

An Analysis Framework for Metric Voting based on LP Duality

David Kempe
University of Southern California

December 17, 2019

Abstract

Distortion-based analysis has established itself as a fruitful framework for comparing voting mechanisms. The assumption is that the m voters and n candidates are jointly embedded in an (unknown) metric space, and the voters submit rankings of candidates by non-decreasing distance from themselves. Based on the submitted rankings, the social choice rule chooses a winning candidate; the quality of the winner is the sum of the (unknown) distances to the voters. Since it is missing the information about the actual distances, the rule’s choice will in general be suboptimal, and the worst-case ratio between the cost of its chosen candidate and the optimal candidate is called the rule’s *distortion*. It was shown in prior work that every deterministic rule has distortion at least 3, while the Copeland rule and related rules guarantee distortion at most 5, and a very recent result gave a generalization of Copeland with distortion $2 + \sqrt{5} \approx 4.236$.

We provide a framework based on LP-duality and flow interpretations of the dual which provides a simpler and more unified way for proving upper bounds on the distortion of social choice rules. Rather than having to reason about all possible metric spaces, to establish an upper bound, it is sufficient to exhibit a certain type of flow with small cost. We illustrate the utility of this approach with three examples. First, we give a fairly simple proof of a strong generalization of the upper bound of 5 on the distortion of Copeland, to social choice rules with short paths from the winning candidate to the optimal candidate in generalized weak preference graphs. A special case of this result recovers the recent $2 + \sqrt{5}$ guarantee. Next, we use this generalization to show that the Ranked Pairs and Schulze rules have distortion $\Theta(\sqrt{n})$. Finally, our framework naturally suggests a combinatorial rule that is a strong candidate for achieving distortion 3, which had also been proposed in recent work. We prove that the distortion bound of 3 would follow from any of three combinatorial conjectures we formulate (and have verified by computer for $n \leq 7$ candidates).

1 Introduction

Voting is an important and widespread way for a group to choose one out of multiple available candidate options.¹ The group could be a country, academic department, or other organization, and the n candidate options they choose from could be courses of action or human candidates. Typically, each voter submits a total order of all options, called a *ranking* or *preference order*. Based on all the submitted rankings, a *social choice rule* (or *mechanism*) determines the winning option.

Different mechanisms will have different desirable and undesirable properties, and it is important to articulate and analyze these properties to guide an organization’s choice of mechanisms. The axiomatic approach, dating back at least several centuries [17, 18], articulates natural axioms about

¹In this submission, we do not consider the equally important and widely studied problem of a group ranking all of the available options.

the properties that the mapping from rankings to a winner should satisfy, and has led to extensive work (see, e.g., [12] for an overview). Unfortunately, many of the key results are impossibility results, in particular the famous Gibbard-Satterthwaite Theorem [21, 31] showing that there is no truthful mechanism satisfying very minimal additional properties.

An alternative that has gained much recent popularity, in particular in the computer science community, is to view social choice through the lens of *optimization and approximation*. In this line of work (e.g., [10, 14, 29, 30]), it is assumed that one can quantify the utility (or cost) that a voter derives from a candidate. These individual utilities or costs can then be aggregated into a social welfare or cost, e.g., by taking the average or median. The social welfare/cost captures how good of a choice a candidate is for the voter population overall.

The problem with this approach, articulated clearly in [11, 3], is that voting mechanisms typically allow voters only to communicate a *ranking* of candidates, but not the actual utilities/costs; furthermore, even if a mechanism provided a way to communicate numerical scores, it is not clear that voters could compute or estimate them accurately. In other words, “one can quantify” is more of an abstract statement than one referring to any decision maker involved in the process. Thus, even though the voting mechanism must optimize a *cardinal* objective function, it only receives *ordinal* information as input, namely, for each voter, whether her² utility/cost for candidate x is larger or smaller than that for candidate y .

As a result, mechanisms must optimize the social welfare *robustly*, choosing a candidate that has high welfare regardless of what the actual cardinal objective values are — so long as they are consistent with the reported ordinal rankings. The *distortion* of a mechanism is the worst-case ratio between the welfare/cost of the mechanism’s selected (based only on ordinal information) candidate and the optimum (with full knowledge of the cardinal values) candidate, over all possible inputs. (Formal definitions of this concept and all other terms can be found in Section 2.)

Our discussion so far has been in terms of general utilities/cost. While some positive results can be obtained for fairly general classes of utility functions (e.g., [10, 14, 29, 30]), stronger results are achievable when the functions take more specific forms. A particularly natural way of defining costs is in terms of a joint *metric space* defined on candidates and voters, where the distance $d(v, x)$ between voter v and candidate x captures their difference in opinion, and hence the cost. Voters then rank candidates by non-decreasing distance from themselves.³ The approach of using the distances explicitly as the cost objective for optimization was proposed in [3]; [2] is an expanded/improved journal version, and [1] provides a broader overview of the area and its results. While [2] consider both the average and median of all voters’ costs as the overall objective, here, we focus solely on the average/total cost.

The main result of [3, 2] is that under the model of metric costs, many widely used voting rules (including Plurality, Veto, Borda count, and others) have distortion linear in the number of candidates or worse. Furthermore, even with just 2 candidates and a 1-dimensional metric space, *every* deterministic voting mechanism has distortion at least 3. On the positive side, [3, 2] show that any rule which always outputs a candidate from the uncovered set of candidates has distortion at most 5, for all metric spaces and numbers of candidates. Uncovered sets are defined in terms of the tournament graph G on n candidates in which the directed edge (x, y) is present iff a (weak) majority of voters prefer x to y . The uncovered set is the set of candidates that have a directed path of length at most 2 in G to every other candidate (see [27]). Very recently, Munagala and Wang [28] gave a voting rule based on uncovered sets in a weighted tournament graph which improves

² For ease of presentation, we use female pronouns for voters and male pronouns for candidates throughout.

³Such distance-based rankings had been considered in a large body of earlier work, e.g., [8, 9, 19, 26, 25, 7, 6], though most of the listed papers studied such rankings specifically when the metric is the line; such preference orders are often called *single-peaked*.

the upper bound from 5 to $2 + \sqrt{5} \approx 4.236$.

There is an obvious gap between the lower bound of 3 for the distortion of every mechanism, and the upper bound of $2 + \sqrt{5}$. In the original version of [3], it was conjectured that a mechanism called *Ranked Pairs* (defined in Section 2) achieves a distortion of 3. This conjecture was disproved by [22], who showed a lower bound of 5 on the distortion of Ranked Pairs (and the Schulze rule, also defined in Section 2).

The proof of the upper bound of 5, the recent upper bound of $2 + \sqrt{5}$, and many other proofs in the literature are based on reasoning about all metric spaces that are consistent with assumed rankings. They often involve intricate case distinctions and rather ad hoc arguments. So far, a more solid foundation and framework for distortion proofs has been missing from the literature.

1.1 Our Contribution

Our main contribution, presented in Section 3, is an analysis framework based on LP duality and flows for proving upper bounds on the metric distortion of voting mechanisms. Our point of departure is a well-known linear program for the following problem: given the rankings of all voters, a winning candidate (presumably selected by a mechanism) and an “optimum” candidate, find a metric space maximizing the distortion of this choice; that is, find a metric that makes the selected winner as expensive as possible, subject to the “optimum” candidate having cost 1.⁴ We show that the dual of the cost minimization LP can be interpreted as a flow problem with an unusual objective function. Using this framework, in order to show an upper bound on the metric distortion of a particular mechanism, rather than having to explicitly consider all possible metric spaces, it is enough to exhibit a flow of small cost meeting certain demands. We illustrate the power of this analysis framework with three applications.

First, in Section 4, we give a strong generalization of the key lemmas from [3] (Theorem 7) and [28] (Lemma 3.7), used to prove distortions of 5 and $2 + \sqrt{5}$ for the respective mechanisms under consideration. The common idea of both is that when a large enough fraction of voters prefer x to y , and a large enough fraction prefer y to z , then the cost of x can be bounded in terms of the cost of z . Theorem 7 of [3] is the special case where both fractions are $\frac{1}{2}$, while Lemma 3.7 of [28] is the case when the first fraction is $\frac{3-\sqrt{5}}{2}$, and the second is $\frac{\sqrt{5}-1}{2}$. These bounds immediately imply the upper bounds on the distortion for any candidate in the uncovered set of a suitably defined tournament graph. We give a generalization to arbitrary chains of preferences, and upper-bound the cost of x_1 in terms of the cost of x_ℓ when a p_i fraction of voters prefer x_i over x_{i+1} , for each $i = 1, \dots, \ell - 1$. For the specific case when all $p_i = p$, the bound can be stated very cleanly: the cost of x_1 is at most $\frac{\ell}{p} - 1$ times that of x_ℓ if ℓ is even, and at most $\frac{\ell-1}{p} + 1$ times that of x_ℓ if ℓ is odd. Our results fully recover and generalize the bounds of [3] and [28]. The generalization to longer path lengths can be useful in analyzing voting mechanisms that are missing information. This can happen if the environment restricts the communication between voters and the mechanism, so that parts of the rankings remain unknown, as in [24]. In fact, the results of Section 4 can be used to significantly improve the upper bounds on the performance of “Copeland-like” mechanisms with missing information, compared to the bounds in [24].

As a direct application of this generalized bound, in Section 5, we resolve the distortion of the Ranked Pairs and Schulze rules (defined in Section 2): we show that both have distortion $\Theta(\sqrt{n})$. The upper bound is a clean application of the lemma bounding distortion via longer chains of preferences, while the lower bound is obtained with a generalization of the example which

⁴This approach can of course immediately be leveraged into an optimal polynomial-time voting mechanism; we discuss this more in Section 3.1.

[22] used to lower-bound the distortion of both rules by 5. The distortion of both rules is thus significantly higher than the distortions of 5 and $2 + \sqrt{5}$ achieved by the uncovered set mechanisms. Understanding the distortion of the Schulze rule in particular is of importance because it is widely used in practice.

As a third application, the flow interpretation naturally suggests a candidate mechanism that might achieve distortion 3, which we present in Section 6. The analysis points to a sufficient condition for distortion 3: that for every given preference profile of the voters, there be a candidate x such that for all other candidates y , a certain bipartite graph on the voters have a perfect matching. In fact, the mechanism itself can be phrased in this terminology, leading to a purely combinatorial polynomial-time mechanism.

This mechanism was independently discovered and presented in [28]. In [28], it is also shown — again with a case distinction proof over metric spaces — that if such a candidate x exists, the mechanism guarantees distortion 3. Our duality framework gives a cleaner and simpler proof of this fact. The main question is then whether the desired candidate x always exists.

