Condorcet Winners and Anscombe's Paradox Under Weighted Binary Voting

Carmel Baharav
ETH Zürich
Zürich, Switzerland
cbaharav@ethz.ch

Andrei Constantinescu
ETH Zürich
Zürich, Switzerland
aconstantine@ethz.ch

Roger Wattenhofer
ETH Zürich
Zürich, Switzerland
wattenhofer@ethz.ch

ABSTRACT

We consider voting on multiple independent binary issues. In addition, a weighting vector for each voter defines how important they consider each issue. The most natural way to aggregate the votes into a single unified proposal is *issue-wise majority* (IWM): taking a majority opinion for each issue. However, in a scenario known as *Ostrogorski's Paradox*, an IWM proposal may not be a Condorcet winner, or it may even fail to garner majority support in a special case known as *Anscombe's Paradox*.

We show that it is co-NP-hard to determine whether there exists a Condorcet-winning proposal even without weights. In contrast, we prove that the *single-switch* condition provides an Ostrogorski-free voting domain under identical weighting vectors. We show that verifying the condition can be achieved in linear time and no-instances admit short, efficiently computable proofs in the form of forbidden substructures. On the way, we give the simplest linear-time test for the *voter/candidate-extremal-interval* condition in approval voting and the simplest and most efficient algorithm for recognizing single-crossing preferences in ordinal voting.

We then tackle Anscombe's Paradox. Under identical weight vectors, we can guarantee a majority-supported proposal agreeing with IWM on strictly more than half of the overall weight, while with two distinct weight vectors, such proposals can get arbitrarily far from IWM. The severity of such examples is controlled by the maximum average topic weight $\tilde{w}_{max}$: a simple bound derived from a partition-based approach is tight on a large portion of the range $\tilde{w}_{max} \in (0, 1)$. Finally, we extend Wagner's rule to the weighted setting: an average majority across topics of at least $\frac{3}{4}$'s precludes Anscombe's paradox from occurring.

KEYWORDS

Condorcet; Anscombe; Ostrogorski; Multiple Referenda; Complexity; Forbidden Substructures; Restricted Domains; Single-Crossing

ACM Reference Format:
Carmel Baharav, Andrei Constantinescu, and Roger Wattenhofer. 2025. Condorcet Winners and Anscombe's Paradox Under Weighted Binary Voting. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

There are numerous scenarios in which people must decide on a slate of binary issues and come up with a single outcome for each topic. When political parties form a platform, they must aggregate their base's opinions and provide a unified set of stances on numerous separate issues. Similarly, when voters head to the ballot box for local elections, they typically vote yes or no on a series of initiatives. The election results in one outcome for each individual topic. On a smaller scale, a group of flatmates might decide on a series of unrelated topics and generate a plan for living together. For example: Should the kitchen be cleaned once a week or twice a week? Should we get the red couch or the yellow couch?

The most natural way to decide the final outcome in all of these scenarios is to take the majority opinion on each individual topic and aggregate them into a unified party platform, legislative agenda, or roommate contract. However, this approach can yield a surprisingly undesirable outcome: a majority of the voters may actually be more unhappy with this result than if the opposite decision were made on every issue (known as *Anscombe's Paradox* [1]). How can this arise? Consider a setting with 5 voters and 3 independent binary issues. The following table illustrates the preferences of each voter on each of the 3 issues: +1 is in favor and −1 is against:

	Issue 1	Issue 2	Issue 3
v_1	+1	−1	−1
v_2	−1	+1	−1
v_3	−1	−1	+1
v_4	+1	+1	+1
v_5	+1	+1	+1

Now, assume each voter would only vote in favor of proposals that they agree with on more than half of the issues (in the paper, a voter will abstain when agreeing with a proposal on exactly half of the issues). Taking the majority on each topic yields the proposal $(+1, +1, +1)$. However, voters 1, 2, and 3 all disagree with a majority of this proposal. Therefore, if we posed this proposal for a vote, a majority of voters would vote against it. If, instead, we posed the opposite proposal $(-1, -1, -1)$, then voters 1, 2, and 3 would support it, and it would win the majority vote. Hence, in this scenario, the proposal comprising the *minority* opinion on each topic wins the majority vote, whereas the proposal comprising the majority opinions fails to get majority support.

An equivalent view on the previous scenario positions Anscombe's paradox in a broader context: instead of assuming a vote on a single proposal with people voting for/against it, let us assume that the vote happens between two competing proposals p and p' and each voter votes for whichever of p and p' agrees with their views on more topics, abstaining in case of equality. Seen as such, Anscombe's paradox is the situation where an *issue-wise* *majority* (IWM) proposal p loses the majority vote against $p' = \overline{p}$, defined as the opposite proposal of p. A less extreme variant of the paradox, known as Ostrogorski's paradox [32] happens when an IWM proposal p loses against some proposal p', not necessarily $\overline{p}$. Settling on the IWM proposal in such cases can lead to daunting situations where one of its opposers calls a final vote between p and p' that "surprisingly" unveils general dissatisfaction with what was otherwise a perfectly democratically chosen outcome.

Consequently, multi-issue aggregation mechanisms need to balance the tension between two majoritarian processes: majority on the individual topics and majority when proposals are compared to one another. In terms of the first, the chosen proposal should ideally stay somewhat close to IWM. In terms of the second, the chosen proposal should not be easily refuted by calling a vote against some other proposal.

Even when voters consider the issues to be of equal importance in their decision-making, we get paradoxical situations. However, in reality, voters rarely consider all issues to be equally important and often disagree on their importance; e.g., a Pew Research study from June 2023 indicated that in the United States, there were massive differences in perceived issue importance along partisan lines [33]. Some voting advice applications already attempt to account for personalized issue-importance, such as Smartvote [3]. Data from these applications can not only help assess how the current parties are aligning with the populace [5], it can also suggest potential new party platforms. Such pre-existing infrastructure to get data on both voter opinions and issue importance underscores the pertinence of issue weights to modeling this problem setting.

1.1 Our Contribution

We study the aggregation of opinions on multiple independent binary issues with respect to two measures of majoritarianism: agreement with issue-wise majority and success in pairwise proposal comparisons. Our analysis considers two weighting models: *external weights* and *internal weights*. In the former, the policy-maker sets a weight to each issue reflecting its relative importance, and voters use weighted agreement when comparing any two proposals. The latter is the same, but each voter is free to choose their own weighting vector. We use the "*unweighted setting*" to refer to the edge case where issues are equally-weighted.

1.1.1 *Condorcet winners.*

In the first part of the paper, we focus on the complexity of determining a *Condorcet-winning proposal*: a proposal that does not lose in a direct vote against any other proposal. Under external weights, we find that any Condorcet winner has to be an IWM proposal, while this does not extend to internal weights. However, even in the unweighted setting with an odd number of voters, where the IWM proposal is unambiguous, checking whether this proposal is a Condorcet winner is co-NP-hard (answering an open question in [13]).

An Ostrogorski-free domain. To mitigate this hardness result, it would be appealing to identify a large set of instances for which IWM proposals are Condorcet winners (i.e., Ostrogorski's paradox does not occur). If membership to this set could also be efficiently verified, this would allow for practically certifying "safe instances" where issue-wise majority is the right choice. We achieve this by the *single-switch* condition of Laffond and Lainé [24]: a preference matrix over ± 1 is single-switch if it admits a *single-switch presentation* — a way to permute and potentially negate some columns such that +1 entries on each row form a prefix or a suffix. They show that for the unweighted case, this condition implies that Ostrogorski's paradox does not occur. We extend and simplify their analysis to show that the same holds under external weights (but not always for internal weights). We then provide a linear-time algorithm for checking whether the preference matrix is single-switch and prove that no-instances admit short proofs of this fact in the form of small forbidden subinstances (that can also be identified in linear time by a black-box reduction to the recognition problem which we have not encountered before).

