preference profile a socially preferred alternative. Thus, a voting rule is a function $f \colon \mathcal{R}^{\mathcal{E}(\mathcal{N})} \to A$. Note that this definition requires f to be *resolute*, meaning that f returns exactly one winner for each profile R.

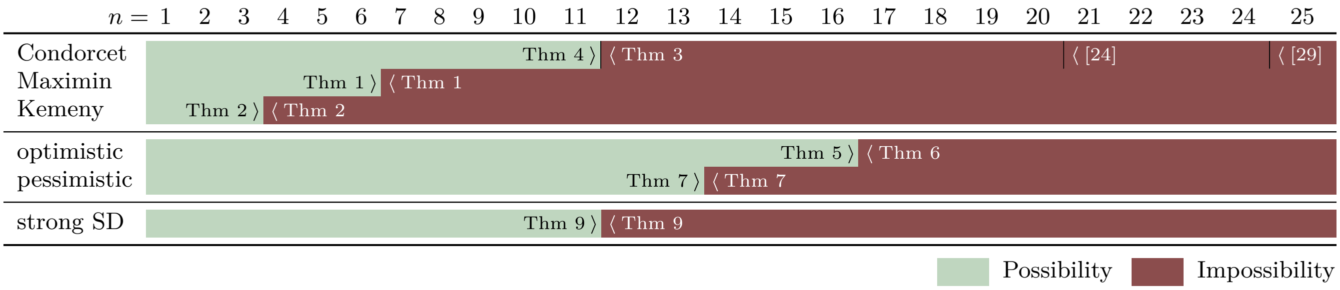

Table 1: Bounds on the number of voters for which Condorcet extensions can satisfy participation. Light green cells indicate the existence of a Condorcet extension satisfying participation (for $m = 4$). Dark red cells indicate that no Condorcet extension satisfies participation (for $m \geqslant 4$).

In this paper, we study voting rules that satisfy *Condorcet-consistency* and *participation*.

Definition 1. *A* Condorcet extension *is a voting rule that selects the Condorcet winner whenever it exists. Thus, f is a Condorcet extension if for every preference profile R that admits a Condorcet winner $\mathtt{x}$, we have $f(R) = \mathtt{x}$. We say that f is* Condorcet-consistent*.*

Definition 2. *A voting rule f satisfies* participation *if all voters always weakly prefer voting to not voting, i.e., if $f(R) \succcurlyeq_i f(R-i)$ for all $R \in \mathcal{R}^N$ and $i \in N$ with $N \in \mathcal{E}(\mathcal{N})$.*

Equivalently, participation requires that for all preference profiles R not including voter j, we have $f(R+\succcurlyeq_j) \succcurlyeq_j f(R)$.

4. MAXIMIN AND KEMENY'S RULE

The proofs of both positive and negative results in this paper were obtained through automated techniques that we describe in Section 5. To become familiar with the kind of arguments produced in this way, we will now study a more restricted setting which is of independent interest.

Specifically, let us consider voting rules that select winners in accordance with the popular *maximin* and *Kemeny* rules. For a preference profile R, an alternative $\mathtt{x}$ is a *maximin winner* if it maximizes $\min_{\mathtt{y} \in A \setminus \{\mathtt{x}\}} g_R(\mathtt{x},\mathtt{y})$; thus, $\mathtt{x}$ never gets defeated too badly in pairwise comparisons. An alternative $\mathtt{x}$ is a *Kemeny winner* if it is ranked first in some *Kemeny ranking*. A Kemeny ranking is a preference relation $\succcurlyeq_{\mathrm{K}} \in \mathcal{R}$ maximizing agreement with voters' individual preferences, i.e., it maximizes the quantity $\sum_{i \in N} |\succcurlyeq_{\mathrm{K}} \cap \succcurlyeq_i|$.

We call a (resolute) voting rule a *maximin extension* (resp. *Kemeny extension*) if it always selects a maximin winner (resp. Kemeny winner). Whenever there are multiple maximin (resp. Kemeny) winners in a profile, such an extension thus needs to break these ties and select exactly one of the winners. Since a Condorcet winner, if it exists, is always the unique maximin and Kemeny winner of a preference profile, any such voting rule is also a Condorcet extension. We can now prove an easy version of Moulin's theorem for these more restricted voting rules.

To this end, we first prove a useful lemma allowing us to extend impossibility proofs for 4 alternatives to also apply if there are more than 4 alternatives. Its proof gives a first hint on how Condorcet-consistency and participation interact.

Lemma 1. *Suppose that f is a Condorcet extension satisfying participation. Let R be a preference profile and $B \subsetneq A$ a set of* bad *alternatives such that each voter ranks every $\mathtt{x} \in B$ below every $\mathtt{y} \in A \setminus B$. Then $f(R) \notin B$.*

Proof. By induction on the number of voters $|N|$ in R. If R consists of a single voter i, then, since f is a Condorcet extension, $f(R)$ must return i's top choice,[1] which is not bad. If R consists of at least 2 voters, and $i \in N$, then by participation $f(R) \succcurlyeq_i f(R-i)$. If $f(R)$ were bad, then so would be $f(R-i)$, contradicting the inductive hypothesis. $\square$

The following computer-aided proofs, just like the more complicated proofs to follow, can be understood solely by carefully examining the corresponding 'proof diagram'. An arrow such as $R \xrightarrow{\;+\,\mathtt{abcd}\;} R'$ indicates that profile R' is obtained from R by adding a voter $\mathtt{abcd}$, and is read as "if one of the bold green alternatives (here $\mathtt{ab}$) is selected at R, then one of them is selected at R'" (by participation). The 'leaves' in the diagrams are profiles admitting a Condorcet winner, and we then print the weighted tournament associated with this profile (unlabelled arcs have weight 1, arcs not printed have weight 0). The winning alternative is shown in a bold circle. In each case, the winning alternative contradicts what is required by the participation axiom. Note that, in this section, we show (unique) maximin winners and Kemeny rankings in these circles. Elsewhere, of course, we show Condorcet winners in the diagrams.

Theorem 1. *There is no maximin extension that satisfies participation for $m \geqslant 4$ and $n \geqslant 7$. (For $m = 4$ and $n \leqslant 6$, such a maximin extension exists.)*

Proof. Let f be a maximin extension which satisfies participation. Consider the 6-voter profile R given in Figure 1. (The numbers at the top of each column describe how often a preference order appears in R. For example, R contains exactly 2 voters with ranking $\mathtt{bdca}$.)

Suppose $f(R) \in \{\mathtt{a},\mathtt{b}\}$. Adding an $\mathtt{abcd}$ vote leads to a weighted tournament in which alternative $\mathtt{c}$ is the unique maximin winner, so that f must choose $\mathtt{c}$. But this contradicts participation since the added voter would benefit from abstaining the election.

Symmetrically, if $f(R) \in \{\mathtt{c},\mathtt{d}\}$, then adding a $\mathtt{dcba}$ vote leads to a weighted tournament in which $\mathtt{b}$ is the unique maximin winner, again contradicting participation. The symmetry of the argument is due to an automorphism

[1] We are not using the full force of Condorcet-consistency here, and only require a notion of unanimity or faithfulness.

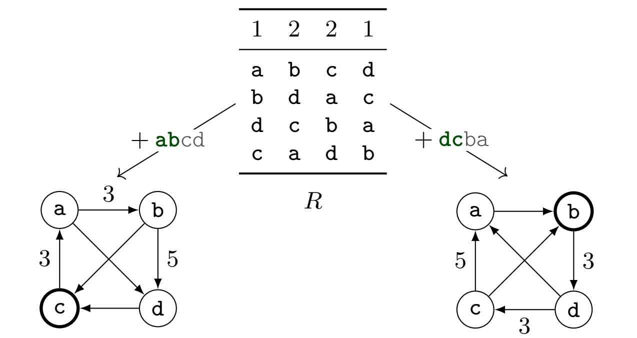

Figure 1: Proof diagram for Theorem 1 (maximin).

of R, namely the relabelling of alternatives according to $\mathtt{abcd} \mapsto \mathtt{dcba}$.