Munagala and Wang [28] conjecture — as do we — that it does. They phrase a conjecture which is essentially a restatement of the fact that the algorithm succeeds in finding a candidate x . In Section 6, we present a slight rephrasing of this conjecture, along with two more very different-looking (in fact, much more self-contained) conjectures, each of which would resolve the question positively, i.e., establish a distortion of 3. One of the two new conjectures is phrased in terms of certain preferences between candidates and sets under randomly drawn preference orders, while another talks about cycles in certain induced subsets of a type of graph we define. The fact that they are sufficient to establish distortion 3 is based on Hall’s Marriage Theorem for bipartite graphs. We have verified the conjecture by hand for $n \leq 4$ candidates, and using exhaustive computer search for $n \leq 7$. Resolving any of the three conjectures positively would answer the key open question of the field of metric voting, closing the gap between the upper bound of $2 + \sqrt{5}$ on the best distortion of any deterministic mechanism, and the lower bound of 3.

1.2 Additional Related Work

The observation that mechanisms may have to optimize a cardinal objective function while only given ordinal information (i.e., rankings) extends beyond just voting mechanisms, to more general problems. See, e.g., [5, 1] for results on other optimization problems under ordinal information.

The lower bound of 3 on the distortion of any mechanism is based on worst-case input instances. Better bounds can be obtained when additional assumptions are placed on the instances. As one example, [4, 23] show that when instances are *decisive*, in the sense that each voter has a candidate she strongly prefers over all others, better upper bounds on the distortion are obtained. As another example, when the candidates are drawn i.i.d. from the set of all voters, [15] gives improved constant distortion bounds in the case of two candidates, while [16] shows that many position-based scoring rules now achieve constant distortion (instead of linear).

The lower bound of 3 on the distortion of voting mechanisms only applies to deterministic mechanisms. Randomization can lead to lower distortion [4]. For example, it is known that the Randomized Dictatorship mechanism, which outputs the first choice of a uniformly random voter, has distortion strictly smaller than 3.

Our work ignores the issue of incentives, i.e., whether voters truthfully report their preferences. The connection between strategy proofness and distortion in metric voting is studied in [20].

The use of LP duality for analyzing the performance of optimization algorithms has a long history, e.g., in approximation algorithms (see [34]). Another more recent example is the duality framework of Cai, Devanur, and Weinberg [13] (see also references in [13] to prior, less general,

work) for analyzing the revenue of Bayesian Incentive Compatible mechanisms. In their case as well, dual solutions can be interpreted as flows, and Cai et al. obtain performance guarantees by exhibiting particular types of “canonical” flows that can be interpreted as witnesses for the revenue guarantees. While this work and ours have the use of duality, and the interpretation as flows, in common, the specific technical details are very different.

2 Preliminaries

2.1 Voters, Candidates, and Social Choice Rules

An *instance* (X, \mathcal{P}) consists of a set of n candidates X , and the voters’ preferences \mathcal{P} among these candidates. Candidates will always be denoted by lowercase letters w, x, y, z (and their variations), with w specifically reserved for a candidate chosen as winner by a mechanism (which will be clear from the context). Sets of candidates are denoted by uppercase letters X, Y, Z . The m voters are denoted by v, v' and variants thereof, and the set of all voters is V .

Each voter v has a *total order* (or *preference order* or *ranking* — we use the three terms interchangeably) \succ_v over the n candidates. We write $x \succ_v y$ to denote that v (strictly) prefers x over y , and $x \succeq_v y$ to denote that v weakly prefers x over y ; the difference is that the latter allows $x = y$. We extend this notation to sets, writing, for instance, $Y \succ_v Z$ to denote that v (strictly) prefers all candidates in Y over all candidates in Z . We write $[x \succ Y] = \{v \in V \mid x \succ_v Y\}$ for the set of voters who rank x strictly ahead of all candidates in Y , and $[Y \succ x] = \{v \in V \mid Y \succ_v x\}$ for the set of voters who rank x strictly behind all candidates in Y .

A *vote profile* \mathcal{P} is the vector of the rankings of all voters $\mathcal{P} = (\succ_v)_{v \in V}$. A *social choice rule* (we use the term *mechanism* interchangeably) $f : (X, \mathcal{P}) \mapsto w$ is given the rankings of all voters, i.e., \mathcal{P} , and deterministically produces as output one *winning* candidate $w = f(X, \mathcal{P}) \in X$.

2.2 (Pseudo-)Metric Space and Distortion

The voter preferences are assumed to be derived from distances between voters and candidates. The distance $d(v, x)$ between voter v and candidate x captures how similar their positions on key issues are. The distances d form a *pseudo-metric*, i.e., they are non-negative and satisfy the triangle inequality⁵ $d(v, x) \leq d(v, y) + d(v', y) + d(v', x)$ for all voters v, v' and candidates x, y .

A vote profile \mathcal{P} is *consistent* with the pseudo-metric d if and only if each voter ranks the candidates by non-decreasing distance from herself; that is, if $x \succ_v y$ whenever $d(v, x) < d(v, y)$. When \mathcal{P} is consistent with d , we write $d \sim \mathcal{P}$. If there are ties among distances, several vote profiles will be consistent with d .

Definition 2.1 (Social Cost, Distortion) 1. The social cost of candidate x is the sum of distances from x to all voters: $C(x) = \sum_v d(v, x)$.

2. A candidate is an optimum candidate iff he minimizes⁶ the social cost: $x_d^* \in \operatorname{argmin}_{x \in X} C(x)$.

3. The distortion of a mechanism f is the largest possible ratio between the cost of the candidate chosen by f , and the optimal (with respect to the pseudo-metric d , which f does not know)

⁵Distances between pairs of voters, or between pairs of candidates, could be defined using shortest-path distances; however, they are irrelevant for our analysis. *Symmetry*, another defining property of a pseudo-metric, would arise automatically when using this definition.

⁶There could be multiple optimum candidates — for our analysis, it will never matter which of them is designated as “the” optimum.

candidate x_d^* :

$$\rho(f) = \max_{\mathcal{P}} \sup_{d: d \sim \mathcal{P}} \frac{C(f(X, \mathcal{P}))}{C(x_d^*)}.$$

The main cause for (large) distortion is that while the social choice rule knows the voter preferences \mathcal{P} , it does not know the pseudo-metric d . We can think of the pseudo-metric d as being chosen adversarially, based on the winning candidate $w = f(\mathcal{P})$ chosen by the mechanism. However, the adversary is constrained by having to ensure that d is consistent with \mathcal{P} .

2.3 Ranked Pairs and the Schulze Rule

In Section 5, we will characterize the distortion of the Ranked Pairs and Schulze Rules. We briefly review these rules here. Both are based on a weighted directed graph on the set of candidates X . The weight $p_{x,y}$ of the edge from candidate x to y is the fraction of voters who have $x \succ y$. As a result, $p_{x,y} + p_{y,x} = 1$ for all x, y .

In Ranked Pairs [33], the (ordered) pairs (x, y) are considered in non-increasing order of $p_{x,y}$. When the pair (x, y) is considered, the directed edge (x, y) is inserted into the graph if and only if doing so creates no cycle. When the insertion process terminates, the graph has a unique source node, which is returned as the winner.

In the Schulze Method [32], a directed weighted graph is created in which each ordered pair (x, y) has an edge with weight $p_{x,y}$. Then, for each pair (x, y) , let $s_{x,y}$ be the width of the widest path from x to y , that is, the largest p such that there is a path from x to y on which all edges (x', y') have $p_{x',y'} \geq p$. It has been shown [32] that there is a candidate node x such that $s_{x,y} \geq s_{y,x}$ for all other candidates y . Any such candidate x is returned as the winner.

For the purposes of our analysis, the only property of these methods that matters is captured by the following lemma, which is well known. (We prove it only for completeness.)

Lemma 2.2 *Let w be the candidate selected by the rule (either Ranked Pairs or Schulze), and y any other candidate. Then, there exists a p and a sequence of (distinct) candidates $x_1 = w, x_2, \dots, x_\ell = y$ with the property that at least a p fraction of voters prefer x_i over x_{i+1} (for each i), and at most a p fraction of voters prefer y over w .*

Proof. For the Ranked Pairs rule, because w was selected, it has no incoming edges in the DAG that is constructed. In particular, this means that Ranked Pairs did not insert the edge (y, w) , so when it was considered for insertion, it would have caused a cycle, meaning that there was a path from w to y all of whose edges had been inserted earlier. By the definition of the Ranked Pairs insertion order, this means that all edges on this path had a higher fraction of voters agreeing with them, giving us the path claimed above.

For the Schulze rule, recall that the winner has the property that $s_{w,x} \geq s_{x,w}$ for all candidates x . Let $p = s_{w,y}$. Then, there is a path from w to y in which each edge corresponds to a preference by at least a p fraction of voters. On the other hand, because (y, w) is a path from y to w , at most an $s_{y,w} \leq s_{w,y}$ fraction of voters can prefer y to w . ■

3 The LP Duality Approach and Flows

In this section, we develop the key tool for our analysis: the dual linear program for distortion in metric voting.

The voters' preferences $\mathcal{P} = (\succ_v)_v$ are given. Let w be a candidate that the mechanism is considering as a potential winner, and x^* the optimal candidate. Following [2, 22], we phrase the

adversary’s problem of finding the distortion-maximizing metric as a linear program whose variables $d_{v,x}$ denote distances between voters v and candidates x . These distances must be non-negative, obey the triangle inequality, and be consistent with the reported preferences of the voters. The objective is to maximize the distortion, i.e., the ratio between the cost of w and the cost of x^* .

$$\begin{array}{ll}
\text{Maximize} & \sum_v d_{v,w} \\
\text{subject to} & d_{v,x} \leq d_{v',x} + d_{v',y} + d_{v,y} \quad \text{for all } x, y, v, v' \text{ (Triangle Inequality)} \\
& d_{v,x} \leq d_{v,y} \quad \text{for all } x, y, v \text{ such that } x \succ_v y \text{ (consistency)} \\
& \sum_v d_{v,x^*} = 1 \quad \text{(normalization)} \\
& \sum_v d_{v,x} \geq 1 \quad \text{for all } x \text{ (optimality of } x^*) \\
& d_{v,x} \geq 0 \quad \text{for all } x, v.
\end{array} \tag{1}$$

As is standard in the use of LPs for optimizing a ratio, the normalization side-steps the issue of having to write a ratio: for any worst-case metric, one could simply rescale all distances by a constant so that the normalization holds — this does not change any ratios, and thus also not the distortion.