Secondary implications. Along the way, in this part, we make multiple secondary contributions: (i) we uncover an interesting topological connection: the set of single-switch presentations of a single-switch matrix can be compactly represented as the union of two mirror-image Möbius strips; (ii) our recognition algorithm for single-switch matrices proceeds by reducing to checking whether the columns of a matrix can be permuted so that the ones on each row form a prefix or a suffix — while a linear-time algorithm is known for this [15],[1] it relies on rather complex machinery — we instead give a much simpler direct algorithm with the same guarantees; (iii) our simpler algorithm can be adapted to yield the simplest and at the same time most efficient algorithm for checking the *single-crossing condition* in ranked social choice [16]. Similarly to the single-switch condition, the latter also admits a characterization in terms of small forbidden substructures [10], and finding such forbidden substructures can be achieved within the same time complexity using our black-box technique, a result which to the best of our knowledge is new.

1.1.2 *Representative majority-supported proposals.*

Settling on a Condorcet-winning proposal would be ideal, especially under external weights where such proposals are by default IWM proposals, but in the absence of Condorcet winners, a compromise is needed. In fact, the hardness of checking whether an IWM proposal is a Condorcet winner can be seen positively: it is computationally demanding to find the proposal that defeats it, so we need not fear a vote being called against the defeating proposal. Hence, it is reasonable to relax the demanding Condorcet condition: the chosen proposal should, at the least, not lose against its opposite — or, in the language of our first formulation of Anscombe's paradox above, should garner majority support. In the second part of the paper, we explore existence guarantees for majority-supported proposals that are as close as possible to an IWM proposal p_{IWM}. So far, this has been studied in the unweighted model [13, 20]: a weakly majority-supported proposal agreeing in strictly more than half of the issues with p_{IWM} exists and can be found in polynomial time, while achieving better guarantees is NP-hard. The word "weakly" can be dropped if majority is strict/unambiguous on at least one issue, i.e., some column of the preference matrix has differing numbers of +1's and −1's. We will be interested in the more complex weighted case.

External weights. We provide a matching guarantee to the unweighted case, showing that there always exists a weakly majority-supported proposal with strictly more than half the total weight in

[1] Under the name of recognizing *voter/candidate-extremal-interval* preferences.

topics agreeing with p_{IWM}. Under a simple condition on certain higher-weight issues, we can also drop the word "weakly."

Internal weights. In sharp contrast, with as few as two different weight vectors, we construct families of instances where the distance between every weakly majority-supported proposal and the unique IWM proposal gets arbitrarily large. The severity of such examples is controlled by the maximum average topic weight $\tilde{w}_{max}$: we give a simple bound derived from a partition-based approach that is tight on a large portion of the range $\tilde{w}_{max} \in (0, 1)$.

More paradox-free instances. Finally, we generalize Wagner's Rule of Three-Fourths [36] for both external and internal weights: if the average weighted majority on the issues is at least $\frac{3}{4}$, then Anscombe's Paradox cannot occur. Without loss of generality, if +1 is a majority opinion on each topic, this translates to the total weight of +1's in the preference matrix being at least $\frac{3}{4}$ of the total weight. A stronger condition precludes Ostrogorski's paradox under external weights: if on each column the relative weight of +1's is at least $\frac{3}{4}$ of that column's total weight. This surprisingly simple check is a counterpart to the single-switch condition, once again giving a convenient characterization for a whole class of instances in which returning an IWM proposal is always a good choice.

1.2 Further Related Work

Variations on the question of how best to reach consensus on a series of issues have been studied thoroughly. We first go over models where all topics are considered equally important.

Approval voting is a popular mechanism that is frequently used for single-winner and multi-winner elections alike [7, 18]. Here, each participant indicates their approval for a subset of candidates. In contrast to our setting, not expressing the approval of a candidate does not give the same signal as voting for the "no" stance on an issue (which is a vote for the logical negation of the issue) [23].

Another related field of study is judgment aggregation, where a series of judges have viewpoints on multiple topics, but there is external logical consistency required between the topics [29]. As in our problem, a reasonable method of reaching consensus is to take the majority opinion on each topic. However, the outcome may fail to be logically consistent — this is the *Discursive Dilemma*, and can occur with as few as 3 judges and 3 topics [22]. There has been some investigation into conditions that avoid this paradox, like List's *unidimensional alignment* [28],[2] and other similar paradoxes under the name of *compound majority paradoxes* [30].

Our problem can also be viewed as a special instance of voting in combinatorial domains: multiple referenda with separable topics [9]. Multiple works explored generalizations of Anscombe's paradox and gave further impossibility results [4, 21], e.g., relating to the Pareto optimality of aggregation rules [31].

Significant work has also been done to characterize when such paradoxes *cannot* occur. Wagner proposed the Rule of Three-Fourths [36], preventing Anscombe's paradox, as well as a generalization [37]. Laffond and Lainé showed that if no two voters disagree on too many issues, then Anscombe's is prevented [25], and for *single-switch* preferences, Ostrogorski's does not occur [24].

[2]The unidimensional alignment condition might appear to closely resemble the single-switch condition, as it essentially requires that the transposed preference matrix be single-switch. However, this is not equivalent, as rows and columns play different roles — issue-wise majority aggregates along columns, not rows.

We now survey proposals to augment various voting systems with weights, allowing voters to express their degrees of interest or investment in the topics. Storable voting allows participants to delay using their vote in a given election, and accumulate votes to use in later elections that they have more stake in [12]. Quadratic voting proposes a somewhat similar system in which people are given an allotment of vote credits, and before a given election can buy a certain number of votes [26]. Both of these systems maintain that voters will use more votes for elections in which they feel strongly and believe they are likely to be pivotal in. Uckelman introduces a framework using goalbases to express cardinal (numeric) preferences over a combinatorial voting domain [35]. This, however, loses information by abstracting away the separability of issues: for us, the cardinal preferences are induced by the weighted Hamming distance. Lang also considers augmenting combinatorial voting with preference weights and provides several computational complexity results [27]. Satisfaction approval voting [8] modifies approval voting by spreading a voter's total weight equally over all of the candidates they approve of. Finally, there is recent interest in studying how voters have varying stakes in elections and how to accommodate these stakes to limit distortion [11, 19].

2 MODEL AND NOTATION

For any non-negative integer m, write $[m] := \{1, \ldots, m\}$. Given a real number x, write $sgn(x) \in \{-1, 0, 1\}$ for its *sign*. Note that for any two reals x, y, we have that $sgn(x \cdot y) = sgn(x) \cdot sgn(y)$.

We consider a setting with n voters and t independent, binary issues/topics. The decision space for each issue is $\mathbb{B} := \{\pm 1\}$. Each voter $i \in [n]$ is modeled as a dimension-t vector $v_i \in \mathbb{B}^t$ indicating for each issue $j \in [t]$ the opinion/preference $v_{i,j} \in \mathbb{B}$ of voter i on issue j. We call the matrix $\mathcal{P} = (v_{i,j})_{i\in[n],j\in[t]}$ the *preference profile*. We also write $\mathcal{P} = (c_1, \ldots, c_t)$, where $c_1, \ldots, c_t \in \mathbb{B}^n$ are the columns of the matrix.