If $m > 4$, we add new 'bad' alternatives $\mathtt{x}_1, \mathtt{x}_2, \dots, \mathtt{x}_{m-4}$ to the bottom of each voter of R and to the bottom of every other voter used in the proof. By Lemma 1, f chooses from $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ at each proof step, and so the proof works as before.

The existence result for $n \leqslant 6$ is established by the methods described in Section 5. $\square$

For 3 alternatives, Moulin [29] proved that the voting rule that chooses the lexicographically first maximin winner satisfies participation. Theorem 1 shows that this is not the case for 4 alternatives, even if there are only 7 voters and no matter how we pick among maximin winners.

Theorem 2. *There is no Kemeny extension that satisfies participation for $m \geqslant 4$ and $n \geqslant 4$. (For $m = 4$ and $n \leqslant 3$, such a Kemeny extension exists.)*

Proof. Let f be a Kemeny extension which satisfies participation. Consider the 4-voter profile R given in Figure 2.

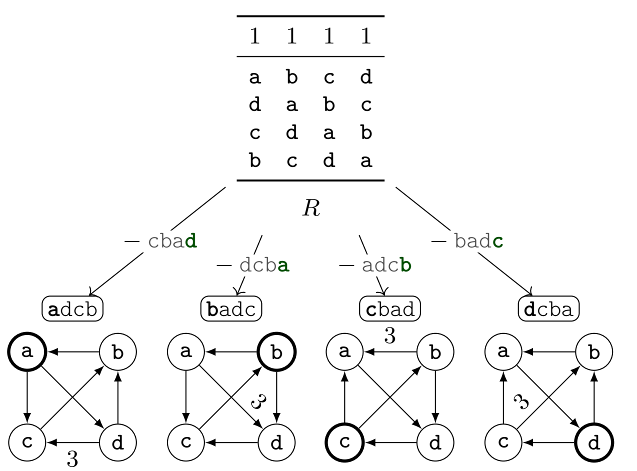

Figure 2: Proof diagram for Theorem 2 (Kemeny). The rankings shown above the tournaments are unique optimal Kemeny rankings.

Suppose $f(R) = \mathtt{d}$. Then removing cbad from R yields a weighted tournament in which adcb is the unique Kemeny ranking (with score 12), and thus a is the unique Kemeny winner, which contradicts participation. Analogously, we can exclude the other three possible outcomes for R by letting a voter abstain, which always leads to a unique Kemeny winner and a contradiction with participation. The arguments are identical because R is completely symmetric in the sense that for any pair of alternatives x and y, there is an automorphism of R that maps x to y.

Just like for Theorem 1, if $m > 4$, we add new bad alternatives $\mathtt{x}_1, \mathtt{x}_2, \dots, \mathtt{x}_{m-4}$ to the bottom of R and of the additional voters. By Lemma 1, f chooses from $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ at each step, completing the proof. $\square$

One remarkable and unexpected aspect of the computer-aided proofs above is that their simplicity is due to automorphisms of the underlying preference profiles. Similar automorphisms will also be used in the proofs of the stronger theorems in Sections 6, 7, and 8. We emphasize that these symmetries are not hard-coded into our problem specification and, to the best of our knowledge, have not been exploited in previous proofs of similar statements.

5. METHOD: SAT SOLVING FOR COMPUTER-AIDED PROOFS

The bounds in this paper were obtained with the aid of a computer. In this section, we describe the method that we employed. The main tool in our approach is an encoding of our problems into propositional logic. We then use SAT solvers to decide whether (in a chosen setting) there exists a Condorcet extension satisfying participation. If the answer is yes, the solver returns an explicit voting rule with the desired properties. If the answer is no, we use a process called *MUS extraction* to find a short certificate of this fact which can be translated into a human-readable proof. By successively proving stronger theorems and using the insights obtained through MUS extraction, we arrived at results as presented in their full generality in this paper.

5.1 SAT Encoding

"For n voters and 4 alternatives, is there a voting rule f that satisfies Condorcet-consistency and participation?"

A natural encoding of this question into propositional logic proceeds like this: Generate all profiles over 4 alternatives with at most n voters. For each such profile R, introduce 4 propositional variables $x_{R,\mathtt{a}}$, $x_{R,\mathtt{b}}$, $x_{R,\mathtt{c}}$, $x_{R,\mathtt{d}}$, where the intended meaning of $x_{R,\mathtt{a}}$ is

$$x_{R,\mathtt{a}} \text{ is set true} \iff f(R) = \mathtt{a}.$$

We then add clauses requiring that for each profile R, $f(R)$ takes *exactly one* value, and we add clauses requiring f to be Condorcet-consistent and satisfy participation.

Sadly, the encoding sketched above is not tractable for the values of n that we are interested in: the number of variables and clauses used grows as $\Theta(24^n)$, because there are $4! = 24$ possible preference relations over 4 alternatives and thus 24^n profiles with n voters. For $n = 7$, this leads to more than 400 billion variables, and for $n = 15$ we exceed 10^{22} variables.

To escape this combinatorial explosion, we will temporarily restrict our attention to *pairwise* voting rules. This means that we assign an outcome alternative $f(T)$ to every weighted tournament T. We then define a voting rule that assigns the outcome $f(T_R)$ to each preference profile R, where T_R is the weighted tournament induced by R.

The number of tournaments induced by profiles with n voters grows much slower than the number of profiles—our computer enumeration suggests a growth of order about 1.5^n. This much more manageable (yet still exponential) growth allows us to consider problem instances up to $n \approx 16$ which turns out to be *just* enough.

Other than referring to (weighted) tournaments instead of profiles, our encoding into propositional formulas now proceeds exactly like before. For each tournament T, we introduce the variables $x_{T,\mathtt{a}}$, $x_{T,\mathtt{b}}$, $x_{T,\mathtt{c}}$, $x_{T,\mathtt{d}}$ and define the formulas

$$\mathsf{non\text{-}empty}_T := x_{T,\mathtt{a}} \vee x_{T,\mathtt{b}} \vee x_{T,\mathtt{c}} \vee x_{T,\mathtt{d}}$$

$$\mathsf{mutex}_T := \bigwedge_{\mathtt{x} \neq \mathtt{y}} (\neg x_{T,\mathtt{x}} \vee \neg x_{T,\mathtt{y}})$$

With our intended interpretation of the variables $x_{T,\mathtt{x}}$, all models of $\bigwedge_T \mathsf{non\text{-}empty}_T \wedge \mathsf{mutex}_T$ are functions from tournaments into $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$. (The word mutex abbreviates 'mutual exclusion' and corresponds to the voting rule selecting a unique winner.)

Since we are interested in voting rules that satisfy participation, we also need to encode this property. To this end, let $T = T_R$ be a tournament induced by R and let $\succcurlyeq$ be a preference relation. Define $T + \succcurlyeq := T_{R+\succcurlyeq}$. (The tournament $T + \succcurlyeq$ is independent of the choice of R.) We define

$$\mathsf{participation}_{T,\succcurlyeq} := \bigwedge_{\mathtt{x}} \left(x_{T,\mathtt{x}} \rightarrow \bigvee_{\mathtt{y} \succcurlyeq \mathtt{x}} x_{T+\succcurlyeq,\mathtt{y}} \right).$$

Requiring f to be Condorcet-consistent is straightforward: if tournament T admits $\mathtt{b}$ as the Condorcet winner, we add

$$\mathsf{condorcet}_T := \neg x_{T,\mathtt{a}} \wedge x_{T,\mathtt{b}} \wedge \neg x_{T,\mathtt{c}} \wedge \neg x_{T,\mathtt{d}},$$

and we add similar formulas for each tournament that admits a Condorcet winner. Then the models of the conjunction of all the $\mathsf{non\text{-}empty}$, mutex, $\mathsf{participation}$, and $\mathsf{condorcet}$ formulas are precisely the pairwise voting rules satisfying Condorcet-consistency and participation.