3.1 An Efficient Optimal Mechanism

As already observed in [2, 22], the LP (1) can be leveraged to immediately yield an instance-optimal polynomial-time mechanism for minimizing distortion, as follows. Given the voter preferences \succ_v , let c_{w,x^*} denote the maximum LP objective of the LP (1) for the winner w and putative optimum x^* . The distortion for w as a winner is then $\bar{c}_w = \max_{x^*} c_{w,x^*}$. The mechanism returns as winner any candidate in $\text{argmin}_w \bar{c}_w$.

Because the algorithm only involves solving n^2 linear programs, it runs in polynomial time. By definition (and correctness of the LP (1)), the distortion for a given vote profile \mathcal{P} and winner w is \bar{c}_w ; thus, the mechanism does indeed minimize distortion. Unfortunately, as also observed in [2], it is not immediate from the mechanism and the LP formulation how to bound the distortion for all vote profiles; though [22] conjecture that the LP-based algorithm guarantees distortion at most 3.

The dual program provides a very useful tool towards making the LP-based algorithm combinatorial, and for reducing an analysis of its distortion to simpler combinatorial conjectures. More generally (and perhaps importantly), the dual program provides a general approach for bounding the metric distortion of other voting rules, too.

3.2 The Dual Linear Program

Rearranging the primal LP into normal form, taking the dual, and switching the signs of the α_x variables (for clarity) yields the following dual LP (2). In this LP, the $\psi_{x,y}^{(v,v')}$ are the dual variables for the triangle inequality constraints, $\phi_{x,y}^{(v)}$ are the dual variables for the consistency constraints, and the α_x are the dual variables for the normalization/optimality constraints.

$$\begin{aligned}
& \text{Minimize} && \sum_x \alpha_x \\
& \text{subject to} && \alpha_x + \sum_{y: x \succ_v y} \phi_{x,y}^{(v)} - \sum_{y: y \succ_v x} \phi_{y,x}^{(v)} + \sum_{y,v'} \left(\psi_{x,y}^{(v,v')} - \psi_{y,x}^{(v,v')} - \psi_{x,y}^{(v',v)} - \psi_{y,x}^{(v',v)} \right) \\
& && \geq \begin{cases} 1 & \text{if } x = w \\ 0 & \text{if } x \neq w \end{cases} \quad \text{for all } v, x \\
& && \psi_{x,y}^{(v,v')} \geq 0 \quad \text{for all } v, v', x, y \\
& && \phi_{x,y}^{(v)} \geq 0 \quad \text{for all } v, x, y \\
& && \alpha_x \leq 0 \quad \text{for all } x \neq x^*.
\end{aligned} \tag{2}$$

Notice that α_{x^*} is in fact unconstrained, due to the equality constraint in the normalization.

The advantage of studying the dual linear program instead of the primal (or reasoning about the distortion directly) is that it omits any reference to any metric space. Rather than having to reason about all candidate metric spaces consistent with given voting patterns, by weak duality, we only have to exhibit one setting of the dual variables that yields a small dual objective value. Thus, our goal in analyzing a mechanism will be to show that for any voter preferences \mathcal{P} , with a suitably chosen winner w , there is a setting of dual variables giving a small objective value.

3.3 Using the Dual by Exhibiting Flows

The LP (2) looks rather unwieldy, mostly due to the terms involving the $\psi_{x,y}^{(v,v')}$ variables. However, by making some specific choices for these variables, it can be interpreted as a flow⁷ problem on a suitably defined graph, with a somewhat unusual objective function. This is captured by the following lemma:

Lemma 3.1 *Let $H = (U, E)$ be a directed graph with vertex set $U = V \times X$, and edges defined as follows:*

- *Whenever $x \succ_v y$, E contains the directed edge $(v, x) \rightarrow (v, y)$. We call such edges preference edges.*
- *For all x and $v \neq v'$, E contains the directed edge $(v, x) \rightarrow (v', x)$. We call such edges sideways edges.*

Let f be a flow on H such that exactly one unit of flow originates at the node (v, w) for each voter v , and flow is only absorbed at nodes (v, x^) for voters v . Define the cost of f at voter v to be $\gamma_v^{(f)} = \sum_{e \text{ into } (v, x^*)} f_e + \sum_{x \neq x^*} \sum_{v' \neq v} (f_{(v', x) \rightarrow (v, x)} + f_{(v, x) \rightarrow (v', x)})$.*

Then, $C(w) \leq C(x^) \cdot \max_v \gamma_v^{(f)}$.*

The graph H has two types of edges. For any fixed voter v , the preference edges $(v, x) \rightarrow (v, y)$ (over all candidate pairs x, y) exactly correspond to v 's preference order. For any fixed candidate x , the sideways edges $(v, x) \rightarrow (v', x)$ (over all voter pairs v, v') form a complete directed graph.

The flow's cost function has two terms for each voter v . The first is fairly standard in the study of multi-commodity flows: the capacity required at the sink node (v, x^*) to be able to absorb all of the flow. The second one is rather non-standard: for each voter v , there is an additional penalty

⁷Some sources use the word "flow" only when there is a single source and a single sink; here, we will have multiple sources and sinks. We will still use the word "flow" in a generic sense.

term for all incoming and outgoing flows of nodes (v, x) for $x \neq x^*$ along sideways edges. In other words, using preference edges is much less costly than using sideways edges: the former just route flow, while the latter route the flow, but also incur a cost penalty at both endpoints.

Proof. Let f be a flow with one unit of flow originating at each node (v, w) , such that flow is only absorbed at nodes (v, x^*) . We define dual variables, and show that these dual variables are feasible. Then, we will obtain the statement of the lemma by weak LP duality.

For each triple v, x, x' , we set $\phi_{x, x'}^{(v)} = f_{(v, x) \rightarrow (v, x')}$. For each triple v, v', x , we set $\psi_{x, x^*}^{(v, v')} = f_{(v, x) \rightarrow (v', x)}$; notice that we carefully choose x^* as the additional candidate for the dual variable. Finally, we set $\alpha_{x^*} = \max_v \gamma_v^{(f)}$. All other dual variables are set to 0.

We now verify that this assignment satisfies all dual constraints. First, because $\psi_{x, y}^{(v, v')} = 0$ and $\alpha_y = 0$ whenever $y \neq x^*$, the dual constraints for all $x \neq x^*$ are exactly circulation constraints; that is, they require that (at least) one unit of flow originate with (v, w) , and that flow be conserved (or appear) at each node $x \neq x^*, x \neq w$. Thus, all of these constraints are satisfied for the given dual variable assignment.

For pairs (v, x^*) , the left-hand side of the dual constraint totals the flow into (v, x^*) along any edges (these are the $\phi_{x, x^*}^{(v)}$ variables and the $\psi_{x^*, x^*}^{(v', v)}$ variables), as well as all the $\psi_{x, x^*}^{(v, v')}$ and $\psi_{x, x^*}^{(v', v)}$ variables, for all $v' \neq v$. By definition of the dual variables, this is exactly the flow into (v, x^*) , plus the flow into and out of all nodes (v, x) for $x \neq x^*$ along edges of the form $(v, x) \rightarrow (v', x)$ and $(v', x) \rightarrow (v, x)$. Thus, it is exactly $\gamma_v^{(f)}$. Because we set $\alpha_{x^*} = \max_v \gamma_v^{(f)}$, the dual constraints for all pairs (v, x^*) are also satisfied.

Since we have a dual feasible solution of objective value $\alpha_{x^*} = \max_v \gamma_v^{(f)}$, by weak duality, for every metric, the cost of the primal is at most $\max_v \gamma_v^{(f)}$. This completes the proof. ■

4 Generalization of Distortion Bounds for Undominated Nodes

As a corollary of Lemma 3.1, we obtain a strong generalization of Theorem 7 in [3] and Lemma 3.7 of [28] (which are given below for comparison). The most general version can be stated as follows:

Corollary 4.1 *Let x_1, x_2, \dots, x_ℓ be (distinct) candidates such that for each $i = 2, \dots, \ell$, at least a $p_i > 0$ fraction of voters prefer candidate x_{i-1} over candidate x_i . Define $\lambda_1 = 1, \lambda_2 = \frac{2}{p_2} - 1$, and $\lambda_i = \frac{2}{p_i}$ for $2 < i \leq \ell$. Let $\Lambda = \max_{S \subseteq \{1, \dots, \ell\}, S \text{ indep.}} \sum_{i \in S} \lambda_i$. (Here, independence of a set S of natural numbers means that the set contains no two consecutive numbers.) Then, $C(x_1) \leq \Lambda \cdot C(x_\ell)$.*

Proof. We define a flow f and analyze its cost. For each i , we call the nodes (v, x_i) (for all voters v) layer i . Let V_i be the set of voters v with $x_{i-1} \succ_v x_i$, with $V_1 := V$ for notational simplicity.

We construct the flow layer by layer; our construction will ensure that each node (v, x_i) with $v \in V_i$ has exactly $\frac{m}{|V_i|}$ units of flow entering. This holds in the base case $i = 1$, because each node in layer 1 is the source node of one unit of flow.

For the i^{th} step of the construction, we first route all the flow within layer i using sideways edges, from nodes (v, x_i) with $v \in V_i$ to nodes (v', x_i) with $v' \in V_{i+1}$. We then route it to nodes (v', x_{i+1}) in layer $i + 1$ using preference edges. More specifically, to route the flow within layer i , we first consider voters $v \in V_i \cap V_{i+1}$. For those voters, $\min(\frac{m}{|V_i|}, \frac{m}{|V_{i+1}|})$ units of flow simply stay at (v, x_i) . The node (v, x_i) for such v will have additional incoming flow from other nodes (if $\frac{m}{|V_{i+1}|} > \frac{m}{|V_i|}$) or additional outgoing flow to other nodes (if $\frac{m}{|V_{i+1}|} < \frac{m}{|V_i|}$). The remaining flow is routed arbitrarily using sideways edges from nodes (v, x_i) with $v \in V_i$ to nodes (v', x_i) with $v' \in V_{i+1}$, of course ensuring that each such node (v', x_i) has in total $\frac{m}{|V_{i+1}|}$ units of flow entering.

After this redistribution within layer i , each (v, x_i) routes its flow to (v, x_{i+1}) . Notice that this is always possible, because $x_i \succ_v x_{i+1}$ for all $v \in V_{i+1}$. The construction is illustrated with an example in Figure 1.