For each issue $j \in [t]$, we are consistent with previous literature [13, 20, 36, 37] and define the *majority* $m_j \in [0, 1]$ on issue j to be the fraction of voters that prefer +1 on it; i.e., the number of +1's in c_j, divided by n. If $m_j > 0.5$, then the *majority opinion* on issue j is +1; if $m_j < 0.5$, then it is −1, and if $m_j = 0.5$, then both +1 and −1 are majority opinions on issue j.

A *proposal* is a vector $p \in \mathbb{B}^t$ that consists of a decision for each issue. We write $\overline{p}$ for the *complement* of proposal p, which simply flips each bit of p; i.e., $\overline{p} = -p$. An *issue-wise majority (IWM)* is a proposal p where the decision on each topic is a majority opinion for the topic.

We study two weighting models: *external weights* and *internal weights*. In the former, an externally supplied vector of non-negative weights $w = (w_1, \ldots, w_t)$ summing up to 1 is available, denoting the importance of each issue as seen collectively by the voters. The internal weights model generalizes this by having each voter $i \in [n]$ report an individual vector of weights $w_i = (w_{i,1}, \ldots, w_{i,t})$; i.e., there need no longer be consensus on the importance of any fixed issue. For internal weights, we write W for the matrix with rows $w_1, \ldots, w_n$. We call the *voting instance* the pair $\mathcal{I} = (\mathcal{P}, W)$ for internal weights and $\mathcal{I} = (\mathcal{P}, w)$ for external weights. We will also talk about the *unweighted* model, which is simply external weights with $w = (1/t, \ldots, 1/t)$, and directly write $\mathcal{I} = \mathcal{P}$ for it.

For the remainder of this section, we assume external weights — the internal weights model requires substantial additional notation so we postpone it to later on.

For any positive integer m, given two vectors $u, v \in \mathbb{B}^m$ and a vector of weights $w \in [0,1]^m$ with unit sum, we write $d_H(u, v, w) := \sum_{j=1}^{m} w_j \cdot \mathbb{I}(u_j \neq v_j)$ for the *w-weighted Hamming distance* between u and v. We omit the w argument when referring to the unweighted Hamming distance. For convenience, we write $\langle u, v\rangle_w := \sum_{j=1}^{m} w_j \cdot u_j \cdot v_j$ for the standard *w-weighted inner/dot-product.* One can easily show that $\langle u, v\rangle_w = 1 - 2 \cdot d_H(u, v, w)$.

Fix an instance $\mathcal{I} = (\mathcal{P}, w)$ in the external weights model. For each voter i with vote v_i we define their *individual* preference relation $\succcurlyeq_i$ between proposals. In particular, given two proposals $p, p' \in \mathbb{B}^t$, voter i weakly prefers p over p', written $p \succcurlyeq_i p'$, iff $d_H(v_i, p, w) \leq d_H(v_i, p', w)$. Note that this is equivalent to $\langle v_i, p\rangle_w \geq \langle v_i, p'\rangle_w \iff \langle v_i, p - p'\rangle_w \geq 0$. We write $\succ_i$ and $\approx_i$ for the strict and symmetric parts of $\succcurlyeq_i$, respectively. We define the *collective* preference relation $\succcurlyeq_{\mathcal{I}}$ between proposals: given two proposals $p, p' \in \mathbb{B}^t$, the voters collectively weakly prefer p over p', written $p \succcurlyeq_{\mathcal{I}} p'$, iff $|\{i \in [n] : p \succ_i p'\}| \geq |\{i \in [n] : p' \succ_i p\}|$. Note that this is equivalent to $\sum_{i=1}^{n} sgn(\langle v_i, p - p'\rangle_w) \geq 0$. We write $\succ_{\mathcal{I}}$ and $\approx_{\mathcal{I}}$ for the strict and symmetric parts of $\succcurlyeq_{\mathcal{I}}$, respectively. A proposal $p \in \mathbb{B}^t$ is a *Condorcet winner* if for any other proposal $p' \in \mathbb{B}^t$ we have $p \succcurlyeq_{\mathcal{I}} p'$.

For a voting instance $\mathcal{I}$, Ostrogorski's paradox occurs if some IWM proposal p_{IWM} is not a Condorcet winner, Anscombe's paradox occurs if for some IWM proposal p_{IWM} we have $\overline{p_{IWM}} \succ_{\mathcal{I}} p_{IWM}$, and the Condorcet paradox happens if there is no Condorcet-winning proposal.

Full Version. The full version of the paper [2] contains the material omitted due to space constraints.

3 COMPLEXITY OF DETERMINING A CONDORCET WINNER

In this section, we prove that it is co-NP-hard to determine whether an instance $\mathcal{I}$ admits a Condorcet-winning proposal, even in the unweighted setting with odd n:

Theorem 1. *Deciding whether an instance $\mathcal{I} = \mathcal{P}$ admits a Condorcet winner is co-NP-hard in the unweighted setting with odd n.*

This could be surprising given the following observation of [24] for the unweighted model, which we extend to external weights:

Lemma 2. *Consider an external-weights instance $\mathcal{I}$ such that $p \in \mathbb{B}^t$ is a Condorcet winner for $\mathcal{I}$. Then, p is an IWM for $\mathcal{I}$.*

Proof. Assume the contrary, then there is an issue $j \in [t]$ such that p_j is not a majority opinion on issue j. Consider the proposal p^* obtained from p by flipping p_j. Then, $p^* \succ_{\mathcal{I}} p$, a contradiction. □

Lemma 2 shows that one can restrict the search space for Condorcet winners to IWM proposals. In the unweighted setting with odd n, there is a single such proposal, which we can assume without loss of generality to be $\mathbf{1} \in \mathbb{B}^t$. Nevertheless, even under these conditions, we will show that checking whether $\mathbf{1}$ is a Condorcet winner is co-NP-hard, or, equivalently, checking whether $\mathbf{1}$ is *not* a Condorcet winner is NP-hard. The latter occurs if and only if there is a proposal $p \in \mathbb{B}^t$ such that $p \succ_{\mathcal{I}} \mathbf{1}$, which, recall, means that strictly more voters $i \in [n]$ prefer $p \succ_i \mathbf{1}$ than $\mathbf{1} \succ_i p$. Hence, it suffices to prove that the following problem is NP-hard:

Problem "Major"
Input: Instance $\mathcal{I} = \mathcal{P}$ in the unweighted setting with odd n such that $\mathbf{1}$ is the issue-wise majority.
Output: Does there exist a proposal $p \in \mathbb{B}^t$ s.t. $p \succ_{\mathcal{I}} \mathbf{1}$?

To show its hardness, we need the following auxiliary problem:

Problem "Unanim"
Input: Voting instance $\mathcal{I} = \mathcal{P}$ in the unweighted setting.
Output: Does there exist a proposal $p \in \mathbb{B}^t$ s.t. $p \succ_i \mathbf{1}$ for all $i \in [n]$ (to be read "p *unanimously* defeats $\mathbf{1}$")?

Unanim is NP-hard [14, Theorem 2], but the proof in [14] is relatively complicated: we give a simpler one in the full version by noting the equivalence to choosing a subset of columns of $\mathcal{P}$ that sum up to a negative amount on each row (we also give a similar reformulation of Major for the interested reader).