By adapting the $\mathsf{condorcet}$ formulas, we can impose more stringent conditions on f—this is how our results for maximin and Kemeny extensions are obtained. We can also use this to exclude Pareto-dominated alternatives, and to require f to always pick from the top cycle.

For some purposes it will be useful not to include the mutex clauses in our final formula. Models of this formula then correspond to *set-valued* voting rules that satisfy participation interpreted according to the optimistic preference extension. See Section 7 for results in this setting.

5.2 SAT Solving and MUS Extraction

The formulas we have obtained above are all given in *conjunctive normal form* (CNF), and thus can be evaluated without further transformations by any off-the-shelf SAT solver. In order to physically produce a CNF formula as described, we employ a straightforward Python script that performs a breadth-first search[2] to discover all weighted tournaments with up to n voters (see Algorithm 1 for a schematic overview of the program). The script outputs a CNF formula in the standard DIMACS format, and also outputs a file that, for each variable $x_{T,\mathtt{x}}$, records the tournament T and alternative $\mathtt{x}$ it denotes. This is necessary because the DIMACS format uses uninformative variable descriptors (consecutive integers) and memorizing variable meanings allows us to interpret the output of the SAT solver.

Algorithm 1 Generate formula for up to n voters

$T_0 \leftarrow$ {weighted tournament on $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ with all edges having weight 0}.
for $k = 1, \dots, n$ **do**
  $T_k \leftarrow \emptyset$
  for $T \in T_{k-1}$ **do**
    for $\succcurlyeq \in \mathcal{R}$ **do**
      Calculate $T' := T + \succcurlyeq$
      if T' has not been seen previously **then**
        Add T' to T_k
        Write $\mathsf{non\text{-}empty}_{T'}, \mathsf{mutex}_{T'}, \mathsf{condorcet}_{T'}$
      Write $\mathsf{participation}_{T,\succcurlyeq}$

As an example, the output formula for $n = 15$ in DIMACS format has a size of about 7 GB and uses 50 million variables and 2 billion clauses, taking 6.5 hours to write. Plingeling [2], a popular SAT solver that we used for all results in this paper, solves this formula in 50 minutes of wall clock time, half of which is spent parsing the formula.

In case a given instance is satisfiable, the solver returns a satisfying assignment, giving us an existence proof and a concrete example for a voting rule satisfying participation (and any further requirements imposed). In case a given instance in unsatisfiable, we would like to have short certificate of this fact as well. One possibility for this is having the SAT solver output a resolution proof (in DRUP format, say). This yields a machine-checkable proof, but has two major drawbacks: the generated proofs can be uncomfortably large,[3] and they do not give human-accessible insights about *why* the formula is unsatisfiable.

We handle this problem by computing a *minimal unsatisfiable subset (MUS)* of the unsatisfiable CNF formula. An MUS is a subset of the clauses of the original formula which itself is unsatisfiable, and is minimally so: removing any clause from it yields a satisfiable formula. We used the tools MUSer2 [1] and MARCO [28] to extract MUSes. If an unsatisfiable formula admits a very small MUS, it is often possible to obtain a human-readable proof of unsatisfiability from it [7, 3].

Note that for purposes of extracting human-readable proofs, it is desirable for the MUS to be as small as possible, and also to refer to as few different tournaments as possible. The first issue can be addressed by running the MUS extractor repeatedly, instructing it to order clauses randomly (note that clause sets of different cardinalities can be minimally unsatisfiable with respect to set inclusion); similarly, we can use tools like MARCO to enumerate all MUSes and look for small ones. The second issue can be addressed by computing a *group MUS*: here, we partition the clauses of the CNF formula into *groups*; then, we look for a minimal set of groups that together are unsatisfiable. In our case, the clauses referring to a given tournament T form a group.

[2]The large memory requirements of a BFS were the largest bottleneck in our approach. Sadly, a DFS does not work in our context.

[3]For example, there is a 13 GB large DRUP proof of a special case of the Erdős Discrepancy Conjecture [25], and a solution to the Boolean Pythagorean Triples problem that takes 200 TB in DRAT format [21].

In practice, finding a group MUS first and then finding a standard (clause-level) MUS within the group MUS yielded sets of size much smaller than MUSes returned without the intermediate group-step (often by a factor of 10).

To translate an MUS into a human-readable proof, we created another program that visualizes the MUS in a convenient form.[4] Indeed, this program outputs the 'proof diagrams' like Figure 3 that appear throughout this paper (though we re-typeset them). We think that interpreting these diagrams is quite natural (and is perhaps even easier than reading a textual translation). More importantly, the automatically produced graphs allowed us to quickly judge the quality of an extracted MUS.

5.3 Incremental Proof Discovery

The SAT encoding described in Section 5.1 only concerns pairwise voting rules, yet none of the (negative) results in this paper require or use this assumption. This is the product of multiple rounds of generating and evaluating SAT formulas, extracting MUSes, and using the insights generated by this as 'educated guesses' to solve more general problems.

Following the process as described so far led to a proof that for 4 alternatives and 12 voters, there is no pairwise Condorcet extension that satisfies participation. That proof used the assumption of pairwiseness, i.e., it assumed that the voting rule returns the same alternative on profiles inducing the same weighted tournament. However, intriguingly, the preference profiles mentioned in the proof did not contain all $4! = 24$ possible preference relations over $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$. In fact, the proof only used 10 of the possible orders. Further, each profile appearing in the proof included $R_0 = \{\mathtt{abdc}, \mathtt{bdca}, \mathtt{cabd}, \mathtt{dcab}\}$ as a subprofile. As we argued at the start of Section 5.1, it is intractable to search over the entire space of preference profiles. On the other hand, it is much easier to merely search over all extensions of R_0 that contain at most $n = 12$ voters and only contain copies of the 10 orders previously identified. The SAT formula produced by doing exactly this turned out to be unsatisfiable, and a small MUS extracted from it gave rise to Theorem 3 below.

The proof of Theorem 6 for 17 voters was obtained by running Algorithm 1 with T_0 initialized to the weighted tournament induced by the initial profile R used in the proof of Theorem 3. Before finding this tournament, we tried various other tournaments as T_0, including ones featuring in Moulin's original proof, and ones occurring at other steps in the proof of Theorem 3, but R turned out to give the best (and indeed a tight) bound, and additionally exhibits a lot of symmetry that was also present in the MUS we extracted.

6. MAIN RESULT

We are now in a position to state and prove our main claim that Condorcet extensions cannot avoid the no-show paradox for 12 or more voters (when there are at least 4 alternatives), and that this result is optimal.

[4]Roughly, the visualization program proceeds as follows: an MUS contains two types of clauses. Some (like $\mathsf{non\text{-}empty}, \mathsf{mutex}, \mathsf{condorcet}$) refer to a single tournament; we draw a vertex for each such tournament mentioned in the MUS. On the other hand $\mathsf{participation}_{T,\succcurlyeq}$ clauses *connect* the outcomes at two tournaments together; we draw an arc between them for each $\mathsf{participation}$ clause occurring in the MUS. We also annotate vertices with their Condorcet winners if $\mathsf{condorcet}$ clauses are mentioned.

Theorem 3. *There is no Condorcet extension that satisfies participation for $m \geqslant 4$ and $n \geqslant 12$.*

Proof. We first consider the case $m = 4$. The proof follows the structure depicted in Figure 3. Let R be the preference profile shown there.

Since R remains fixed after relabelling alternatives according to $\mathtt{abcd} \mapsto \mathtt{dcba}$ and reordering voters, we may assume without loss of generality that $f(R) \in \{\mathtt{a}, \mathtt{b}\}$. (An explicit proof in case $f(R) \in \{\mathtt{c}, \mathtt{d}\}$ is indicated in Figure 3.)