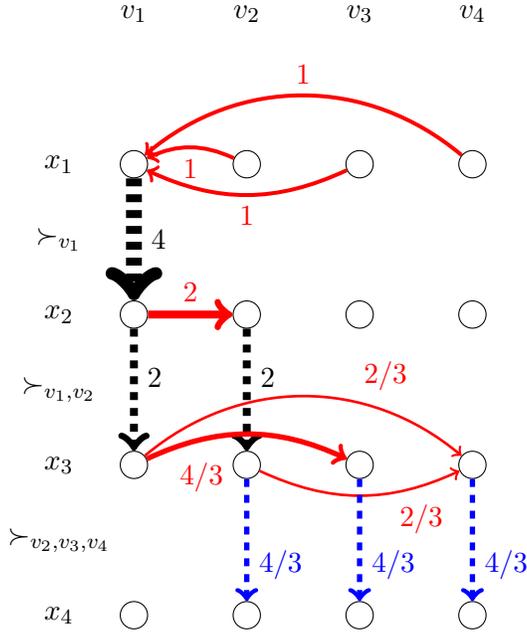


Figure 1: An illustration of the flow construction. In the example, there are 4 voters and 4 relevant candidates, with voter preferences shown on the left. The preference fractions are $p_1 = 1/4, p_2 = 1/2, p_3 = 3/4$. Sideways flows are shown in solid red, while flow along preference edges is shown in dashed lines. The dashed lines into nodes for candidate x_4 are shown in blue (instead of black), to emphasize that they contribute to the objective function. The amount of flow is given numerically, and also shown using the width of the lines/arcs. Edges that are not used by the flow are not shown.

We now analyze the cost associated with any fixed voter v . The cost has two components: the incoming flow at (v, x_ℓ) (shown in blue in Figure 1), and the cost associated with incoming/outgoing flow using sideways edges incident on (v, x_i) for $i < \ell$ (shown in red in Figure 1). We begin with the incoming flow at (v, x_ℓ) : if $v \in V_{\ell-1}$, the incoming flow is $\frac{m}{|V_\ell|} = \frac{1}{p_\ell}$; otherwise, it is 0.

Next, we consider the cost associated with sideways edges. As a general guideline (subtleties will be discussed momentarily), when $v \in V_i$, the node (v, x_i) has $\frac{m}{|V_i|} = \frac{1}{p_i}$ units of flow coming in along sideways edges, and the node (v, x_{i+1}) has the same amount of flow leaving along sideways edges. The associated cost of both together is $\frac{2}{p_i}$. Two obvious exceptions are layers $i = 1$ and $i = \ell - 1$. For $i = 1$, one unit of flow simply originates with (v, x_1) , resulting in no cost. For $i = \ell - 1$, no sideways edge is used to route outgoing flow; however, this is compensated by the incoming flow at (v, x_ℓ) (discussed in the preceding paragraph), which adds the same cost term.

However, simply adding up the bounds from the preceding paragraph over all steps i with $v \in V_i$ is too pessimistic, because our flow construction avoids routing more flow than necessary when $v \in V_i \cap V_{i+1}$. A tighter bound is captured by the following lemma, proved below:

Lemma 4.2 *Let I be the set of all indices i with $v \in V_i$. I can be partitioned into disjoint intervals of integers $\{L_1, L_1 + 1, \dots, R_1\}, \{L_2, L_2 + 1, \dots, R_2\}, \dots, \{L_K, L_K + 1, \dots, R_K\}$ (for some $K \geq 1$)*

such that:

1. For each k , there exists an index $M_k \in \{L_k, \dots, R_k\}$ such that $p_{L_k} \geq p_{L_k+1} \geq \dots \geq p_{M_k} > 0$ and $p_{M_k} \leq p_{M_k+1} \leq \dots \leq p_{R_k}$; that is, the p_i are monotone non-increasing from L_k to M_k , and monotone non-decreasing from M_k to R_k .
2. The total cost of flow (both sideways flow and flow into (v, x_ℓ) in case $R_k = \ell$) associated with nodes (v, x_i) with $L_k \leq i \leq R_k$ is at most λ_{M_k} .

To apply Lemma 4.2, the key observation is that the set $\{M_1, M_2, \dots, M_K\}$ is independent, i.e., contains no two consecutive integers. If it did — say, $i = M_k$ and $i + 1 = M_{k'}$ — then both $i, i + 1 \in I$. If $p_{i+1} \leq p_i$, this would contradict the maximality of i in the definition of M_k ; on the other hand, if $p_{i+1} > p_i$, then $i + 1 \leq R_k$ by the definition of R_k , so it is impossible that $i + 1 = M_{k'}$.

Now, summing up the costs for each of the disjoint intervals, we obtain that the total cost of the flow at nodes associated with v (both sideways flow and flow into (v, x_ℓ)) is at most $\sum_{k=1}^K \lambda_{M_k}$; because the set of M_k is independent, this sum is at most Λ . Using Lemma 3.1, this completes the proof. \blacksquare

Proof of Lemma 4.2. We inductively define L_k, R_k, M_k satisfying the first property, then show that they also guarantee the second property. For the base case, we set (for convenience) $R_0 := 0$. For the inductive step, focus on any $k \geq 1$. Define $L_k = \min\{i \mid i \in I, i > R_{k-1}\}$. (The construction terminates when there is no such i .) Let $M_k = \max\{i \mid \{L_k, \dots, i\} \subseteq I, p_{L_k} \geq p_{L_k+1} \geq \dots \geq p_i\}$. In words, M_k is the largest index i such that all indices between L_k and i are in I , and the p_j values are monotone non-increasing all the way to i . Notice that because $M_k \in I$, we also have $p_{M_k} > 0$. Now, let $R_k = \max\{i \mid \{L_k, \dots, i\} \subseteq I, p_{M_k} \leq p_{M_k+1} \leq \dots \leq p_i\}$. In words, R_k is the largest index i such that all indices between M_k and i (and thus also between L_k and i) are in I , and the p_j values are monotone non-decreasing from M_k to i . This definition explicitly ensures that each interval $\{L_k, \dots, R_k\}$ is entirely contained in I , and satisfies the given monotonicity conditions. We now verify the second property.

We first consider the case $M_k \notin \{1, 2, \ell - 1\}$, where $\lambda_{M_k} = 2/p_{M_k}$. The important observation for the proof, also visible in Figure 1, is that when $v \in V_i \cap V_{i+1}$, this eliminates sideways flow to and from nodes associated with v . Specifically, none of the nodes (v, x_i) for $L_k \leq i < M_k$ have outgoing flow along sideways edges (since all their flow stays for the next step). The incoming flow at (v, x_i) along sideways edges is exactly $1/p_i - 1/p_{i-1}$ (with $1/p_{L_k-1} := 0$ for convenience); the remaining flow at (v, x_i) is what is kept from step $i - 1$. Similarly, none of the nodes (v, x_i) for $M_k < i \leq R_k$ have incoming flow along sideways edges, since the node (v, x_{i-1}) has enough flow to meet all of the needs of (v, x_i) ; the outgoing flow at such nodes (v, x_i) along sideways edges is exactly $1/p_i - 1/p_{i+1}$ (with $1/p_{R_k+1} := 0$). Summing up all the incoming flows for $i = L_k, \dots, M_k$ (a telescoping series), and the outgoing flows for $i = M_k, \dots, R_k$ (another telescoping series), the total flow on sideways edges for all (v, x_i) with $i \in \{L_k, \dots, R_k\}$ is at most $2/p_{M_k} = \lambda_{M_k}$. Because $M_k < \ell$, there is no other cost associated with these nodes. Next, we consider the remaining cases $M_k \in \{1, 2, \ell - 1\}$.

1. If $M_k = 1$, then by construction, $2 \notin I$; otherwise, the fact that $p_1 = 1 \geq p_2$ would rule out setting $M_k = 1$. Thus, the entire interval is just $\{L_1, \dots, R_1\} = \{1\}$. Because there is no incoming sideways flow into (v, x_1) , the only cost is for one unit of outgoing sideways flow, i.e., the cost is $1 = \lambda_1$.
2. If $M_k = 2$, then we must have $k = 1$ and $L_k = 1$, because $1 \in I$ always by definition. We can directly apply the general analysis, except that we can subtract one unit of cost, the reason

being (as in the case $M_k = 1$) that there is no one unit of sideways flow into (v, x_1) . Thus, the total cost associated with the interval $\{L_1, \dots, R_1\}$ is at most $2/p_2 - 1 = \lambda_2$.

3. If $M_k = \ell - 1$, then $k = K$. There is no sideways flow out of $(v, x_{\ell-1})$ (and there are no nodes (v, x_i) for $i > \ell - 1$ to consider). Thus, the total cost of the sideways flows associated with $\{L_K, \dots, R_K\}$ is at most $1/p_{M_k}$. On the other hand, in this case, there is also a cost of $1/p_{M_k}$ for flow into (v, x_ℓ) (the blue flow in Figure 1); however, the total cost is still bounded by $2/p_{M_k} = \lambda_{M_k}$.

This shows that the bound holds for all cases of the interval. ■

4.1 Special Cases

Lemma 3.7 of [28] is the special case of Corollary 4.1 with $\ell = 3, x_1 = w, x_3 = x^*$, and $p_1 = \frac{3-\sqrt{5}}{2}, p_2 = \frac{\sqrt{5}-1}{2}$. Our Corollary 4.1 then exactly recovers the bound of $2 + \sqrt{5}$.

When we have a uniform lower bound on the p_i , Corollary 4.1 can be simplified significantly. (A direct proof of Corollary 4.3 would also be simpler than the proof of the more general Corollary 4.1.)

Corollary 4.3 *Let x_1, x_2, \dots, x_ℓ be (distinct) candidates such that for each $i = 2, \dots, \ell$, at least a $p > 0$ fraction of the voters prefer candidate x_{i-1} over candidate x_i . Then, if ℓ is even, $C(x_1) \leq (\frac{\ell}{p} - 1) \cdot C(x_\ell)$; if ℓ is odd, $C(x_1) \leq (\frac{\ell-1}{p} + 1) \cdot C(x_\ell)$.*

Proof. We substitute $p_i = p$ for all i in Corollary 4.1; then, we observe that for even ℓ , the independent set of integers giving the largest sum is $\{2, 4, \dots, \ell\}$, while for odd ℓ , it is $\{1, 3, \dots, \ell\}$. ■

The asymmetry between even and odd ℓ disappears when $p = \frac{1}{2}$ (i.e., in the case of the majority graph), where the bound simply becomes $2\ell - 1$. The result thus strongly generalizes Theorem 7 in [3], which is the special case of $p = \frac{1}{2}$ and $\ell = 3$.