Lemma 3. *Major is NP-hard.*

Proof. We reduce from the NP-hard problem Unanim. Consider an instance $\mathcal{I} = \mathcal{P}$ of Unanim with n voters. If there is an issue $j \in [t]$ disapproved by all voters in $\mathcal{P}$, then $\mathcal{P}$ is a yes-instance of Unanim: all voters prefer the proposal with +1 in all coordinates except the j-th to proposal $\mathbf{1}$. This case can be easily detected in polynomial time, so we henceforth assume the contrary.

We build an instance $\mathcal{I}' = \mathcal{P}'$ of Major from $\mathcal{P}$ by adding $n-1$ voters approving all issues. For $\mathcal{P}'$ to be a valid instance for Major we need that $2n-1$ is odd (which it is) and that $\mathbf{1}$ is the issue-wise majority. The latter holds because at least $n-1+1 = n$ voters approve of each issue: the $n-1$ added ones and at least one from the first n by our assumption. It remains to show that a proposal $p \in \mathbb{B}^t$ unanimously defeats $\mathbf{1}$ in $\mathcal{P}$ iff it majority-defeats $\mathbf{1}$ in $\mathcal{P}'$.

Assume $p \in \mathbb{B}^t$ unanimously defeats $\mathbf{1}$ in $\mathcal{P}$. Then, each of the first n voters in $\mathcal{P}'$ prefers p to $\mathbf{1}$. Since there are only $n-1 < n$ other voters in $\mathcal{P}'$, a majority of the voters in $\mathcal{P}'$ prefer p to $\mathbf{1}$.

Conversely, assume $p \in \mathbb{B}^t$ majority-defeats $\mathbf{1}$ in $\mathcal{P}'$. Clearly, $p \neq \mathbf{1}$ has to hold, so all of the $n-1$ added voters prefer $\mathbf{1}$ to p. To counteract this, since $p \succ_{\mathcal{I}'} \mathbf{1}$, the first n voters in $\mathcal{P}'$ must prefer p to $\mathbf{1}$, meaning that p unanimously defeats $\mathbf{1}$ in $\mathcal{P}$. □

For completeness, we put the pieces together to give a self-contained proof of Theorem 1 in the full version.

4 AN OSTROGORSKI-FREE DOMAIN

As we have seen, at least for external weights, a Condorcet-winning proposal has to be an issue-wise majority proposal. Yet, we proved that determining whether one of them is actually Condorcet-winning is co-NP-hard, even in the unweighted case with odd n, where there is only one such proposal to check. To mitigate this hardness result, it would be useful if we could identify a large set of instances for which IWM proposals are guaranteed to be Condorcet-winning, i.e., Ostrogorski's paradox does not occur. Laffond and Lainé [24] introduced the *single-switch* condition, which achieves exactly this goal for the unweighted setting. Furthermore, they showed that it is the

1	2	3	4	5	6
+1	+1	+1	−1	+1	−1
+1	−1	+1	+1	−1	+1
+1	+1	−1	−1	+1	−1

(a) Profile $\mathcal{P}$.

2	1	3	4	$\overline{5}$	6
+1	+1	+1	−1	−1	−1
−1	+1	+1	+1	+1	+1
+1	+1	−1	−1	−1	−1

(b) Single-switch presentation of $\mathcal{P}$.

Figure 1: The profile $\mathcal{P}$ in Fig. 1a is single-switch because its columns can be permuted and flipped as in Fig. 1b to ensure that ones on each row form a prefix or a suffix.

most general condition preventing Ostrogorski's paradox among conditions that do not consider the multiplicities of the votes (i.e., conditions defining a *domain*) or whether a vote is negated or not (i.e., they only look at the set $\{\{v_i, \overline{v_i}\} \mid i \in [n]\}$ and not at how many times each v_i or $\overline{v_i}$ is repeated). In particular, if an instance in the unweighted model is not single-switch, then it is possible to add copies of some of the votes v_i (or their negations $\overline{v_i}$) so that some issue-wise majority proposal is not a Condorcet winner. Two important questions underpinning their condition are: (i) Does it still guarantee the existence of a Condorcet winner in the (at least externally) weighted setting? (ii) Is it possible to check whether it applies in polynomial time? If not, are there short proofs of this fact? Here, we answer all these questions in the affirmative.

A preference profile (matrix) $\mathcal{P} = (c_1, \ldots, c_t)$ is *single-switch* (SSW) if we can flip (multiply by -1 all entries in) some columns and then permute the columns to get a new profile $\mathcal{P}'$ such that $+1$ entries on every row form either a prefix or a suffix, in which case we say that $\mathcal{P}'$ is an SSW presentation of $\mathcal{P}$. We allow flipping no columns or leaving all columns in their original place. Intuitively, issues are arranged along a left-right axis. Left-wing voters approve a prefix of issues, with the length depending on their tolerance, while right-wing voters similarly approve a suffix of issues.[3] See Fig. 1 for an illustration of the notion. A voting instance $\mathcal{I}$ is single-switch if its preference profile $\mathcal{P}$ is single-switch.

4.1 For External Weights Single-Switch Prevents Ostrogorski's Paradox

We find that, assuming external-weights, the single-switch condition guarantees that all IWM proposals are Condorcet winners. To show this, we first show that every issue-wise majority proposal does not lose against its opposite, i.e., Anscombe's paradox does not occur. We do this by streamlining and adapting the argument in [24] (which was only for the unweighted model). Because the single-switch condition is closed under removing issues, the general statement then follows easily by noting that, under external weights, Ostrogorski's paradox happens if and only if there is a subset of issues inducing an instance where Anscombe's paradox happens. The details are deferred to the full version.

Theorem 4. *In the external-weights model, every issue-wise majority proposal of a single-switch instance is a Condorcet winner.*

[3]This shares similarities with several related concepts, such as single-peaked and single-crossing preferences. However, unlike other notions, we allow issues to be flipped before ordering them, as they can be logically negated without changing meaning.

4.2 Recognizing Single-Switch Profiles

The result in the previous section is particularly appealing: in the external-weights model, if the preferences are single-switch, any issue-wise majority proposal is a Condorcet winner. This bypasses our previous hardness result in the case of single-switch preferences. However, this is only useful provided one can quickly tell whether a given profile $\mathcal{P}$ is single-switch or not. In this section, we show that this can be determined in linear time, i.e., $O(nt)$. For yes-instances, our algorithm also determines an SSW presentation $\mathcal{P}'$ (implicitly also the permutation and flips used to obtain it). Given $\mathcal{P}'$, we also characterize the set of all SSW presentations as the union of two "orbits" around $\mathcal{P}'$ and its column-reversal. These orbits can be attractively interpreted topologically as two mirror-image Möbius strips. To begin, we need the following observation following easily from the case $n = 1$. See the full version for the proof.