By participation, it follows from $f(R) \in \{\mathtt{a}, \mathtt{b}\}$ that also $f(R_\alpha := R + 2 \cdot \mathtt{abcd}) \in \{\mathtt{a}, \mathtt{b}\}$ since the voters with preferences $\mathtt{abcd}$ cannot be worse off by joining the electorate. If $f(R_\alpha) = \mathtt{a}$, again by participation, removing 2 voters with preferences $\mathtt{bdca}$ does not change the winning alternative (so $f(R_\alpha - 2 \cdot \mathtt{bdca}) = \mathtt{a}$), and neither does adding $\mathtt{acdb}$, so $f(R_\alpha - 2 \cdot \mathtt{bdca} + \mathtt{acdb}) = \mathtt{a}$, which, however, is in conflict with $R_\alpha - 2 \cdot \mathtt{bdca} + \mathtt{acdb}$ having a Condorcet winner, $\mathtt{c}$.

Thus we must have $f(R_\alpha) = \mathtt{b}$, which implies that $f(R_\alpha - \mathtt{dcab}) = \mathtt{b}$, and thus $f(R_\beta := R_\alpha - \mathtt{dcab} - 2 \cdot \mathtt{cabd}) \in \{\mathtt{b}, \mathtt{d}\}$.

We again proceed by cases: If $f(R_\beta) = \mathtt{b}$, we can add 2 voters $\mathtt{badc}$ to obtain a profile with Condorcet winner $\mathtt{a}$, which contradicts participation. But then, if $f(R_\beta) = \mathtt{d}$, we get that $f(R_\beta - \mathtt{abcd}) = \mathtt{d}$ and, by another application of participation, that $f(R_\beta - \mathtt{abcd} + 3 \cdot \mathtt{dcba}) = \mathtt{d}$ in contrast to the existence of Condorcet winner $\mathtt{b}$, a contradiction.

If $m > 4$, we add bad alternatives $\mathtt{x}_1, \mathtt{x}_2, \dots, \mathtt{x}_{m-4}$ to the bottom of R and all other voters. By Lemma 1, f chooses from $\{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}\}$ at each step, allowing the proof to go through as for the case $m = 4$. $\square$

The following result establishes that our bound on the number of voters is tight. A very useful feature of our computer-aided approach is that we can easily add additional requirements for the desired voting rule. Two common requirements for voting rules are that they should only return alternatives that are *Pareto-optimal* and contained in the *top cycle* (also known as the *Smith set*) (see, e.g., [17]).

Theorem 4. *There is a Condorcet extension f that satisfies participation for $m = 4$ and $n \leqslant 11$. Moreover, f is pairwise, Pareto-optimal, and a refinement of the top cycle.*

The Condorcet extension f is given as a look-up table, which is derived from the output of a SAT solver. The look-up table lists all 1,204,215 weighted tournaments inducible by up to 11 voters and assigns each an output alternative (see Figure 4 for an excerpt). The relevant text file has a size of 28 MB (gzipped 4.5 MB) and is publicly available in the Harvard Dataverse, together with a Python script verifying that it describes a voting rule that satisfies participation [9].

Comparing this voting rule with known voting rules, it turns out that it selects one of the maximin winners in 99.8% and one of the Kemeny winners in 98% of all weighted tournaments. Note that there can be multiple maximin and Kemeny winners in a given profile. However, the rule still agrees with the maximin rule with lexicographic tie-breaking on 95% of weighted tournaments. The similarity with the maximin rule is interesting insofar as a well-documented flaw of the maximin rule is that it fails to be a refinement of the top cycle (and may even return Condorcet losers). Our computer-generated rule always picks from the top cycle while it remains very close to the maximin rule.

80% of the considered weighted tournaments admit a Condorcet winner, which uniquely determines the output of the rule; this can be used to reduce the size of the look-up table.

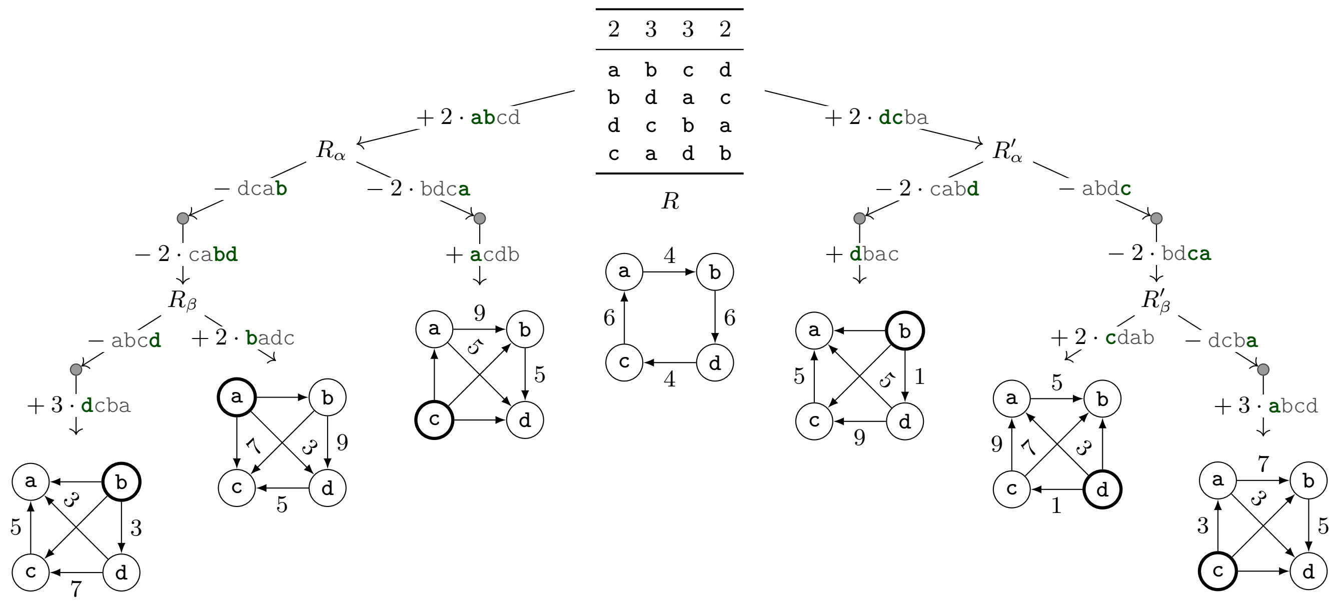

Figure 3: Computer-aided proof of Theorem 3 in graphical form, showing that there is no Condorcet extension that satisfies participation for $m \geqslant 4$ and $n \geqslant 12$. See Section 4 for an explanation of how to read this diagram.

```
a,#1,(1,1,1,1,1,1)
a,#1,(1,1,1,1,1,-1)
a,#1,(1,1,1,-1,1,1)
a,#1,(1,1,1,-1,-1,1)
a,#1,(1,1,1,1,-1,-1)
a,#1,(1,1,1,-1,-1,-1)
b,#1,(-1,1,1,1,1,1)
b,#1,(-1,1,1,1,1,-1)
b,#1,(-1,-1,1,1,1,1)
b,#1,(-1,-1,-1,1,1,1)
b,#1,(-1,1,-1,1,1,-1)
b,#1,(-1,-1,-1,1,1,-1)
c,#1,(1,-1,1,-1,1,1)
c,#1,(1,-1,1,-1,-1,1)
a,#11,(9,11,3,9,1,-9)
a,#11,(11,9,3,7,1,-9)
c,#11,(5,-9,-1,-11,-1,7)
c,#11,(5,-9,-1,-11,-1,5)
c,#11,(3,-11,-1,-9,1,7)
c,#11,(3,-11,-3,-9,1,7)
c,#11,(3,-11,-3,-11,-1,7)
b,#11,(-1,3,-5,-3,5,-3)
b,#11,(-3,3,-7,-3,5,-3)
b,#11,(-3,1,-7,-3,5,-3)
c,#11,(-3,1,-5,-5,5,-1)
a,#11,(3,7,11,-3,9,11)
a,#11,(3,7,11,-3,9,9)
a,#11,(3,7,11,-5,9,11)
```

Figure 4: Excerpt of look-up table giving a pairwise Condorcet extension satisfying participation for $n \leqslant 11$ voters (from Theorem 4). Each row lists a weighted tournament as $(g_R(\mathtt{a},\mathtt{b}), g_R(\mathtt{a},\mathtt{c}), g_R(\mathtt{a},\mathtt{d}), g_R(\mathtt{b},\mathtt{c}), g_R(\mathtt{b},\mathtt{d}), g_R(\mathtt{c},\mathtt{d}))$ with a chosen alternative, and with the number of voters inducing the tournament.