5 Distortion of Ranked Pairs and Schulze

Corollary 4.3 allows us to pin down the distortion of the Ranked Pairs and Schulze rules to within small constant factors.

Corollary 5.1 *Both the Ranked Pairs mechanism and the Schulze rule asymptotically have distortion at most $2\sqrt{2} \cdot \sqrt{n} + o(\sqrt{n})$ and at least $\frac{\sqrt{2}}{2}\sqrt{n}$.*

Proof. We begin by proving the upper bounds. Let w be the candidate selected by the rule, and x^* the optimum candidate. By Lemma 2.2, applied with $y = x^*$, there exists a p and a sequence of distinct candidates $x_1 = w, x_2, \dots, x_\ell = x^*$ with the property that for each i , at least a p fraction of voters prefer x_i over x_{i+1} , and at most a p fraction of voters prefer x^* over w . The existence of x_1, \dots, x_ℓ with these properties is all that we need from the specific voting rules. The rest of the proof will be completely generic, and would thus also apply to any other voting rule satisfying Lemma 2.2.

We consider two cases, depending on the value of p . The case $p \leq 1 - \frac{1}{\sqrt{2n}}$ is easy. In this case, at least a $1 - p \geq \frac{1}{\sqrt{2n}}$ fraction of voters prefer w over x^* . Lemma 6 from [2] states that if at least a q fraction of voters prefer x over x' , then $C(x) \leq (1 + \frac{2(1-q)}{q}) \cdot C(x')$. Applying this lemma to w and x^* , the distortion of w is at most $\frac{2}{1-p} - 1 \leq 2\sqrt{2} \cdot \sqrt{n}$.

When $p > 1 - \frac{1}{\sqrt{2n}}$, we use Corollary 4.3. Let $\lambda = \lfloor \sqrt{n/2} \rfloor$, and $B = \lceil \ell/\lambda \rceil$. Because $\ell \leq n$, we get that $B \leq \sqrt{2n} + o(\sqrt{n})$. Consider the $B + 1$ candidates $y_j := x_{j\lambda+1}$ for $j = 0, 1, \dots, B - 1$, and $y_B := x_\ell$. Because for each i , at least a $1 - \frac{1}{\sqrt{2n}}$ fraction of voters prefer x_i to x_{i+1} , a union bound over the candidates $x_{j\lambda+1}, x_{j\lambda+2}, \dots, x_{(j+1)\lambda}$ shows that for each $j < B$, at least a $1 - \lambda \cdot \frac{1}{\sqrt{2n}} \geq \frac{1}{2}$ fraction of voters prefer y_j over y_{j+1} . By using Corollary 4.3 (applied with $p = \frac{1}{2}$ and $\ell = B + 1$), we obtain that

$$C(w) = C(y_0) \leq (2B + 1) \cdot C(y_B) \leq (2\sqrt{2} \cdot \sqrt{n} + o(\sqrt{n})) \cdot C(x^*),$$

so the distortion is upper-bounded by $2\sqrt{2} \cdot \sqrt{n} + o(\sqrt{n})$.

We now turn to a lower bound. Our lower-bound construction is a straightforward generalization of the construction that [22] used to show a lower bound of 5 on the distortion of the two rules. Let m be given (assumed even), and set (with foresight⁸) $B = m/2$. Our construction has $m + 2$ voters and $n = mB$ candidates x_1, x_2, \dots, x_n . Voters $v = m + 1, m + 2$ have $x_i \succ_v x_{i+1}$ for all i . We call this the *default order*. To define the order (and later: distances) for voters $v = 1, \dots, m$, we define the following *blocks* of consecutive (according to the default order) candidates. Block $\mathcal{B}_{v,b}$ for $1 \leq b < B$ consists of the m candidates $\{x_{(b-1)m+v}, x_{(b-1)m+v+1}, \dots, x_{bm+v-1}\}$. Block $\mathcal{B}_{v,0}$ consists of the $v - 1$ candidates $\{x_1, \dots, x_{v-1}\}$. (Notice that $\mathcal{B}_{1,0} = \emptyset$.) Finally, block $\mathcal{B}_{v,B}$ consists of the $m + 1 - v$ candidates $\{x_{(B-1)m+v}, x_{(B-1)m+v+1}, \dots, x_n\}$. Voter $v = 1, \dots, m$ ranks the blocks in reverse order $\mathcal{B}_{v,m}, \mathcal{B}_{v,m-1}, \mathcal{B}_{v,m-2}, \dots, \mathcal{B}_{v,0}$; within each block, v ranks the candidates by the default order. An example of the block structure and ordering is shown in Figure 2.

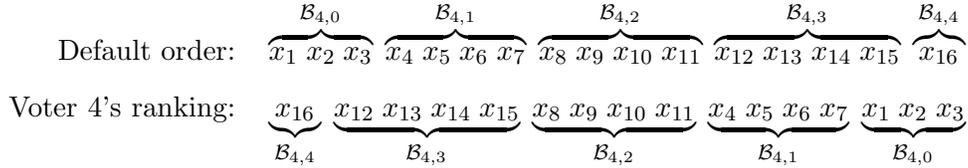


Figure 2: An illustration of the blocks $\mathcal{B}_{4,b}$ and the preference ranking of voter 4. In this example, $m = B = 4$.

Intuitively, this means that on a “global” scale, voters $v = 1, \dots, m$ completely disagree with the default order, but locally, it looks like they agree. More formally, each voter v disagrees with the default order only for pairs (x_i, x_{i+1}) when $i \equiv v - 1 \pmod{m}$. In particular, this means that for every pair (x_i, x_{i+1}) , only one voter disagrees with the default order. For every pair (x_i, x_j) with $i < j$, at least the voters $v = m + 1, m + 2$ have $x_i \succ_v x_j$, so the edges (x_i, x_{i+1}) have highest weight. In particular, this means that they are inserted first by Ranked Pairs, and hence Ranked Pairs chooses $w = x_1$. Similarly, x_1 has a path of width $1 - \frac{1}{m+2}$ to each x_i , exceeding the width for any other node. Thus, the Schulze method also chooses $w = x_1$.

We will now define a metric d which is consistent with these rankings. Voters $v = m + 1, m + 2$ have distance B to each candidate.⁹ Each voter $v = 1, \dots, m$ has distance $d(v, x_i) = 2(B - b) + 1$ from all candidates $x_i \in \mathcal{B}_{v,b}$. First, these distances explicitly ensure consistency with the voters’

⁸The choice $B = m/2$ gives the tightest lower bound for this type of construction. Other choices work as well; for instance, $B = m$ gives a lower bound of $\frac{2}{3}\sqrt{n}$ instead of $\frac{\sqrt{2}}{2}\sqrt{n}$.

⁹To avoid tie breaking issues, one can easily perturb these and other equal distances by small ϵ values so that all distances are unique and the desired order is uniquely induced. We avoid doing so to not overload the proof with inessential notation.

rankings. It remains to verify the triangle inequality. Consider two voters $v \neq v'$ and candidates $x_i \neq x_j$. We need to show that $d(v, x_i) \leq d(v, x_j) + d(v', x_j) + d(v', x_i)$. This is trivial if $v \in \{m+1, m+2\}$, because v has distance B to all candidates. If $v' \in \{m+1, m+2\}$, then the right-hand side contains two terms of B , and one term $d(v, x_j)$ that is at least 1. Hence, the inequality holds. Finally, if both $v, v' \leq m$, then let b, b' be the blocks such that $x_i \in \mathcal{B}_{v,b}, x_j \in \mathcal{B}_{v',b'}$. The definition of the block structure ensures that $|b - b'| \leq 1$; in particular, $d(v, x_i) \leq d(v', x_i) + 2$. Because $d(v, x_j) \geq 1, d(v', x_j) \geq 1$, the triangle inequality again holds.

Finally, we compute the social costs of x_1 (the winner in Schulze and Ranked Pairs) and x_n . x_1 is in block 0 for all voters $v = 2, \dots, m$, and in block 1 for voter 1. Hence, his combined distance from these voters is $(m-1)(2B+1) + (2B-1) = 2Bm + m - 2$. With the added distance of B from each of $v = m+1, m+2$, the total cost of x_1 is $2Bm + 2B + m - 2 \geq 2Bm$. On the other hand, x_n is at distance 1 from all voters $v = 1, \dots, m$ and at distance B from $v = m+1, m+2$, for a total cost of $2B + m$. The ratio is thus at least $\frac{2Bm}{2B+m} \stackrel{B=m/2}{=} \frac{m}{2} \stackrel{n=m^2/2}{=} \frac{\sqrt{2}}{2} \cdot m$. This completes the proof. ■

Remark 5.2 The upper bound in Corollary 5.1 was a direct application of our flow-based framework. While the lower bound did not explicitly use the framework, the counter-example was in fact discovered after failed attempts to improve the upper bound. The failure to find ways to route flow very clearly suggested the types of rankings that were obstacles (i.e., reversed block structures). In turn, the distances were found essentially using the primal linear program.

6 A Candidate Algorithm for Distortion 3

As a third application, we derive a purely combinatorial (i.e., not LP-based) voting mechanism, which we conjecture to have distortion 3. We show that this conjecture would follow from any of three different-looking combinatorial conjectures we will formulate.

The point of departure for the derivation of the mechanism is Corollary 6.1, which simplifies Lemma 3.1, reducing it to a purely combinatorial property of a certain graph. Corollary 6.1 was proved as Theorem 4.4 in [28], using a significantly more complex proof.

For any two candidates x, y , we consider the following bipartite graph $H_{x,y}$ on the node set (V, V) ; that is, there is one node on the “left” for each voter v , and one node on the “right” for each voter v' . (We will use “left” and “right” to distinguish the two vertex sets.) There is an edge (v, v') if and only if there exists a candidate $z \in X$ ($z = x$ or $z = y$ are explicitly allowed) such that $x \succeq_v z$ and $z \succeq_{v'} y$.

Corollary 6.1 *Let $x \neq y$ be two candidates. If $H_{x,y}$ has a perfect matching, then $C(x) \leq 3C(y)$.*

Proof. Assume that there is a perfect matching in $H_{x,y}$; for each voter v , let μ_v be the voter v is matched with. By definition of $H_{x,y}$, there is a candidate \tilde{x}_v such that $x \succeq_v \tilde{x}_v$ and $\tilde{x}_v \succeq_{\mu_v} y$.