Lemma 5. *Consider a profile $\mathcal{P}$ admitting an SSW presentation $\mathcal{P}' = (c_1, \ldots, c_t)$. Then, $\mathcal{P}'_r := (c_2, \ldots, c_t, \overline{c_1})$ is also a SSW presentation of $\mathcal{P}$. Furthermore, any t (circularly) consecutive columns in $\mathcal{P}'' := (c_1, \ldots, c_t, \overline{c_1}, \ldots, \overline{c_t})$ form an SSW presentation of $\mathcal{P}$.*

Hence, any SSW presentation $\mathcal{P}'$ of a profile $\mathcal{P}$ corresponds to a set of $2t$ such presentations that we call the *orbit* $O_{\mathcal{P}'}$ of $\mathcal{P}'$. Formally, these are the $2t$ profiles that can be obtained by taking t (circularly) consecutive columns in $\mathcal{P}''$ in the above. Note that the orbits of any two SSW presentations either coincide or are disjoint, so the set of all orbits partitions the set of SSW presentations of $\mathcal{P}$. Also, the $2t$ profiles in $O_{\mathcal{P}'}$ are pairwise distinct, which can be easily seen by considering the case $n = 1$, under which $\mathcal{P}''$ is circularly equivalent to a list of t minus ones followed by t ones. This reasoning additionally allows us to assign to each orbit a *representative*, namely the profile with all -1's on the first row:

Corollary 6. *Every orbit contains exactly one profile where the first row is all -1's.*

Orbits can be understood through a topological lens: For the orbit $O_{\mathcal{P}'}$ of $\mathcal{P}' = (c_1, \ldots, c_t)$ take an $n \times t$ rectangular piece of paper and write the columns $c_1, \ldots, c_t$ on the front and $\overline{c_1}, \ldots, \overline{c_t}$ on the back, such that for each $i \in [t]$, column c_i on the front aligns with column $\overline{c_i}$ on the back. Then, give the paper a length-wise half-twist and glue the left and right sides to form a surface known as a Möbius strip: see Fig. 2. Cutting along the width of the strip between *any* two columns recovers an $n \times t$ piece of paper with one SSW presentation on one side and its opposite on the other side. In high-level terms, each orbit is topologically a Möbius strip.

To check whether a profile $\mathcal{P}$ is single-switch, by Corollary 6, it suffices to check for presentations with all -1's in the first row: all other presentations are generated by the orbits of such presentations. There is a simple strategy to achieve this: flip columns in $\mathcal{P}$ to make the first row all -1's, and then check whether columns in the resulting profile can be permuted to ensure that ones on each row form a prefix or a suffix. This amounts to recognizing single-switch-no-flips profiles: A profile $\mathcal{P}$ is *single-switch-no-flips* (SSWNF) if its columns can be permuted to get a new profile $\mathcal{P}'$ such that $+1$ entries on every row form either a prefix or a suffix, in which case we say that $\mathcal{P}'$ is an SSWNF presentation of $\mathcal{P}$.

Recognizing single-switch-no-flips profiles. Telling whether a profile $\mathcal{P} = (c_1, \ldots, c_t)$ is single-switch-no-flips can be achieved

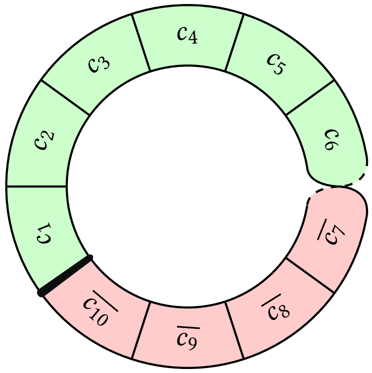

Figure 2: Möbius strip of orbit $O_{\mathcal{P}'}$ for $\mathcal{P}' = (c_1, \ldots, c_{10})$. We start with a rectangular piece of paper of length 10 and write $(c_1, \ldots, c_{10})$ on the (green) front side and $(\overline{c_1}, \ldots, \overline{c_{10}})$ on the (red) backside. We then give the paper a length-wise half-turn and glue the endpoints (bold strip). This gives raise to a surface with a single continuous side.

by appending a negated copy of $\mathcal{P}$ underneath [15] and running a solver for the *Consecutive Ones Problem* (C1P), which can be solved in $O(nt)$ time [6], implying the same about our problem. However, such solvers are complicated and notoriously error-prone: most available implementations fail on at least some edge cases [17]. Moreover, reducing to C1P does not utilize the additional structure present in our problem and hence does not shed light on the structure of all solutions, as we set out to do. We give a much simpler algorithm achieving the $O(nt)$ time-bound: Find an index x maximizing $d_H(c_1, c_x)$. Then, sort (using Counting Sort) the columns based on their Hamming distance from c_x to get a profile $\mathcal{P}' = (c'_1, \ldots, c'_t)$ where $d_H(c_x, c'_i) \leq d_H(c_x, c'_{i+1})$ for $i \in [t-1]$ (i.e., ties in Hamming distance can be broken arbitrarily). We claim that either $\mathcal{P}'$ is the unique SSWNF presentation of $\mathcal{P}$ (up to reversing the order of the columns), or there is no such presentation, so we can easily check in additional $O(nt)$ time whether the candidate solution works. All required claims are shown in the full version:

Theorem 7. *There is a* simple *$O(nt)$ algorithm computing (or deciding the inexistence of) an SSWNF presentation of a profile $\mathcal{P}$. Moreover, if it exists, this presentation is unique up to reversing column order.*

The full version also provides a much ampler discussion of related work for this sub-problem, including the relation between our algorithm and previous algorithms for recognizing *single-crossing preferences*. As a bonus, it gives a similar simpler, more efficient algorithm for recognizing single-crossing preferences, running in time $O(nt\sqrt{\log n})$, improving state of the art [16, Algorithm 4].

Putting it together. To decide whether a profile $\mathcal{P}$ is single-switch, we flip columns in $\mathcal{P}$ to get a profile $\mathcal{P}'$ with only -1's in the first row and then use the algorithm in Theorem 7 to find an SSWNF presentation $\mathcal{P}''$ of $\mathcal{P}'$ (and hence also $\mathcal{P}$). If it exists, this presentation is unique up to column reversal, so we can also characterize the set of all SSW presentations of $\mathcal{P}$ by unioning the orbits of $\mathcal{P}''$ and its column-reversal. Note that these two orbits may coincide for pathological input profiles $\mathcal{P}$.

Theorem 8. *There is an $O(nt)$ algorithm computing (or deciding the inexistence of) an SSW presentation of a profile $\mathcal{P}$. If the algorithm returns a presentation $\mathcal{P}''$, let $\mathcal{P}''_r$ be $\mathcal{P}''$ with the order of the columns reversed, then the set of all SSW presentations of $\mathcal{P}$ is $O_{\mathcal{P}''} \cup O_{\mathcal{P}''_r}$.*

4.3 Forbidden Subprofiles Characterization of Single-Switch Preferences

Whenever the single-switch condition is not satisfied, it would be useful if there were a short proof of this fact: a small subprofile that is not single-switch. Formally, a profile/matrix $\mathcal{P}$ contains a profile/matrix $\mathcal{P}'$ as a *subprofile/submatrix* if we can remove (possibly zero) rows and columns from $\mathcal{P}$ to get $\mathcal{P}'$ up to permuting rows and columns. Note that existence is not immediate: there could exist arbitrarily large matrices not satisfying the condition but all of whose proper submatrices do. We show that this is not the case: either the condition holds, or there is a 3×4 or 4×3 submatrix witnessing that this is not the case, as in the following:

Theorem 9. *A profile $\mathcal{P}$ is single-switch if and only if it does not contain as a subprofile $\mathcal{P}^a_1, \mathcal{P}^a_2$ and any profile that can be obtained from them by flipping rows and columns:*

$$\mathcal{P}^a_1 = \begin{bmatrix} -1 & -1 & -1 & -1 \\ +1 & +1 & -1 & -1 \\ +1 & -1 & +1 & -1 \end{bmatrix} \quad \mathcal{P}^a_2 = \begin{bmatrix} -1 & -1 & -1 \\ +1 & -1 & -1 \\ -1 & +1 & -1 \\ -1 & -1 & +1 \end{bmatrix}$$

We prove Theorem 9 in the full version by combining a similar characterization for single-switch-no-flips profiles given in [34] (under the name *voter/candidate-extremal-interval* preferences) with our insight that to go to the no-flips version it suffices to make one row all -1's. Henceforth, we call the 3×4 and 4×3 preference profiles in the theorem above *forbidden subprofiles*. Then, the theorem says that $\mathcal{P}$ is single-switch if and only if it contains no forbidden subprofiles. Note how this implies that single-switch profiles are relatively rare: the probability that a random binary $n \times t$ matrix is single-switch tends to zero as n and t tend to infinity.