7. SET-VALUED VOTING RULES

A drawback of voting rules, as we defined them so far, is that the requirement to always return a single winner is in conflict with basic fairness conditions, namely anonymity and neutrality. A large part of the social choice literature therefore deals with set-valued voting rules, where ties are usually assumed to be eventually broken by some tie-breaking mechanism.

A *set-valued voting rule* (sometimes known as a voting *correspondence* or as an *irresolute* voting rule) is a function $F\colon \mathcal{R}^{\mathcal{E}(\mathcal{N})} \to 2^A \setminus \{\emptyset\}$ that assigns each preference profile R a non-empty set of alternatives. The function F is a *(set-valued) Condorcet extension* if for every preference profile R that admits a Condorcet winner $\mathtt{x}$, we have $F(R) = \{\mathtt{x}\}$.

Following the work of Pérez [30] and Jimeno et al. [23], we seek to study the occurrence of the no-show paradox in this setting. To do so, we need to define appropriate notions of participation, and for this we will need to specify agents' preferences over *sets* of alternatives. Here, we use the *optimistic* and *pessimistic preference extensions.* An optimist prefers sets with better most-preferred alternative, while a pessimist prefers sets with better least-preferred alternative. For example, if $U = \{\mathtt{a},\mathtt{b},\mathtt{d}\}$ and $V = \{\mathtt{b},\mathtt{c}\}$, then an optimist with preferences abcd prefers U to V, while a pessimist prefers V to U. With these notions, we extend the participation property to set-valued voting rules.

Given a set $U \subseteq A$, we write $\max_{\succcurlyeq_i} U$ (resp. $\min_{\succcurlyeq_i} U$) for the most-preferred (resp. least-preferred) alternative in U according to $\succcurlyeq_i$, so that for all $\mathtt{x} \in U$ we have $\max_{\succcurlyeq_i} U \succcurlyeq_i \mathtt{x} \succcurlyeq_i \min_{\succcurlyeq_i} U$.

Definition 3. *A set-valued voting rule F satisfies* optimistic participation *if* $\max_{\succcurlyeq_i} F(R + \succcurlyeq_i) \succcurlyeq_i \max_{\succcurlyeq_i} F(R)$.
A set-valued voting rule F satisfies pessimistic participation *if* $\min_{\succcurlyeq_i} F(R) \succcurlyeq_i \min_{\succcurlyeq_i} F(R - i)$.

One may wonder why we defined optimistic participation in terms of *adding* voters, and pessimistic participation in terms of *subtracting* voters, since the two operations give equivalent definitions. However, it will turn out to be useful to associate optimistic participation with "+", and pessimistic participation with "−". This is because the + and − operations as used in a proof diagram are similarly associated with these set extension: Suppose F satisfies optimistic participation, R is a profile, and the only thing we know about $F(R)$ is that $\mathtt{a} \in F(R)$. We can then invoke optimistic participation for voter i to deduce that there exists some $\mathtt{x}$ with $\mathtt{x} \in F(R + \succcurlyeq_i)$ and $\mathtt{x} \succcurlyeq_i \mathtt{a}$ (because optimist i must weakly prefer joining). On the other hand, given only this knowledge, we cannot use optimistic participation to similarly deduce anything about the elements of $F(R-j)$; for such a deduction we would need to know j's *most-preferred* element in $F(R)$, which we do not know. There is a similar connection between pessimistic participation and subtracting voters: if F' satisfies pessimistic participation and we know only that $\mathtt{a} \in F'(R)$, then we can deduce that there is some $\mathtt{y} \in F'(R-j)$ with $\mathtt{a} \succcurlyeq_i \mathtt{y}$ (because pessimist j must weakly dislike abstaining); but no similar deductions are possible for $F(R + \succcurlyeq_i)$.

With these observations, we can anticipate that proof diagrams about impossibility results for optimistic [pessimistic] participation only use additions [subtractions] of voters.

Before we go on to our results, we need an analogue of Lemma 1 for the set-valued notions of participation we have defined.

Lemma 2. *Suppose that F is a set-valued Condorcet extension. Let R be a preference profile and $B \subsetneq A$ a set of* bad *alternatives such that each voter ranks every $\mathtt{x} \in B$ below every $\mathtt{y} \in A \setminus B$.*

If F satisfies optimistic participation, then $F(R)$ contains a non-bad alternative: $F(R) \not\subseteq B$.

If F satisfies pessimistic participation, then $F(R)$ does not contain any bad alternatives: $F(R) \cap B = \emptyset$.

Proof. *Optimistic participation.* By induction on the number of voters $|N|$ in R. If R consists of a single voter i, then, since F is a Condorcet extension, $F(R) = \{\mathtt{x}\}$, where $\mathtt{x}$ is i's top choice, which is not bad. If R consists of at least 2 voters, and $i \in N$, then by optimistic participation $\max_{\succcurlyeq_i} F(R) \succcurlyeq_i \max_{\succcurlyeq_i} F(R-i)$ (since $R = R - i + \succcurlyeq_i$). By inductive hypothesis, $F(R-i)$ contains a non-bad alternative $\mathtt{x}$, and so $\max_{\succcurlyeq_i} F(R-i)$ is not bad. Thus, $\max_{\succcurlyeq_i} F(R)$ is not bad either, and so $F(R)$ contains a non-bad alternative.

Pessimistic participation. By induction on the number of voters $|N|$ in R. If R consists of a single voter i, then, since F is a Condorcet extension, $F(R) = \{\mathtt{x}\}$, where $\mathtt{x}$ is i's top choice, which is not bad. If R consists of at least 2 voters, and $i \in N$, then by pessimistic participation $\min_{\succcurlyeq_i} F(R) \succcurlyeq_i \min_{\succcurlyeq_i} F(R-i)$. By inductive hypothesis, $F(R-i)$ does not contain a bad alternative, so $\min_{\succcurlyeq_i} F(R-i)$ is not bad. Thus, $\min_{\succcurlyeq_i} F(R)$ is not bad either, and so by definition of min, we see that $F(R)$ contains no bad alternatives. □

A set-valued voting rule F is called *resolute* if it always selects a single alternative, so that for all R we have $|F(R)| = 1$. A (single-valued) voting rule f is naturally identified with a resolute set-valued voting rule F; if f satisfies participation, then this F satisfies both optimistic and pessimistic participation. Hence, by Theorem 4, there is a (resolute) set-valued Condorcet extension F that satisfies both optimistic and pessimistic participation for up to 4 alternatives and 11 voters. However, there might be hope that allowing voting rules to be irresolute also allows for participation to be attainable for more voters, and indeed this is the case.

Theorem 5. *There is a set-valued Condorcet extension F that satisfies optimistic participation for $m = 4$ and $n \leqslant 16$, and also is Pareto-optimal and a refinement of the top cycle.*

The SAT solver indicates that no such set-valued voting rule is pairwise. The voting rule we found is publicly available [9]. Theorem 5 is optimal in the sense that optimistic participation cannot be achieved if we allow for one more voter.

Theorem 6. *There is no set-valued Condorcet extension that satisfies optimistic participation for $m \geqslant 4$ and $n \geqslant 17$.*

Proof. Let F be such a function, and consider the 10-voter profile R given in Figure 5.

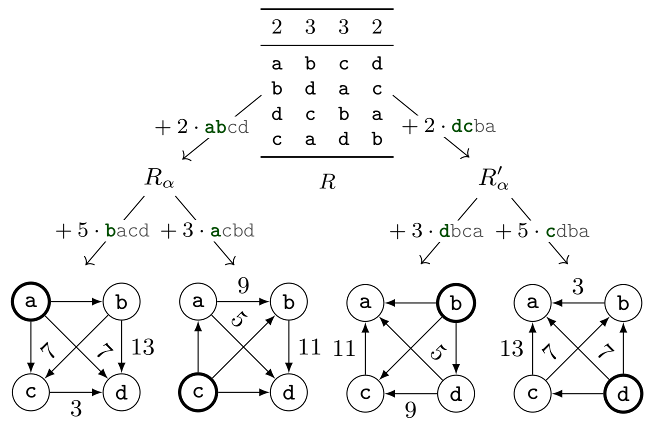

Figure 5: Proof diagram for Theorem 6 (optimist).