We now define the flow f from each source node (v, x) . We route one unit of flow along the path $(v, x) \rightarrow (v, \tilde{x}_v) \rightarrow (\mu_v, \tilde{x}_v) \rightarrow (\mu_v, y)$. Notice that by definition of \tilde{x}_v , the first and third edge always exist. Also, if $\tilde{x}_v = x$, then the first two nodes are the same, and we omit the first edge. Similarly, if $\mu_v = v$, we omit the second edge, and if $\tilde{x}_v = y$, we omit the third edge.

This construction defines a valid flow, routing one unit of flow from each (v, x) to some (v', y) (for some v'). So it only remains to bound $\gamma_v^{(f)} \leq 3$ for all v .

Because μ is a matching, there is exactly one unit of flow arriving at each node (v, y) . For a given voter v , let v' be the unique voter with $\mu_{v'} = v$. Then, the only two edges of the form $(v, z) \rightarrow (v', z)$

or $(v', z) \rightarrow (v, z)$ that can be used by f are $(v, \tilde{x}_v) \rightarrow (\mu_v, \tilde{x}_v)$ and $(v', \tilde{x}_{v'}) \rightarrow (v, \tilde{x}_{v'})$. Hence, the second part of the cost term $\gamma_v^{(f)}$ is at most 2, meaning that $\gamma_v^{(f)} \leq 3$. The claim now follows by applying Lemma 3.1. \blacksquare

Corollary 6.1 immediately suggests a natural mechanism with distortion at most three, which was also given as MATCHINGUNCOVERED in [28]:

ALL BIPARTITE MATCHINGS:

Find a candidate w such that for all other candidates x , the bipartite graph $H_{w,x}$ has a perfect matching.

The mechanism ALL BIPARTITE MATCHINGS sidesteps having to solve the $\Theta(n^2)$ LPs (1), instead solving $\Theta(n^2)$ bipartite matching problems. The question then is whether such a candidate w actually exists. We present three different conjectures, each of which would imply the existence of w , and thus, by Corollary 6.1, that ALL BIPARTITE MATCHINGS has distortion 3.

6.1 The Candidate Comparison Graph \mathcal{G}

A key analysis tool in this section is a directed graph \mathcal{G} on the set of all candidates X , which we call the *Candidate Comparison Graph*. \mathcal{G} contains the directed edge (y, x) if and only if the graph $H_{x,y}$ does *not* have a perfect matching. In a sense, y is a witness that x would be a dangerous choice as winner, since Corollary 6.1 would not apply to bound the cost of x when y is the optimal candidate.¹⁰ Conversely, any candidate w without incoming edges in \mathcal{G} is a safe choice as a winner, because Corollary 6.1 implies a bound of 3 on its cost ratio to x^* .

The following straightforward lemma captures that if we remove some candidates, and leave each voter's ranking of the remaining candidates unchanged, then the edges of the resulting graph \mathcal{G}' are a superset of the edges of the subgraph of \mathcal{G} induced by the remaining candidates. We write $\mathcal{P}[X']$ for the vector of rankings \succ'_v , where each \succ'_v is the ranking \succ_v , restricted to candidates in X' . In other words, $x \succ'_v y$ if and only if $x \succ_v y$, for all $x, y \in X'$.

Lemma 6.2 *Let (X, \mathcal{P}) be a social choice instance, and $X' \subseteq X$. Let $\mathcal{G} = \mathcal{G}(X, \mathcal{P})$ and $\mathcal{G}' = \mathcal{G}(X', \mathcal{P}[X'])$. Then, $\mathcal{G}' \supseteq \mathcal{G}[X']$; that is, \mathcal{G}' contains all edges of the induced subgraph $\mathcal{G}[X']$.*

Proof. Consider two candidates $x, y \in X'$ and their corresponding bipartite graph $H'_{x,y}$ on voters under the restricted instance. By definition, $H'_{x,y}$ contains the edge (v, v') iff there exists a candidate $z \in X'$ such that $x \succeq_v z$ and $z \succeq_{v'} y$. Thus, if $H'_{x,y}$ contains the edge (v, v') , then so does the bipartite graph $H_{x,y}$ for the larger/original instance (X, \mathcal{P}) . In other words, the edges of $H'_{x,y}$ are a subset of the edges of $H_{x,y}$. Therefore, whenever $H'_{x,y}$ contains a perfect matching, so does $H_{x,y}$. Because edges in \mathcal{G} (and \mathcal{G}') correspond to pairs that do *not* have bipartite matchings, the graph \mathcal{G}' is a supergraph of the induced subgraph $\mathcal{G}[X']$. \blacksquare

6.2 Acyclicity of \mathcal{G}

One sufficient condition for the existence of a source node in \mathcal{G} (i.e., a node without incoming edges) is for \mathcal{G} to be acyclic. This gives rise to our first conjecture, which was also given as Conjecture 4.8 in [28] (though it is expressed slightly differently there):

¹⁰Of course, Corollary 6.1 is only a *sufficient* condition, not a necessary one. So even when Corollary 6.1 cannot be applied, it is conceivable that w would achieve a distortion of 3. However, it is not clear which tool we could use to bound the distortion, which is why we focus only on the implications of Corollary 6.1 here.

Conjecture 1 For every instance (X, \mathcal{P}) the graph $\mathcal{G} = \mathcal{G}(X, \mathcal{P})$ is non-Hamiltonian.¹¹

Since our goal here is merely to show the existence of a source node in \mathcal{G} , it appears like overkill to aim for the “stronger” conjecture of being non-Hamiltonian/acyclic. Despite appearances, Conjecture 1 is not in fact stronger than the existence of a source node, as we show in the following proposition:

Proposition 6.3 ALL BIPARTITE MATCHINGS *succeeds* on all inputs *if and only if* Conjecture 1 is true.

Notice that the proposition does not say that whenever a *specific* instance violates Conjecture 1, the algorithm will fail on *that* instance. As will be evident in the proof, it only implies that the algorithm fails on *some* (potentially different) instance.

Proof. 1. For the first direction, if Conjecture 1 is true, then \mathcal{G} is non-Hamiltonian for all inputs. We claim that this implies that in fact, \mathcal{G} is *acyclic* for all inputs. Suppose that we had an instance (X, \mathcal{P}) for which \mathcal{G} contains a directed cycle, say, $C = (x_1, x_2, \dots, x_k, x_1)$ for some k . Let $X' = \{x_1, \dots, x_k\}$. Consider the instance $(X', \mathcal{P}[X'])$. By Lemma 6.2, the graph $\mathcal{G}(X', \mathcal{P}[X'])$ is a supergraph of the induced subgraph $\mathcal{G}[X']$, and must therefore contain a cycle including all vertices X' , i.e., $\mathcal{G}(X', \mathcal{P}[X'])$ is Hamiltonian.

We have thus shown that for all instances, \mathcal{G} is acyclic, meaning that for all inputs, \mathcal{G} has a source node, which is a safe output for ALL BIPARTITE MATCHINGS.

2. For the converse direction, assume that Conjecture 1 is false, and consider an instance (X, \mathcal{P}) for which $\mathcal{G}(X, \mathcal{P})$ is Hamiltonian. Because each node has at least one incoming edge, $\mathcal{G}(X, \mathcal{P})$ cannot have a source node. ■

6.3 Distributions of Permutations

We next derive a much simpler-looking — but actually equivalent — conjecture, which is phrased only in terms of distributions of permutations. The key lemma for deriving this equivalent conjecture is the following:

Lemma 6.4 \mathcal{G} contains the edge (y, x) if and only if there exists a set $Z_{x,y}$ of candidates with $x \in Z_{x,y}$ and $y \notin Z_{x,y}$ such that

$$|[y \succ Z_{x,y}]| + |[\overline{Z_{x,y}} \succ x]| > m. \quad (3)$$

Proof. The proof relies on the well-known Hall Theorem:

Theorem 6.5 (Hall’s Bipartite Matching Theorem) Let $G = (X \cup Y, E)$ be a bipartite graph with $|X| = |Y|$. For any vertex set S , let $\Gamma(S)$ denote the neighbors of S . G has a perfect matching if and only if there is no contracting vertex set, i.e., no set $X' \subseteq X$ with $|\Gamma(X')| < |X'|$.

Fix a pair x, y of candidates. By Hall’s Theorem, $(y, x) \in \mathcal{G}$ (i.e., $H_{x,y}$ has no perfect matching) if and only if there is a contracting set of voters $V_{x,y}$. By definition of \mathcal{G} , for every voter v on the left, the neighborhood $\Gamma(v)$ consists of all voters v' on the right such that there exists a candidate z with $x \succeq_v z$ and $z \succeq_{v'} y$.

¹¹Recall that a directed graph is *Hamiltonian* if it contains a directed cycle of all nodes.

1. For the first direction, we assume that \mathcal{G} contains the edge (y, x) . Let $V_{x,y}$ be a maximal contracting set. Let $Z_{x,y}$ be the set of all candidates z such that at least one voter $v \in V_{x,y}$ has $x \succeq_v z$. Then, $\Gamma(V_{x,y}) = \{v' \mid \text{there exists a } z \in Z_{x,y} \text{ with } z \succeq_{v'} y\}$. This implies two things:

- $\overline{\Gamma(V_{x,y})} = [y \succ Z_{x,y}]$ is the set of all voters who rank y strictly ahead of all of $Z_{x,y}$; this follows directly from the preceding characterization.
- $V_{x,y} = \overline{[Z_{x,y} \succ x]}$. The reason is that every voter $v \in V_{x,y}$ ranks only candidates in $Z_{x,y}$ (weakly) after x , so $V_{x,y} \subseteq \overline{[Z_{x,y} \succ x]}$. Since $\Gamma(V_{x,y}) = \Gamma(\overline{[Z_{x,y} \succ x]})$, the set $\overline{[Z_{x,y} \succ x]}$ is also a candidate for a contracting set, and must equal $V_{x,y}$ by maximality of $V_{x,y}$.

By definition of $Z_{x,y}$, we always have $x \in Z_{x,y}$. Also, we always have $y \notin Z_{x,y}$, because any voter v (on the left) with $x \succ_v y$ has edges to all voters v' on the right, and can therefore not be in a contracting set $V_{x,y}$.

Next, we consider the cardinalities of the sets involved. Because

$$m - |[y \succ Z_{x,y}]| = m - |\overline{\Gamma(V_{x,y})}| = |\Gamma(V_{x,y})| \stackrel{\text{Hall}}{<} |V_{x,y}| = |\overline{[Z_{x,y} \succ x]}|,$$

we have shown that there exists a set $Z_{x,y}$ with $x \in Z_{x,y}$ and $y \notin Z_{x,y}$ such that $|[y \succ Z_{x,y}]| + |\overline{[Z_{x,y} \succ x]}| > m$.