Finding forbidden subprofiles. So far, we have seen that non-membership to the class of single-switch preferences admits short proofs, but can such proofs also be constructed efficiently? Given some no-instance, it is straightforward to determine which forbidden subprofiles occur in it in time $O(n^3t^4 + n^4t^3)$. In contrast, our recognition algorithm runs in time $O(nt)$, but does not identify a forbidden subprofile. We will now assume our $O(nt)$ recognition algorithm as a black box and show how to identify a forbidden subprofile for a given no-instance $\mathcal{P}$ in time $O(nt)$.

Let us first describe an $O(n^2t + nt^2)$ approach: one at a time, try to remove each row and each column of $\mathcal{P}$, i.e., $n + t$ removal attempts; if doing so makes the resulting profile a yes-instance, undo the removal, and otherwise let it persist. At the end, the ensuing no-instance $\mathcal{P}'$ is a subprofile of $\mathcal{P}$ whose proper subprofiles are yes-instances, so $\mathcal{P}'$ is a forbidden subprofile, completing the argument.

We now modify the previous idea to run in time $O(nt)$ by removing multiple rows/columns at a time. We will first only remove rows, and then, starting from the resulting profile, only columns. The reasoning for columns is entirely analogous, so we only describe the procedure for rows: partition the rows into 5 groups $G_1, \dots, G_5$, each of size roughly $n/5$. Because all forbidden subprofiles are of size 3×4 or 4×3, any occurrence of a forbidden subprofile in $\mathcal{P}$ only uses rows from at most 4 of the 5 groups. Consequently, we can find a group G_i such that removing all rows in G_i from $\mathcal{P}$ keeps the property that $\mathcal{P}$ is a no-instance. Doing so requires at most 5 calls to the recognition algorithm, so it can be done in overall time $O(nt)$. Ignoring for brevity the cases where n is not divisible by 5, this reasoning shows how to reduce n to $4n/5$ in time $O(nt)$. Applying the same reasoning iteratively until n goes below 5 takes total time $O(nt)$ because the geometric series $\sum_{i=0}^{\infty}(4/5)^i$ converges.

Theorem 10. *Given a non-single-switch profile $\mathcal{P}$, a forbidden subprofile of $\mathcal{P}$ can be determined in time $O(nt)$.*

We note that the previous idea applies more broadly; e.g., for single-crossing preferences, which admit a characterization in terms of two small *forbidden subinstances* [10], our $O(nt\sqrt{\log n})$ recognition algorithm can be bootstrapped to also produce a forbidden subinstance for no-instances within the same time bound. A formal statement and more details can be found in the full version.

5 ANSCOMBE'S PARADOX

When preferences are not single-switch, determining whether an IWM proposal is a Condorcet winner is co-NP hard. In light of this, we focus on the most diabolical subset of Ostrogorski paradox instances: those inducing Anscombe's paradox (where an IWM proposal is defeated by its complement, or, equivalently, an IWM proposal fails to get majority support). If Anscombe's paradox occurs, a natural question is: "How close can we get to any given IWM while still requiring that the proposal gets majority support?"

We first explore this question under external weights, i.e., in instances $\mathcal{I} = (\mathcal{P}, w)$ where all voters share the same, unit-sum weights vector w. Then, we introduce the necessary notation and study it for internal weights. Finally, we give a simple characterization of a broad swath of instances that avoid Anscombe's paradox entirely for internal weights. We assume throughout that $t > 1$, as Anscombe's paradox does not occur with one topic, and without loss of generality that $m_j \geq 0.5$ for all $j \in [t]$ (i.e., that +1 is a majority opinion on all topics).

Formally, some voter i supports (approves of) a proposal p if $d_H(v_i, p, w) < 1/2$, opposes (disapproves of) p if $d_H(v_i, p, w) > 1/2$, and is indifferent to p if $d_H(v_i, p, w) = 1/2$. A proposal is *strictly majority-supported* if more people support it than oppose it and *weakly majority-supported* if no more people oppose it than support it. Our definition of majority support matches [13] but differs from [20] (where indifferent voters count towards the proposal's support).

5.1 External Weights

In the unweighted case, it is straightforward to argue that for any IWM, there exists a weakly majority-supported proposal within distance $\leq \frac{1}{2} + \frac{1}{2t}$ because at least one proposal in every complementary pair $(p, \overline{p})$ gets weak majority support (and at least one pair satisfies the distance bound for both proposals). A slightly better guarantee of distance $< \frac{1}{2}$ holds by a more difficult proof [13, 20]. For external weights, the complementary pairs argument no longer gives a bound close to $\frac{1}{2}$ if no subset of topic weights sum up close to $\frac{1}{2}$. One may hope to reduce to the unweighted case by splitting topics into multiple equal-weight topics and use the $< \frac{1}{2}$ bound there, but the resulting majority-supported proposals may have different values for an original topic's clones, making it hard to translate to proposals in the original instance. Despite these setbacks, we surprisingly find that the $< \frac{1}{2}$ guarantee still holds for external weights. Our proof, deferred to the full version, simplifies and adapts the argument in [13]. We also guarantee *strict* majority support if there is a strict majority in at least one *relevant* topic, roughly meaning topics with high enough weight to be the tipping point in a vote (see the full version for a formal definition).

Theorem 11. *For any $\mathcal{I} = (\mathcal{P}, w)$ and p_{IWM}, there is a weakly majority supported proposal p with $d_H(p, p_{IWM}, w) < 1/2$. If majority is strict in any relevant topic, "weak" can be replaced with "strict".*

5.2 Internal Weights

We now explore a model where individuals can have unique weight vectors, expressing not only diverse preferences on issue outcomes but also differing opinions on relative topic importance.

Internal Weights Model. In the internal weights model, an instance $\mathcal{I} = (\mathcal{P}, W)$ consists of a preference profile $\mathcal{P}$ and a weight profile W with rows $w_1, \dots, w_n$ where each weight vector w_i corresponds to voter i, is non-negative, and sums to 1. The *average weight vector* is defined as $\tilde{w} := \frac{1}{n}\sum_{i=1}^{n} w_i$. Zero entries in the average weight vector correspond to issues that no voters placed any weight on (and hence can be ignored). We assume no such topics exist without loss of generality. We define the *majority* for a given topic j to be $m_j := \frac{1}{n\tilde{w}_j}\sum_{i=1}^{n} w_{i,j} \cdot \mathbb{I}(v_{i,j} = +1)$. This is the fraction of voter weight placed on that issue that prefers +1. Note that this agrees with our previous definition for external weights (where it was just the fraction of voters that prefer +1 on that topic). The average majority for a given preference profile is defined as $\tilde{m} := \sum_{j=1}^{t} \tilde{w}_j m_j$. This naturally weights consensus on issues proportionally to how important those issues are to the population.