Suppose that either $\mathtt{a} \in F(R)$ or $\mathtt{b} \in F(R)$. (The case of $\mathtt{c} \in F(R)$ or $\mathtt{d} \in F(R)$ is symmetric.) Then let $R_\alpha := R + 2 \cdot \mathtt{abcd}$. By optimistic participation, we then have either $\mathtt{a} \in F(R_\alpha)$ or $\mathtt{b} \in F(R_\alpha)$. If we had $\mathtt{a} \in F(R_\alpha)$, then also $\mathtt{a} \in F(R_\alpha + 3 \cdot \mathtt{acbd})$ but this profile has Condorcet winner $\mathtt{c}$, and if $\mathtt{b} \in F(R_\alpha)$ then also $\mathtt{b} \in F(R_\alpha + 5 \cdot \mathtt{bacd})$ but this profile has Condorcet winner $\mathtt{a}$. This is a contradiction.

This argument extends to more than 4 alternatives by appealing to Lemma 2. □

Inspecting Moulin's original proof [29] shows that it also establishes an impossibility for optimistic participation (for 25 voters). This is because Moulin's proof only ever *adds* voters to the profile under consideration. Apparently unaware of this, Jimeno et al. [23] explicitly establish such a result for 27 voters and 5 alternatives. Observe that, as promised, each step of the proof of Theorem 6 involves *adding* voters to the current profile, and we never remove voters. In light of our remarks after Definition 3, this is the reason why the proof establishes a result for optimistic participation. If we restrict ourselves to deleting voters, we obtain a result for pessimistic participation.

Theorem 7. *There is no set-valued Condorcet extension that satisfies pessimistic participation for $m \geqslant 4$ and $n \geqslant 14$. On the other hand, for $m = 4$ and $n \leqslant 13$, there exists such a set-valued voting rule.*

Proof. The proof has a similar structure to the proof of Theorem 3, and is shown in graphical form in Figure 6. The initial profile of this proof is $R^* := R + 2 \cdot \mathtt{abcd} + 2 \cdot \mathtt{dcba}$, taking R to be the profile of Figure 3. We further replace proof steps in which voters are added by similar ones where voters are deleted, and invoke pessimistic participation at each such step to obtain a contradiction.

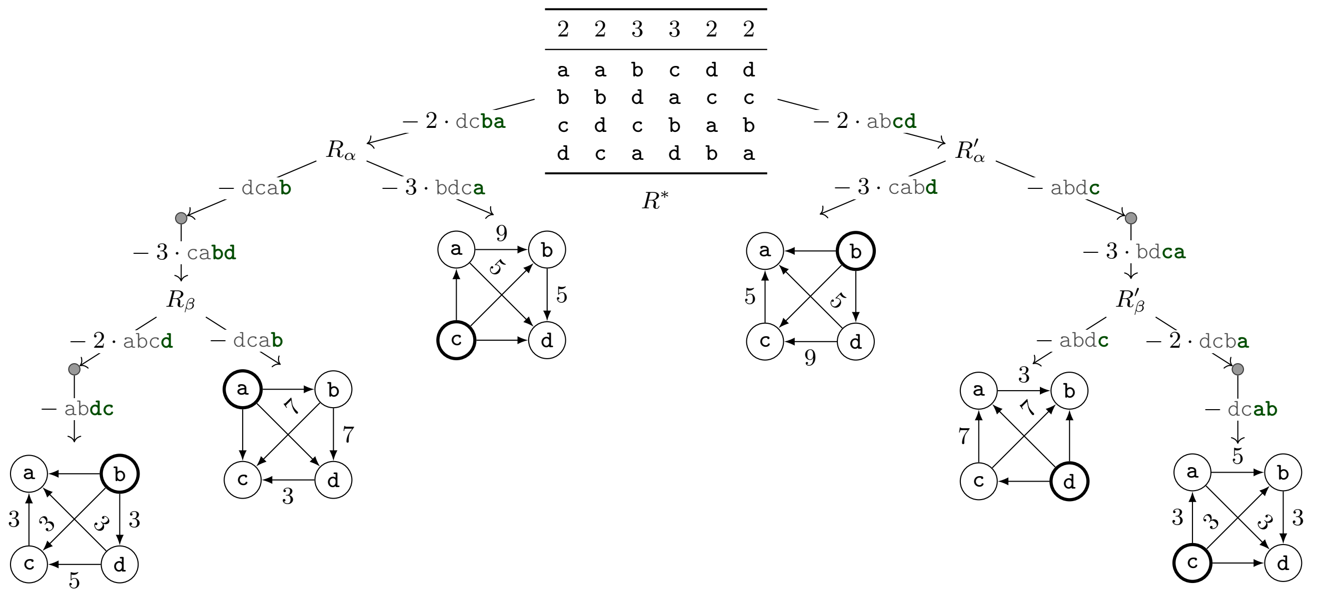

Figure 6: Proof diagram for Theorem 7 (pessimist).

Explicitly, since R^* remains fixed after relabelling alternatives according to $\texttt{abcd} \mapsto \texttt{dcba}$ and reordering voters, we may assume without loss of generality that $\texttt{a} \in f(R^*)$ or $\texttt{b} \in f(R^*)$.

Now let $R_\alpha := R^* - 2\cdot\texttt{dcba}$. By pessimistic participation, it follows that also either $\texttt{a} \in f(R_\alpha)$ or $\texttt{b} \in f(R_\alpha)$ since the voters with preferences abcd should be weakly worse off by leaving the electorate. If $\texttt{a} \in f(R_\alpha)$, again by pessimistic participation, removing 3 voters with preferences bdca still leaves $\texttt{a} \in f(R_\alpha - 3\cdot\texttt{bdca})$, but $R_\alpha - 3\cdot\texttt{bdca}$ has a Condorcet winner, c.

Thus we must have $\texttt{b} \in f(R_\alpha)$, which implies that $\texttt{b} \in f(R_\alpha - \texttt{dcab})$. Let $R_\beta := R_\alpha - \texttt{dcab} - 3\cdot\texttt{cabd}$. Participation implies that either $\texttt{b} \in f(R_\beta)$ or $\texttt{d} \in f(R_\beta)$.

We again proceed by cases: If $\texttt{b} \in f(R_\beta)$, we can remove a voter dcab to obtain a profile with Condorcet winner a, which contradicts participation. But then, if $\texttt{d} \in f(R_\beta)$, we get that $\texttt{d} \in f(R_\beta - 2\cdot\texttt{abcd})$ and, by another application of participation, that $f(R_\beta - 2\cdot\texttt{abcd} - \texttt{abdc})$ contains either c or d, in contrast to the existence of Condorcet winner b, a contradiction. □

This result strengthens a result of Jimeno et al. [23], who show that for $m \geqslant 5$ no set-valued Condorcet extension satisfying a property called "weak translation invariance" can also satisfy pessimistic participation. Our proof does not need the extra assumption, already works for $m = 4$ alternatives, and uses just 14 instead of 971 voters.[5]

As previously observed, adding voters in our impossibility proofs corresponds to optimistic participation, while removing voters corresponds to pessimistic participation. In the proof of Theorem 3, we use both operations, which allows for a tighter bound of just 12 voters.[6] In the set-valued setting, we can formulate this result in a slightly stronger way.[7]

THEOREM 8. *There is no set-valued Condorcet extension that satisfies optimistic and pessimistic participation simultaneously for $m \geqslant 4$ and $n \geqslant 12$. On the other hand, for $m = 4$ and $n \leqslant 11$ such a set-valued rule exists (and also is Pareto-optimal and a refinement of the top cycle).*

PROOF. Use the proof of Theorem 3, invoking optimistic participation for edges labelled with the addition of a voter (+), and invoking pessimistic participation for edges labelled with removal of a voter (−). On the other hand, the (single-valued) voting rule of Theorem 4 clearly satisfies both optimistic and pessimistic participation. □

8. PROBABILISTIC VOTING RULES

A *probabilistic voting rule* (also known as a *social decision scheme*) assigns to each preference profile R a probability distribution (or *lottery*) over A. Thus, a probabilistic voting rule might assign a fair coin flip between a and b as the outcome of an election.