2. For the converse direction, assume that there exists a set $Z_{x,y}$ of candidates with $x \in Z_{x,y}$ and $y \notin Z_{x,y}$ such that $|[y \succ Z_{x,y}]| + |\overline{[Z_{x,y} \succ x]}| > m$.

Let $V_{x,y} = \overline{[Z_{x,y} \succ x]}$ be the set of all voters who rank x behind all of $\overline{Z_{x,y}}$. We will show that $V_{x,y}$ is contracting. Thereto, the important step is to characterize the neighborhood $\Gamma(V_{x,y})$. By definition, it consists of all voters v' such that there exists a voter $v \in V_{x,y}$ and a candidate z with $x \succeq_v z$ and $z \succeq_{v'} y$. Because each voter $v \in V_{x,y}$ ranks x behind all of $\overline{Z_{x,y}}$, the only potential candidates for z are candidates in $Z_{x,y}$. In particular, no voter $v' \in [y \succ Z_{x,y}]$ can be in $\Gamma(V_{x,y})$, implying that $[y \succ Z_{x,y}] \subseteq \overline{\Gamma(V_{x,y})}$.

This implies that $|[y \succ Z_{x,y}]| \leq |\overline{\Gamma(V_{x,y})}|$, so $|\Gamma(V_{x,y})| \leq m - |[y \succ Z_{x,y}]|$. And by definition of $V_{x,y}$, we also have $|V_{x,y}| = |\overline{[Z_{x,y} \succ x]}|$. Together with the assumption that $|[y \succ Z_{x,y}]| + |\overline{[Z_{x,y} \succ x]}| > m$, we get that $m < |V_{x,y}| + (m - |\Gamma(V_{x,y})|)$, implying that $V_{x,y}$ is contracting. By Hall's Theorem, $H_{x,y}$ has no perfect matching, so \mathcal{G} contains the edge (y, x) . ■

Remark 6.6 As an easy corollary of Lemma 6.4, notice that the constraint (3) implies that strictly more than half of the voters prefer y over x . This is because $x \in Z_{x,y}$ and $y \notin Z_{x,y}$ implies that all voters in $[y \succ Z_{x,y}]$ and all voters in $\overline{[Z_{x,y} \succ x]}$ rank y ahead of x . Because the combined cardinalities of the two sets add up to more than m , by the Pigeon Hole Principle, at least one of the two sets $[y \succ Z_{x,y}]$, $\overline{[Z_{x,y} \succ x]}$ must contain more than half of all voters. This proves that \mathcal{G} is a subgraph of the weak comparison graph.

Based on Lemma 6.4, we formulate the following conjecture, and prove it equivalent to Conjecture 1.

Conjecture 2 Let $Z_1, Z_2, \dots, Z_n \subseteq \{1, \dots, n\}$ be arbitrary sets with $i \in Z_i$. Define the following two indicator functions over elements i and total orders \succ :

$$\alpha(i, \succ) = \begin{cases} 1 & \text{if } i+1 \succ Z_i \\ 0 & \text{otherwise;} \end{cases} \quad \beta(i, \succ) = \begin{cases} 1 & \text{if } \overline{Z_i} \succ i \\ 0 & \text{otherwise;} \end{cases} \quad (4)$$

here and for the rest of this proof, all additions/subtractions are modulo n ; that is, $n + 1 := 1$ and $1 - 1 := n$.

Let \mathcal{D} be any distribution over total orders \succ of $\{1, \dots, n\}$. Then, there exists an i such that

$$\mathbb{E}_{\succ \sim \mathcal{D}} [\alpha(i, \succ) + \beta(i, \succ)] \leq 1. \quad (5)$$

Proposition 6.7 *Conjecture 2 is true if and only if Conjecture 1 is true.*

Proof. We first show that Conjecture 2 implies Conjecture 1, by proving the contrapositive. Assume that \mathcal{G} is Hamiltonian, containing a directed cycle $x_n \rightarrow x_{n-1} \rightarrow \dots \rightarrow x_1 \rightarrow x_n$ comprising all n candidates. By Lemma 6.4, there exist sets $Z_i = Z_{x_i, x_{i+1}}$ (with $Z_n = Z_{x_n, x_1}$) with $x_i \in Z_i, x_{i+1} \notin Z_i$, such that for all $i = 1, \dots, n$, we have

$$|[x_{i+1} \succ Z_i]| + |[\overline{Z_i} \succ x_i]| > m. \quad (6)$$

We define a distribution \mathcal{D} over rankings by drawing a uniformly random voter from $\{1, \dots, m\}$, and choosing this voter's ranking. Because the distribution is uniform, we obtain that

$$\mathbb{E}_{\succ \sim \mathcal{D}} [\alpha(i, \succ)] = \frac{1}{m} \cdot |[x_{i+1} \succ Z_i]|, \quad \mathbb{E}_{\succ \sim \mathcal{D}} [\beta(i, \succ)] = \frac{1}{m} \cdot |[\overline{Z_i} \succ x_i]|. \quad (7)$$

The inequality (6) then implies that $\mathbb{E}_{\succ \sim \mathcal{D}} [\alpha(i, \succ) + \beta(i, \succ)] > 1$ for all i , showing that Conjecture 2 is violated.

For the converse direction, assume that \mathcal{D} is a distribution over total orders of $\{1, \dots, n\}$ such that

$$\mathbb{E}_{\succ \sim \mathcal{D}} [\alpha(i, \succ) + \beta(i, \succ)] > 1, \text{ for all } i. \quad (8)$$

Define $\delta := \min_i \mathbb{E}_{\succ \sim \mathcal{D}} [\alpha(i, \succ) + \beta(i, \succ)]$; because the number of candidates is finite, δ is well-defined, and $\delta > 1$. Therefore, with sufficiently small changes to \mathcal{D} , we can ensure that the probability for each total order \succ is a rational number, while preserving all (strict) inequalities (8). Once the probabilities are all rational, we can write them with a common denominator, meaning that we can define \mathcal{D} as a uniform distribution over a finite multi-set of rankings; in turn, we can consider these rankings as voters.

Because the distribution is uniform over voters, we can apply the characterization (7) to conclude that $|[x_{i+1} \succ Z_i]| + |[\overline{Z_i} \succ x_i]| > m$ for all candidates i . By Lemma 6.4, applied to each pair (x_i, x_{i+1}) , the graph \mathcal{G} contains each edge (x_{i+1}, x_i) , so \mathcal{G} is Hamiltonian. \blacksquare

6.4 A Graph-Theoretic Reformulation

Our attempts to prove Conjecture 2 (so far unsuccessful) have been based on proofs by contradiction. The assumed constraints (8) prescribe several constraints on rankings that must hold simultaneously; using transitivity, this leads to a contradiction by forcing preferences to contain cycles. The essence of this approach is captured by another conjecture. To formulate it, we define the following class of directed graphs, which we term *Constraint-Choice Graphs*.

Definition 6.8 (Constraint-Choice Graph) *Let $Y_n = \{y_1, \dots, y_n\}, A_n = \{a_1, \dots, a_n\}, B_n = \{b_1, \dots, b_n\}$ be three disjoint sets of nodes. A constraint-choice graph for a given n contains $3n$ nodes $U_n = Y_n \cup A_n \cup B_n$, and the following edges:*

- For each i , it contains the directed edges¹² $(y_i, a_{i-1}), (y_i, b_{i-1}), (a_i, y_i), (b_i, y_i)$.

¹²As before, we define $1 - 1 := n$ and $n + 1 := 1$.

- For each i, j with $j \neq i, j \neq i-1$, it contains exactly one of the two directed edges $(a_j, y_i), (y_i, b_j)$.

An example of a constraint choice graph is shown in Figure 3. The edges listed first in Definition 6.8 are shown in solid black, while the edges listed second are shown in dashed red lines.

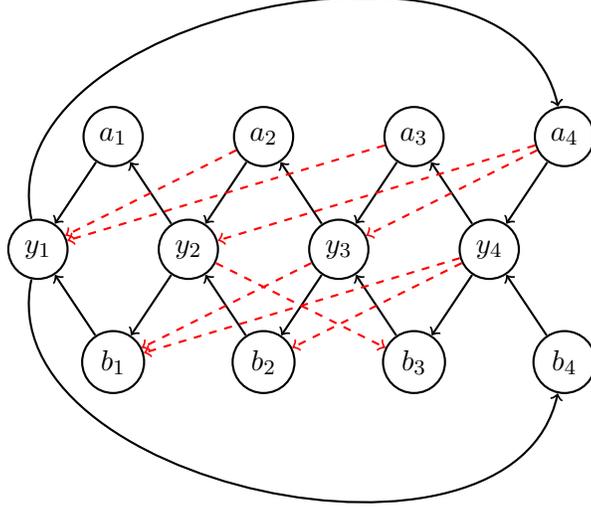


Figure 3: An illustration of a constraint-choice graph for $n = 4$ candidates. The graph depicted here corresponds to the sets $Z_1 = \{1\}, Z_2 = \{1, 2\}, Z_3 = \{1, 3\}, Z_4 = \{2, 3, 4\}$ in the construction of the proof of Proposition 6.10.

Conjecture 3 For every n and every constraint choice graph G_n of $3n$ nodes, there exists a non-empty index set $S \subseteq \{1, \dots, n\}$ with the following property: For every vertex set $T \subseteq \{a_i, b_i \mid i \in S\}$ of size $|T| > |S|$, the induced subgraph $G_n[Y_n \cup T]$ contains a directed cycle.

Remark 6.9 Notice that the conjecture indeed talks about the subgraph induced by *all* nodes y_i (not just those with indices in T), in addition to at least $|S| + 1$ nodes from among the a_i, b_i with $i \in S$.

Proposition 6.10 If Conjecture 3 is true, then Conjecture 2 is true.

Proof. We prove the contrapositive, and assume that Conjecture 2 is false; that is, Inequality (8) holds for all i .

Given the (assumed) sets $Z_1, Z_2, \dots, Z_n \subseteq \{1, \dots, n\}$, we define the following constraint-choice graph G_n . It contains the $3n$ nodes $Y_n \cup A_n \cup B_n$, and all the edges that are prescribed by Definition 6.8; in addition, if $i \in Z_j$, then G_n contains the edge (a_j, y_i) ; otherwise, it contains the edge (y_i, b_j) . This completes the definition of G_n . An example for specific sets Z_j is shown in Figure 3. To gain intuition for the following proof, the reader is encouraged to think of y_i as corresponding to candidate x_i , of a_i as corresponding to Z_i , and of b_i as corresponding to \overline{Z}_i .