Under external weights, we could give a constant upper bound (Theorem 11) on the minimum distance of some majority-supported proposal from an IWM, independent of the weight profile. As we will see in Theorems 12 and 13, the severity of Anscombe's Paradox under *internal* weights is closely related to the maximum average topic weight $\tilde{w}_{max}$ (the maximum entry in $\tilde{w}$). Formally, we will upper bound the worst-case *IWM distance* g_ℓ for instances with maximum average topic weight $\tilde{w}_{max} = \ell \in (0, 1)$ and selections of p_{IWM} for the instance:

$$g_\ell := \max_{\substack{\mathcal{I}=(\mathcal{P},W),\ p_{IWM} \\ s.t.\ \tilde{w}_{max}=\ell}} \left(\min_{p \text{ weakly majority-supported}} d_H(p, p_{IWM}, \tilde{w}) \right)$$

We first give a simple upper bound on g_ℓ for $\ell \in (0, 1)$ derived from a partition-based algorithm. Surprisingly, we then show that this seemingly weak upper bound is tight for a large portion of the range $\tilde{w}_{max} \in (0, 1)$. Our lower-bound constructions more strongly imply the existence of instances where *all* weakly majority-supported proposals are far from *all* IWM's. Fig. 3 provides a summary of the bounds we give on $g_{\tilde{w}_{max}}$.

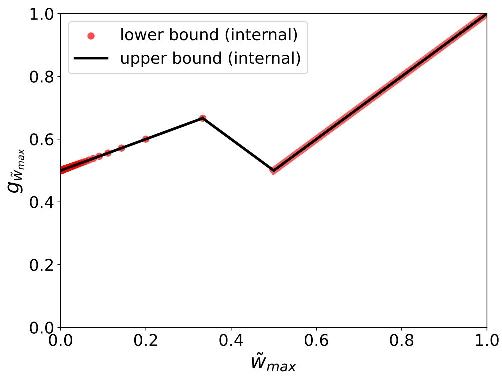

Figure 3: A summary of our bounds on $g_{\tilde{w}_{max}}$.

Partition-based upper bounds. Theorem 12 guarantees both the existence of reasonable majority-supported proposals and provides an algorithm to efficiently recover them.

Theorem 12. *We have the following upper bounds on g_ℓ:*

- *If $\ell \in (0, 1/3)$, then $g_\ell \leq 1/2 + \ell/2$;*
- *If $\ell \in [1/3, 1/2]$, then $g_\ell \leq 1 - \ell$;*
- *If $\ell \in (1/2, 1)$, then $g_\ell \leq \ell$.*

In each case, we can compute a weakly majority-supported proposal p with $d_H(p, p_{IWM}, \tilde{w})$ at most the given bound in polynomial time.

The full proof is deferred to the full version, but the intuition is as follows: for any proposal, either it or its complement will get weak majority support (potentially both), and for any p_{IWM}, $d_H(\overline{p}, p_{IWM}, \tilde{w}) = 1 - d_H(p, p_{IWM}, \tilde{w})$. Therefore, we construct p that keeps $\max\{d_H(\overline{p}, p_{IWM}, \tilde{w}), d_H(p, p_{IWM}, \tilde{w})\}$ small. This is ultimately equivalent to the partition optimization problem with the t entries in the average weight vector as inputs. Our bounds are constructive and give the pair $(p, \overline{p})$ achieving the bound.

Lower bounds. By definition, $g_\ell \leq 1$, so the upper bounds in Theorem 12 might seem fairly weak. However, in Theorem 13, we show that they are actually tight for many values of ℓ. This implies that when $\tilde{w}_{max}$ is large, $g_{\tilde{w}_{max}}$ can get arbitrarily close to 1.

Theorem 13. *The following lower bounds for g_ℓ hold:*

- *If $\ell = 1/(2k+1)$ with $k \in \mathbb{Z}_{\geq 0}$, then $g_\ell \geq 1/2 + \ell/2$;*
- *If $\ell \in (1/2, 1)$, then $g_\ell \geq \ell$.*

We conjecture that the upper bounds in Theorem 12 are tight for the remaining values of ℓ, but leave this to future work. The proof of Theorem 13 is deferred to the full version, but we provide the construction for $\ell \in (1/2, 1)$ and some intuition here. In the instance below, we choose x large enough such that $\tilde{w}_{max} = \ell$ and the first issue holds a strict majority of the weight for all voters. There are x copies of the first voter, and $x+1$ copies of the second.

$$\mathcal{P} = \begin{array}{r}(x)\times\\(x+1)\times\end{array}\begin{bmatrix} +1 & +1 \\ -1 & +1 \end{bmatrix} \quad \mathcal{W} = \begin{array}{r}(x)\times\\(x+1)\times\end{array}\begin{bmatrix} \frac{x+1}{x}\cdot\ell & 1-\frac{x+1}{x}\cdot\ell \\ \frac{x}{x+1}\cdot\ell & 1-\frac{x}{x+1}\cdot\ell \end{bmatrix}$$

In this instance, all voters are essentially "single-issue voters" on the first topic, but the second type of voters split their weight slightly more evenly between the two topics. +1 is the weighted majority opinion on the first topic, but any proposal with +1 for that topic will not get majority support because voters of the second type will oppose it. Notably, **1** is the unique IWM in our constructions, implying there is no majority-supported proposal close to *any* IWM.

Theorem 13 quashes any hope of improving on Theorem 12 and proving a similar result to the external weights setting (where $g_\ell < 1/2$ held for any weights profile). Once voters can have distinct weight vectors, increasing $\tilde{w}_{max}$ can make the distance between all majority-supported proposals and IWM proposals arbitrarily large. We conclude this section by characterizing a group of voting instances in which Anscombe's Paradox will not occur.

Condition precluding Anscombe's Paradox. We find that generalizations of Wagner's Rule of Three-Fourths hold in both the external and internal weights settings:

Theorem 14. *If $\tilde{m} \geq 3/4$ then Anscombe's paradox will not occur. Additionally, if $m_j \geq 3/4$ for all $j \in [t]$ in the external weights setting, then Ostrogorski's paradox will not occur.*

Our proof (deferred to the full version) follows Wagner's original proof strategy of counting agreement with an IWM in an instance in two ways: column-wise and row-wise, but is modified to account for weights. We get the second part of our claim by using the fact that, under external weights, Ostrogorski's paradox occurs if and only if there is a subset of issues inducing an instance where Anscombe's paradox occurs.