Formally, let $\Delta(A) = \{\mathbf{p} : A \to [0,1] : \sum_{\mathtt{x}\in A} \mathbf{p}(\mathtt{x}) = 1\}$ be the set of lotteries over A; a lottery $\mathbf{p} \in \Delta(A)$ assigns probability $\mathbf{p}(\mathtt{x})$ to alternative x. A probabilistic voting rule f is a function $f : \mathcal{R}^{\mathcal{E}(\mathcal{N})} \to \Delta(A)$. In this context, we say that f is a *Condorcet extension* if $f(R)$ puts probability 1 on the Condorcet winner of R whenever it exists: if R admits x as the Condorcet winner, then $f(R)(\mathtt{x}) = 1$.

As in the set-valued case, we need a notion of comparing outcomes in order to extend the definition of participation. Here, we use the concept of *stochastic dominance (SD).*

DEFINITION 4. *Let $\succcurlyeq\ \in \mathcal{R}$ be a preference relation over A, and let $\mathbf{p}, \mathbf{q} \in \Delta(A)$ be lotteries. Then $\mathbf{p}$ is (weakly)*

[5] The large number of voters is due to several applications of the "weak translation invariance" axiom, each of which adds $5! = 120$ voters to the preference profile under consideration.

[6] The proof of Theorem 6 can thus be seen as a simplified version of the proof of Theorem 3, at the expense of a higher bound on the number of voters.

[7] The notion of participation used in Theorem 8 is reminiscent of the strategyproofness notion used by Duggan and Schwartz [16].

SD-preferred over $\mathbf{q}$ *by* $\succcurlyeq$ *if for each alternative* $\mathtt{x}$*, we have*

$$\textstyle\sum_{\mathtt{y}\succcurlyeq \mathtt{x}} \mathbf{p}(y) \geqslant \sum_{\mathtt{y}\succcurlyeq \mathtt{x}} \mathbf{q}(y).$$

In this case, we write $\mathbf{p} \succcurlyeq^{\mathrm{SD}} \mathbf{q}$.

For example, the lottery $\frac{2}{3}\mathtt{a} + \frac{1}{3}\mathtt{c}$ is SD-preferred to the lottery $\frac{1}{3}\mathtt{a} + \frac{1}{3}\mathtt{b} + \frac{1}{3}\mathtt{c}$ by a voter with preferences $\mathtt{abcd}$. A voter with preferences $\mathtt{bacd}$ will feel the other way around. The main appeal of stochastic dominance stems from the following equivalence: $\mathbf{p} \succcurlyeq^{\mathrm{SD}} \mathbf{q}$ if and only if $\mathbf{p}$ yields at least as much von-Neumann-Morgenstern utility as $\mathbf{q}$ under any utility function that is consistent with the ordinal preferences $\succcurlyeq$. Using this notion of comparing lotteries, we can define participation analogously to the previous settings.

Definition 5. *A probabilistic voting rule f satisfies* strong SD-participation *if* $f(R) \succcurlyeq_i^{\mathrm{SD}} f(R-i)$ *for all* $R \in \mathcal{R}^N$ *and* $i \in N$ *with* $N \in \mathcal{E}(\mathcal{N})$*.*

Any (single-valued) voting rule f can be seen as a probabilistic voting rule that puts probability 1 on its chosen alternative. If f satisfies participation, then this derived probabilistic voting rule is easily seen to satisfy strong SD-participation. Hence Theorem 4 gives us a probabilistic Condorcet extension that satisfies strong SD-participation for $n \leqslant 11$ voters and $m = 4$ alternatives.

We now establish a connection between strong SD-participation and the set-valued notions of participation that we considered in Section 7. This connection will allow us to translate the impossibility results we obtained there to the probabilistic setting. To set up this connection, let us define the *support* of a lottery $\mathbf{p} \in \Delta(A)$ to be $\mathrm{supp}(\mathbf{p}) := \{\mathtt{x} \in A : \mathbf{p}(\mathtt{x}) > 0\}$.

Proposition 1. *Let f be a probabilistic voting rule satisfying strong SD-participation. Let $F = \mathrm{supp} \circ f$ be the support of f, i.e., $F(R) = \mathrm{supp}(f(R))$ for all profiles R. Then F satisfies both optimistic and pessimistic participation.*

Proof. Let R be a preference profile with electorate $N \in \mathcal{E}(\mathcal{N})$, and let $i \in \mathcal{N} \setminus N$ be a voter with preferences $\succcurlyeq_i$.

Optimistic participation. Let $\mathtt{x} = \max_{\succcurlyeq_i} F(R)$. We need to show that $\max_{\succcurlyeq_i} F(R+\succcurlyeq_i) \succcurlyeq_i \mathtt{x}$, by finding an alternative $\mathtt{y}$ with $\mathtt{y} \succcurlyeq_i \mathtt{x}$ that is in the support of $f(R+\succcurlyeq_i)$.

But since f satisfies strong SD-participation, we have

$$\sum_{\mathtt{y}\succcurlyeq_i \mathtt{x}} f(R+\succcurlyeq_i)(\mathtt{y}) \geqslant \sum_{\mathtt{y}\succcurlyeq_i \mathtt{x}} f(R)(\mathtt{y}) > 0,$$

where the strict inequality follows from the definition of the support and of $\mathtt{x}$. Hence some alternative $\mathtt{y}$ with $\mathtt{y} \succcurlyeq_i \mathtt{x}$ is in the support of $f(R+\succcurlyeq_i)$, as required.

Pessimistic participation. Let $\mathtt{z} = \min_{\succcurlyeq_i} F(R)$. We need to show that $\min_{\succcurlyeq_i} F(R+\succcurlyeq_i) \succcurlyeq_i \mathtt{z}$, by showing that no alternative $\mathtt{y}$ with $\mathtt{y} \prec_i \mathtt{z}$ is in the support of $f(R+\succcurlyeq_i)$.

But since f satisfies strong SD-participation, we have

$$0 \leqslant \sum_{\mathtt{y}\prec_i \mathtt{z}} f(R+\succcurlyeq_i)(\mathtt{y}) \leqslant \sum_{\mathtt{y}\prec_i \mathtt{z}} f(R)(\mathtt{y}) = 0,$$

where the last equality follows from the definition of the support and of $\mathtt{z}$. Hence the support of $f(R+\succcurlyeq_i)$ consists solely of alternatives $\mathtt{y}$ with $\mathtt{y} \succcurlyeq_i \mathtt{z}$, as required. $\square$

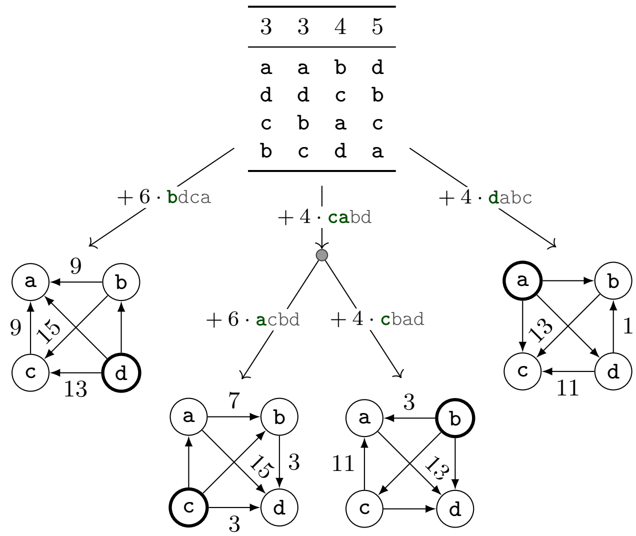

Figure 7: Proof diagram for Moulin's original proof.