To prove that G_n violates Conjecture 3, consider an arbitrary non-empty set of indices $S \subseteq \{1, \dots, n\}$. By linearity of expectations, (8) implies that $\mathbb{E}_{\succ \sim \mathcal{D}} [\sum_{i \in S} (\alpha(i, \succ) + \beta(i, \succ))] > |S|$. Because the maximum is at least the average, this implies that there exists some ranking \succ_S with $\sum_{i \in S} (\alpha(i, \succ_S) + \beta(i, \succ_S)) > |S|$, and because the quantity on the left-hand side is integral, we can strengthen this inequality to

$$\sum_{i \in S} (\alpha(i, \succ_S) + \beta(i, \succ_S)) \geq |S| + 1. \quad (9)$$

Define the node set $T_S := \{a_i \mid \alpha(i, \succ_S) = 1\} \cup \{b_i \mid \beta(i, \succ_S) = 1\}$. By Inequality (9), T_S contains at least $|S| + 1$ nodes from $\{a_i, b_i \mid i \in S\}$. We will show that the induced subgraph $G_n[T_S \cup Y_n]$ is acyclic; since we show this for every S , it proves that G_n violates Conjecture 3.

To show that $G_n[T_S \cup Y_n]$ is acyclic, we use a proof by contradiction, and assume that $G_n[T_S \cup Y_n]$ contains a cycle C . Because edges only go between nodes y_i and either a_j or b_j , this cycle must alternate nodes y_i with nodes a_j or b_j . Let $|C| = 2k$, and let i_1, i_2, \dots, i_k be such that the order of nodes y_i in C is $y_{i_1}, y_{i_2}, \dots, y_{i_k}, y_{i_1}$. Fix some $\ell \in \{1, \dots, k\}$. Between y_{i_ℓ} and $y_{i_{\ell+1}}$ (here, $k+1 := 1$), the cycle must visit either some node $a_j \in T_S$ or some node $b_j \in T_S$. We distinguish two cases:

- If the intermediate node is a_j , observe first that by Definition 6.8, the only incoming edge to a_j is (y_{j+1}, a_j) , so $i_\ell = j + 1$. The specific constraint-choice graph G_n defined in this proof includes outgoing edges from a_j to exactly the y_i with $i \in Z_j$; notice that this includes the case of the edge (a_j, y_j) , because $j \in Z_j$. In particular, it applies to the edge $(a_j, y_{i_{\ell+1}})$, implying that $i_{\ell+1} \in Z_j$.

Because $a_j \in T_S$, we have that $\alpha(i, \succ_S) = 1$, meaning that under \succ_S , candidate $i_\ell = j + 1$ precedes all candidates in Z_j ; this includes, in particular, the candidate $i_{\ell+1}$. In summary, we have inferred that $i_\ell \succ_S i_{\ell+1}$.

- If the intermediate node is b_j , observe that by Definition 6.8, the only outgoing edge from b_j is (b_j, y_j) , implying that $i_{\ell+1} = j$. For the particular graph G_n defined in this proof, the incoming edge (y_{i_ℓ}, b_j) exists exactly when $i_\ell \notin Z_j$.

Because $b_j \in T_S$, we have that $\beta(j, \succ_S) = 1$, so under \succ_S , candidate $j = i_{\ell+1}$ is ranked after all candidates in $\overline{Z_j}$. By the argument of the preceding paragraph, the set Z_j includes, in particular, the candidate i_ℓ . In summary, we have again inferred that $i_\ell \succ_S i_{\ell+1}$.

Thus, we have derived that $i_\ell \succ_S i_{\ell+1}$ for each $\ell = 1, \dots, k$. By transitivity, this results in a cycle in \succ_S , a contradiction to it being a ranking. \blacksquare

Remark 6.11 Conjecture 3 is sufficiently clean and combinatorial that it can be verified by hand for $n \leq 4$. An exhaustive computer search for $n \leq 7$ has verified the conjecture for all such n , recovering and extending computer-assisted results in Theorem 4.11 in [28] (although we did not include some ranges for $n > 7$ which [28] handle with a restricted number of voters). Unfortunately, because it basically involves a search over all possibilities of n subsets Z_i of n elements (represented as graph edge choices), the running time scales roughly as 2^{n^2} ; from $n = 6$ to $n = 7$, the running time increases from less than a minute to roughly a day. Hence, even $n = 8$ is likely out of reach. But the computer search is encouraging in terms of trying to prove the conjecture (rather than disproving it).

7 Conclusions

Our work raises a very obvious question: prove (or possibly disprove) the conjectures stated in Section 6. Based on exhaustive computer search, it seems more likely that the conjectures are true, and the ALL BIPARTITE MATCHINGS mechanism in fact is always able to find a candidate with distortion at most 3.

Going beyond these conjectures, we believe that the duality-based framework may be useful for bounding the performance of other voting mechanisms, in particular, those that may miss information on parts of voters' ranking. For instance, such a situation can occur in the setting of [24], where voters can only name the candidates in a subset of positions on their ballot, rather than giving a complete ranking. The analysis of a mechanism proposed in [24] becomes much simpler (and tighter) using the techniques developed here.

While we have only studied deterministic mechanisms here, the framework can also be extended to randomized mechanisms. When the mechanism selects a candidate x with probability q_x , an upper bound can be obtained by bounding a flow that inserts q_x units of flow at each of the nodes (v, x) , which again have to be routed to x^* .

It is conceivable that duality-based approaches similar to the one we developed could be helpful for the analysis of mechanisms for other problems in the cardinal/ordinal framework: a worst-case metric for a given input can often be characterized in terms of a linear program, and the dual may in general lead to a framework for proving upper bounds on the performance of a chosen mechanism.

Acknowledgements

The author would like to thank Elliot Anshelevich, Sid Banerjee, Shaddin Dughmi, Bobby Kleinberg and Kai Wang for useful conversations and pointers, and anonymous reviewers for useful feedback.

References

- [1] Elliot Anshelevich. Ordinal approximation in matching and social choice. *ACM SIGecom Exchanges*, 15(1):60–64, July 2016.
- [2] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [3] Elliot Anshelevich, Onkar Bhardwaj, and John Postl. Approximating optimal social choice under metric preferences. In *Proc. 29th AAAI Conf. on Artificial Intelligence*, pages 777–783, 2015.
- [4] Elliot Anshelevich and John Postl. Randomized social choice functions under metric preferences. In *Proc. 25th Intl. Joint Conf. on Artificial Intelligence*, pages 46–59, 2016.
- [5] Elliot Anshelevich and Shreyas Sekar. Blind, greedy, and random: Algorithms for matching and clustering using only ordinal information. In *Proc. 30th AAAI Conf. on Artificial Intelligence*, pages 390–396, 2016.
- [6] Salvador Barberà. An introduction to strategy-proof social choice functions. *Social Choice and Welfare*, 18:619–653, 2001.
- [7] Salvador Barberà, Faruk Gul, and Ennio Stacchetti. Generalized median voter schemes and committees. *Journal of Economic Theory*, 61:262–289, 1993.

- [8] Duncan Black. On the rationale of group decision making. *J. Political Economy*, 56:23–34, 1948.
- [9] Duncan Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.
- [10] Craig Boutilier, Ioannis Caragiannis, Simi Haber, Tyler Lu, Ariel D. Procaccia, and Or Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [11] Craig Boutilier and Jeffrey S. Rosenschein. Incomplete information and communication in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 10, pages 223–257. Cambridge University Press, 2016.
- [12] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [13] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A duality based unified approach to bayesian mechanism design. In *Proc. 48th ACM Symp. on Theory of Computing*, pages 926–939, 2016.
- [14] Ioannis Caragiannis and Ariel D. Procaccia. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence*, 175(9):1655–1671, 2011.
- [15] Yu Cheng, Shaddin Dughmi, and David Kempe. Of the people: Voting is more effective with representative candidates. In *Proc. 18th ACM Conf. on Economics and Computation*, pages 305–322, 2017.
- [16] Yu Cheng, Shaddin Dughmi, and David Kempe. On the distortion of voting with multiple representative candidates. In *Proc. 32nd AAAI Conf. on Artificial Intelligence*, pages 973–980, 2018.
- [17] Jean-Charles de Borda. Mémoire sur les élections au scrutin. *Histoire de l’Académie Royale des Sciences, Paris*, pages 657–665, 1784.
- [18] M. J. A. Nicolas de Condorcet. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale, Paris, 1785.
- [19] Anthony Downs. An economic theory of political action in a democracy. *The Journal of Political Economy*, 65(2):135–150, 1957.
- [20] Michal Feldman, Amos Fiat, and Iddan Golomb. On voting and facility location. In *Proc. 17th ACM Conf. on Economics and Computation*, pages 269–286, 2016.
- [21] Alan F. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–601, 1973.
- [22] Ashish Goel, Anilesh Kollagunta Krishnaswamy, and Kamesh Munagala. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proc. 18th ACM Conf. on Economics and Computation*, pages 287–304, 2017.
- [23] Stephen Gross, Elliot Anshelevich, and Lirong Xia. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proc. 31st AAAI Conf. on Artificial Intelligence*, pages 544–550, 2017.

- [24] David Kempe. Communication, distortion, and randomness in metric voting. In *Proc. 34th AAAI Conf. on Artificial Intelligence*, 2020.
- [25] Samuel Merrill and Bernard Grofman. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press, 1999.
- [26] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35:437–455, 1980.
- [27] Hervé Moulin. Choosing from a tournament. *Social Choice and Welfare*, 3(4):271–291, 1986.
- [28] Kamesh Munagala and Kangning Wang. Improved metric distortion for deterministic social choice rules. In *Proc. 20th ACM Conf. on Economics and Computation*, pages 245–262, 2019.
- [29] Ariel D. Procaccia. Can approximation circumvent Gibbard-Satterthwaite? In *Proc. 24th AAAI Conf. on Artificial Intelligence*, pages 836–841, 2010.
- [30] Ariel D. Procaccia and Jeffrey S. Rosenschein. The distortion of cardinal preferences in voting. In *Proc. 10th Intl. Workshop on Cooperative Inform. Agents X*, pages 317–331, 2006.
- [31] Mark A. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [32] Markus Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare*, 36(2):267–303, 2011.
- [33] T. Nicolaus Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- [34] Vijay V. Vazirani. *Approximation Algorithms*. Springer, 2001.