6 CONCLUSION AND FUTURE WORK

We explored how best to represent the will of voters on multiple, separable issues when optimizing for two potentially conflicting ideals: agreement with issue-wise majority and success in pairwise proposal comparisons. Additionally, we augmented previous multi-issue voting models to account for non-uniform and individualized issue importance. We demonstrated that determining whether an IWM is a Condorcet winner is co-NP hard, but provided an efficiently checkable condition under which Ostrogorski's paradox does not occur. We then examined instances where an IWM loses to the opposing proposal (i.e., Anscombe's paradox occurs) and showed how our two weighting models alter our ability to reconcile the two objectives. While we now have a rich understanding of the interaction of these two majoritarian ideals, one could optimize for different notions of representation in the proposal selection. It would be interesting to study variants of maximizing total voter "satisfaction" — the total weight voters have on topics that they agree with the final proposal on (a weighted version of an objective proposed in [20]). On the technical side, our work leaves open a number of interesting questions and gaps: (i) Our Theorem 11 for external weights is only existential. In contrast, in the unweighted setting, [13] also provide a polynomial-time method to derandomize the probabilistic argument. Extending this approach to the weighted setting appears generally more challenging but likely feasible in pseudo-polynomial time with slightly more involved techniques. (ii) Paper [13] also shows a hardness result for the unweighted case: telling whether a proposal achieves more agreement with an IWM than guaranteed by the probabilistic argument is NP-hard. It would be interesting to get a similar result for every fixed weights vector w. (iii) We have only succeeded in proving that our bounds in Theorem 12 are tight for some portion of the range $\tilde{w}_{max} \in (0, 1)$.

REFERENCES

[1] Gertrude Elizabeth Margaret Anscombe. 1976. On frustration of the majority by fulfilment of the majority's will. *Analysis* 36, 4 (1976), 161–168.

[2] Carmel Baharav, Andrei Constantinescu, and Roger Wattenhofer. 2025. Condorcet Winners and Anscombes Paradox Under Weighted Binary Voting. arXiv:2502.14639 [cs.GT] https://arxiv.org/abs/2502.14639

[3] Christine Benesch, Rino Heim, Mark Schelker, and Lukas Schmid. 2023. Do Voting Advice Applications Change Political Behavior? *The Journal of Politics* 85, 2 (2023), 684–700.

[4] Jean-Pierre Benoît and Lewis A Kornhauser. 2010. Only a dictatorship is efficient. *Games and Economic Behavior* 70, 2 (2010), 261–270.

[5] Sebastian Bensland, Karim Saleh, and Pietro Ronchetti. 2020. A Machine Learning Analysis of the Swiss Political Spectrum and Candidate Recommendation Process. ETH Zürich student project.

[6] Kellogg S. Booth and George S. Lueker. 1976. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *J. Comput. System Sci.* 13, 3 (1976), 335–379. https://doi.org/10.1016/S0022-0000(76)80045-1

[7] Steven J Brams and Peter C Fishburn. 1978. Approval voting. *American Political Science Review* 72, 3 (1978), 831–847.

[8] Steven J Brams and D Marc Kilgour. 2015. Satisfaction approval voting. *Mathematical and Computational Modeling: With Applications in Natural and Social Sciences, Engineering, and the Arts* (2015), 273–298.

[9] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.

[10] Robert Bredereck, Jiehua Chen, and Gerhard J. Woeginger. 2013. A characterization of the single-crossing domain. *Social Choice and Welfare* 41, 4 (2013), 989–998. http://www.jstor.org/stable/42001441

[11] Harry Brighouse and Marc Fleurbaey. 2010. Democracy and Proportionality. *The Journal of Political Philosophy* 18, 2 (2010), 137–155.

[12] Alessandra Casella. 2005. Storable votes. *Games and Economic Behavior* 51, 2 (2005), 391–419.

[13] Andrei Constantinescu and Roger Wattenhofer. 2023. Computing the Best Policy That Survives a Vote. *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS '23)* (2023).

[14] Edith Elkind, Abheek Ghosh, and Paul W. Goldberg. 2022. Complexity of Deliberative Coalition Formation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 4975–4982.

[15] Edith Elkind and Martin Lackner. 2015. Structure in dichotomous preferences. In *Proceedings of the 24th International Conference on Artificial Intelligence* (Buenos Aires, Argentina) *(IJCAI'15)*. AAAI Press, 2019–2025.

[16] Edith Elkind, Martin Lackner, and Dominik Peters. 2022. Preference Restrictions in Computational Social Choice: A Survey. arXiv:2205.09092 [cs.GT] https://arxiv.org/abs/2205.09092

[17] Simon D. Fink, Matthias Pfretzschner, and Ignaz Rutter. 2023. Experimental Comparison of PC-Trees and PQ-Trees. *ACM J. Exp. Algorithmics* 28, Article 1.10 (oct 2023). https://doi.org/10.1145/3611653

[18] Peter C Fishburn. 1981. An analysis of simple voting systems for electing committees. *SIAM J. Appl. Math.* 41, 3 (1981), 499–502.

[19] Bailey Flanigan, Ariel D Procaccia, and Sven Wang. 2023. Accounting for Stakes in Democratic Decisions. (2023). FORC'23 (non-archival).

[20] Robin Fritsch and Roger Wattenhofer. 2022. The Price of Majority Support. *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS '22)* (2022).

[21] Umberto Grandi. 2014. The common structure of paradoxes in aggregation theory. *arXiv preprint arXiv:1406.2855* (2014).

[22] Lewis A Kornhauser and Lawrence G Sager. 1993. The one and the many: Adjudication in collegial courts. *Calif. L. Rev.* 81 (1993), 1.

[23] Martin Lackner and Piotr Skowron. 2022. Approval-based committee voting. In *Multi-Winner Voting with Approval Preferences*. Springer, 1–7.

[24] Gilbert Laffond and Jean Lainé. 2006. Single-switch preferences and the Ostrogorski paradox. *Mathematical Social Sciences* 52, 1 (2006), 49–66.

[25] Gilbert Laffond and Jean Lainé. 2013. Unanimity and the Anscombe's paradox. *Top* 21 (2013), 590–611.

[26] Steven P Lalley and E Glen Weyl. 2018. Quadratic voting: How mechanism design can radicalize democracy. In *AEA Papers and Proceedings*, Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 33–37.

[27] Jérôme Lang. 2004. Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence* 42 (2004), 37–71.

[28] Christian List. 2003. A possibility theorem on aggregation over multiple interconnected propositions. *Mathematical Social Sciences* 45, 1 (2003), 1–13.

[29] Christian List. 2012. The theory of judgment aggregation: an introductory review. *Synthese* 187, 1 (2012), 179–207.

[30] Hannu Nurmi. 1999. Compound Majority Paradoxes. In *Voting Paradoxes and How to Deal with Them*. Springer Berlin Heidelberg, 70–86.

[31] Ipek Özkal-Sanver and M Remzi Sanver. 2006. Ensuring pareto optimality by referendum voting. *Social Choice and Welfare* 27 (2006), 211–219.

[32] Douglas W Rae and Hans Daudt. 1976. The Ostrogorski paradox: a peculiarity of compound majority decision. *European Journal of Political Research* 4, 4 (1976), 391–398.

[33] Pew Research. 2023. Inflation, Health Costs, Partisan Cooperation Among the Nation's Top Problems. https://www.pewresearch.org/politics/2023/06/21/inflation-health-costs-partisan-cooperation-among-the-nations-top-problems/ Last accessed 19 September 2023.

[34] Zoi Terzopoulou, Alexander Karpov, and Svetlana Obraztsova. 2021. Restricted Domains of Dichotomous Preferences with Possibly Incomplete Information. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 6 (May 2021), 5726–5733. https://doi.org/10.1609/aaai.v35i6.16718

[35] Joel Uckelman. 2009. *More than the sum of its parts: compact preference representation over combinatorial domains*. University of Amsterdam.

[36] Carl Wagner. 1983. Anscombe's paradox and the rule of three-fourths. *Theory and Decision* 15, 3 (1983), 303.

[37] Carl Wagner. 1984. Avoiding Anscombe's paradox. *Theory and Decision* 16, 3 (1984), 233.