Putting these results together with the impossibility result of Theorem 8, we obtain the following.

Theorem 9. *There is no probabilistic Condorcet extension that satisfies strong SD-participation for $n \geqslant 12$ and $m \geqslant 4$. On the other hand, for $m = 4$ and $n \leqslant 11$, such a probabilistic voting rule exists.*

Theorem 9 resolves an open problem mentioned by Brandl et al. [4, Sec. 6].

9. CONCLUSIONS AND FUTURE WORK

We have given tight results delineating in which situations no-show paradoxes must occur. Our results precisely characterize under which conditions Condorcet-consistency and participation are incompatible and thus nicely complement recent advances to satisfy Condorcet-consistency and participation by exploiting uncertainties of the voters about their preferences or about the voting rule's tie-breaking mechanism [3, 4, 5].

The graphical representation of our proofs can also be used to compactly represent Moulin's original proof [29]. To do this, several applications of Moulin's Claim 3 have to be decoded into the explicit votes that are added to the profiles under consideration. This was already done in the expositions of Schulze [33] and Smith [34] and the resulting proof diagram is shown in Figure 7.

Due to unmanageable branching factors when there are 5 alternatives (and hence $5! = 120$ possible preference relations), we were unable to check using our approach whether no-show paradoxes occur with even fewer voters when the number of alternatives grows. We leave this question for the (possibly far) future. Further, it would be interesting to gain a deeper understanding of the computer-generated Condorcet extension that satisfies participation for up to 11 voters. So far, we only know that it (slightly) differs from all Condorcet extensions that are usually considered in the literature. As a first step, it would be desirable to obtain a representation of this rule that is more concise than a look-up table.

Another interesting topic for future research is to find optimal bounds for a variant of the no-show paradox due to Sanver and Zwicker [32], in which participation is weakened to half-way monotonicity.

Acknowledgments

We thank the reviewers for helpful comments that improved the presentation of our results. We are very grateful to Frano Stanisic for pointing out an error in an earlier version of this paper.

Christian Geist is supported by Deutsche Forschungsgemeinschaft under grant BR 2312/9-1. Dominik Peters is supported by EPSRC. Part of this work was conducted while Dominik Peters visited TUM, supported by the COST Action IC1205 on Computational Social Choice.

Results of this paper were presented at the 15th International Conference on Autonomous Agents and Multiagent Systems (Singapore, May 2016), the 6th International Workshop on Computational Social Choice (Toulouse, June 2016), and the 13th Meeting of the Society of Social Choice and Welfare (Lund, July 2016).

REFERENCES

[1] A. Belov and J. Marques-Silva. MUSer2: An efficient MUS extractor. *Journal on Satisfiability, Boolean Modeling and Computation*, 8:123–128, 2012.

[2] A. Biere. Lingeling, Plingeling and Treengeling entering the SAT competition 2013. In *Proceedings of the SAT Competition 2013*, pages 51–52, 2013.

[3] F. Brandl, F. Brandt, C. Geist, and J. Hofbauer. Strategic abstention based on preference extensions: Positive results and computer-generated impossibilities. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 18–24. AAAI Press, 2015.

[4] F. Brandl, F. Brandt, and J. Hofbauer. Incentives for participation and abstention in probabilistic social choice. In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1411–1419. IFAAMAS, 2015.

[5] F. Brandl, F. Brandt, and J. Hofbauer. Welfare maximization entices participation. 2016. Working paper.

[6] F. Brandt. Set-monotonicity implies Kelly-strategyproofness. *Social Choice and Welfare*, 45 (4):793–804, 2015.

[7] F. Brandt and C. Geist. Finding strategyproof social choice functions via SAT solving. *Journal of Artificial Intelligence Research*, 55:565–602, 2016.

[8] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[9] F. Brandt, C. Geist, and D. Peters. Replication data for: Optimal bounds for the no-show paradox via SAT solving. Harvard Dataverse. http://dx.doi.org/10.7910/DVN/TGIQB7, 2016.

[10] D. E. Campbell and J. S. Kelly. Strategy-proofness and weighted voting. *Mathematical Social Sciences*, 60(1): 15–23, 2010.

[11] D. E. Campbell, J. Graver, and J. S. Kelly. There are more strategy-proof procedures than you think. *Mathematical Social Sciences*, 64(3):263–265, 2012.

[12] S. Chatterjee and A. Sen. Automated reasoning in social choice theory — some remarks. *Mathematics in Computer Science*, 8(1):5–10, 2014.

[13] V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 103–110, 2002.

[14] M. de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, 1785. Facsimile published in 1972 by Chelsea Publishing Company, New York.

[15] C. Duddy. Condorcet's principle and the strong no-show paradoxes. *Theory and Decision*, 77(2):275–285, 2014.

[16] J. Duggan and T. Schwartz. Strategic manipulability without resoluteness or shared beliefs: Gibbard-Satterthwaite generalized. *Social Choice and Welfare*, 17(1):85–93, 2000.

[17] P. C. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

[18] P. C. Fishburn and S. J. Brams. Paradoxes of preferential voting. *Mathematics Magazine*, 56(4):207–214, 1983.

[19] A. Fréchette, N. Newman, and K. Leyton-Brown. Solving the station repacking problem. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2016.

[20] C. Geist and U. Endriss. Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research*, 40:143–174, 2011.

[21] M. J. H. Heule, O. Kullmann, and V. W. Marek. Solving and verifying the Boolean Pythagorean triples problem via cube-and-conquer. In *Proceedings of the 19th International Conference on Theory and Applications of Satisfiability Testing*, volume 9710 of *Lecture Notes in Computer Science (LNCS)*, pages 228–245. Springer-Verlag, 2016.

[22] R. Holzman. To vote or not to vote: What is the quota? *Discrete Applied Mathematics*, 22(2):133–141, 1988.

[23] J. L. Jimeno, J. Pérez, and E. García. An extension of the Moulin No Show Paradox for voting correspondences. *Social Choice and Welfare*, 33(3):343–459, 2009.

[24] K. Kardel. Participation and group-participation in social choice theory. Master's thesis, Technische Universität München, 2014.

[25] B. Konev and A. Lisitsa. A SAT attack on the Erdős discrepancy conjecture. In *Theory and Applications of Satisfiability Testing–SAT 2014*, pages 219–226. Springer, 2014.

[26] J.-F. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.

[27] D. Lepelley and V. Merlin. Scoring run-off paradoxes for variable electorates. *Economic Theory*, 17(1):53–80, 2000.

[28] M. H. Liffiton, A. Previti, A. Malik, and J. Marques-Silva. Fast, flexible MUS enumeration. *Constraints*, pages 1–28, 2015.

[29] H. Moulin. Condorcet's principle implies the no show paradox. *Journal of Economic Theory*, 45:53–64, 1988.

[30] J. Pérez. The strong no show paradoxes are a common flaw in Condorcet voting correspondences. *Social Choice and Welfare*, 18(3):601–616, 2001.

[31] D. Ray. On the practical possibility of a 'no show paradox' under the single transferable vote. *Mathematical Social Sciences*, 11(2):183–189, 1986.

[32] M. R. Sanver and W. S. Zwicker. One-way monotonicity as a form of strategy-proofness. *International Journal of Game Theory*, 38(4):553–574, 2009.

[33] M. Schulze. Condorcet and Participation. Election-Methods mailing list. http://lists.electorama.com/pipermail/election-methods-electorama.com/2003-October/011042.html, 2003. (Accessed: 7 Nov 2015).

[34] W. D. Smith. "Participation failure" is forced in Condorcet methods with at least 4 candidates. The Center for Range Voting. http://rangevoting.org/CondPF.html, 2007. (Accessed: 7 Nov 2015).

[35] P. Tang and F. Lin. Computer-aided proofs of Arrow's and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.

[36] D. R. Woodall. Monotonicity of single-seat preferential election rules. *Discrete Applied Mathematics*, 77(1):81–98, 1997.