On the Approximability of Dodgson and Young Elections

Ioannis Caragiannis* Jason A. Covey† Michal Feldman‡ Christopher M. Homan§
Christos Kaklamanis¶ Nikos Karanikolas‖ Ariel D. Procaccia**
Jeffrey S. Rosenschein††

Abstract

The voting rules proposed by Dodgson and Young are both designed to find the alternative closest to being a Condorcet winner, according to two different notions of proximity; the score of a given alternative is known to be hard to compute under either rule.

In this paper, we put forward two algorithms for approximating the Dodgson score: an LP-based randomized rounding algorithm and a deterministic greedy algorithm, both of which yield an $\mathcal{O}(\log m)$ approximation ratio, where m is the number of alternatives; we observe that this result is asymptotically optimal, and further prove that our greedy algorithm is optimal up to a factor of 2, unless problems in $\mathcal{NP}$ have quasi-polynomial time algorithms. Although the greedy algorithm is computationally superior, we argue that the randomized rounding algorithm has an advantage from a social choice point of view.

Further, we demonstrate that computing any reasonable approximation of the ranking produced by Dodgson's rule is $\mathcal{NP}$-hard. This result provides a complexity-theoretic explanation of sharp discrepancies that have been observed in the Social Choice Theory literature when comparing Dodgson elections with simpler voting rules.

Finally, we show that the problem of calculating the Young score is $\mathcal{NP}$-hard to approximate by any factor. This leads to an inapproximability result for the Young ranking.

Keywords: Computational social choice, Approximation algorithms

*Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece, email: `caragian@ceid.upatras.gr`

†Department of Computer Science, Rochester Institute of Technology, 102 Lomb Memorial Drive, Rochester, NY 14623-5603, email: `jac8687@rit.edu`

‡School of Business Administration and Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: `mfeldman@huji.ac.il`

§Department of Computer Science, Rochester Institute of Technology, 102 Lomb Memorial Drive Rochester, NY 14623-5603, email: `cmh@cs.rit.edu`

¶Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece, email: `kakl@ceid.upatras.gr`

‖Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics University of Patras, 26500 Rio, Greece, email: `nkaranik@ceid.upatras.gr`

**School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: `arielpro@cs.huji.ac.il`

††School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: `jeff@cs.huji.ac.il`

1 Introduction

The discipline of voting theory deals with the following setting: a group of n agents each ranks a set of m alternatives; one alternative is to be elected. The big question is: which alternative best reflects the social good? The French philosopher and mathematician Marie Jean Antoine Nicolas de Caritat, marquis de Condorcet, suggested the following intuitive criterion: the winner should be an alternative that beats every other alternative in a *pairwise election*, i.e., an alternative that a majority of the agents prefers over any other alternative. Sadly, it is fairly easy to see that the preferences of the majority may be cyclic, hence a *Condorcet winner* does not necessarily exist. This unfortunate phenomenon is known as the *Condorcet paradox* (see Black [5]).

In order to circumvent this result, several researchers have proposed to choose an alternative that is "as close as possible" to a Condorcet winner. Different notions of proximity can be considered, and yield different voting rules. One such notion was suggested in 1876 by Charles Dodgson, better known by his pen name Lewis Caroll, author of "Alice's Adventures in Wonderland". The *Dodgson score* [5] of an alternative, with respect to a given set of agents' preferences, is the minimum number of exchanges between adjacent alternatives in the agents' rankings one has to introduce in order to make the given alternative a Condorcet winner. A *Dodgson winner* is any alternative with a minimum Dodgson score.

Young [36] raised a second option: measuring the distance by agents. Specifically, the *Young score* of an alternative is the size of the largest subset of agents such that, if only these ballots are taken into account, the given alternative becomes a Condorcet winner. A *Young winner* is any alternative with the maximum Young score. Alternatively, one can perceive a Young winner as the alternative that becomes a Condorcet winner by removing the least number of agents.

Though these two voting rules sound appealing and straightforward, they are notoriously complicated to resolve. As early as 1989, Bartholdi, Tovey and Trick [2] showed that computing the Dodgson score is $\mathcal{NP}$-complete, and that pinpointing a Dodgson winner is $\mathcal{NP}$-hard. This important paper was one of the first to introduce complexity-theoretic considerations to social choice theory. Hemaspaandra et al. [14] refined the abovementioned result by showing that the Dodgson winner problem is complete for Θ_2^p, the class of problems that can be solved by $\mathcal{O}(\log n)$ queries to an $\mathcal{NP}$ set. Subsequently, Rothe et al. [32] proved that the Young winner problem is also complete for Θ_2^p.

The abovementioned complexity results give rise to the agenda of *approximately* calculating an alternative's score, under the Dodgson and Young schemes. This is clearly an interesting computational problem, as an application area of algorithmic techniques.

However, from the point of view of social choice theory, it is not immediately apparent that an approximation of a voting rule is satisfactory, since an "incorrect" alternative—in our case, one that is not closest to a Condorcet winner—might be elected. Nevertheless, we argue that the use of such an approximation is strongly motivated. Indeed, at least in the case of the Dodgson and Young rules, the winner is an "approximation" in the first place, in instances where no Condorcet winner exists. Moreover, the approximation algorithm is equivalent to a new voting rule, which is guaranteed to elect an alternative that is not far from being a Condorcet winner. In other words, a perfectly sensible definition of a "socially good" winner, given the circumstances, is simply the alternative chosen by the approximation algorithm. Note that the approximation algorithm can be designed to satisfy the Condorcet criterion, i.e., always elect a Condorcet winner if one exists. This is always true for an approximation of the Dodgson score, as the Dodgson score of a Condorcet winner is zero. Moreover, approximation algorithms can be designed to satisfy other, less trivial, social choice desiderata, and hence may ultimately be considered socially sensible voting rules.

Related work. The agenda of approximating voting rules was recently pursued by Ailon et al. [1], Coppersmith et al. [8], and Kenyon-Mathieu and Schudy [18]. These works deal, directly or indirectly, with the Kemeny rank aggregation rule, which chooses a ranking of the alternatives instead of a single winning alternative. The Kemeny rule picks the ranking that has the maximum number of agreements with the agents' individual rankings regarding the correct order of pairs of alternatives. Ailon et al. improve the trivial 2-approximation algorithm to an involved, randomized algorithm that gives an 11/7-approximation; Kenyon-Mathieu and Schudy further improve the approximation, and obtain a PTAS.

Two recent works have directly put forward algorithms for the Dodgson winner problem [15, 24]. Both papers independently build upon the same basic idea: if the number of agents is significantly larger than the number of alternatives, and one looks at a uniform distribution over the preferences of the agents, with high probability one obtains an instance on which it is trivial to compute the Dodgson score of a given alternative. This directly gives rise to an algorithm that can usually compute the Dodgson score (under the assumption on the number of agents and alternatives). However, this is not an approximation algorithm in the usual sense, since the algorithm *a priori* gives up on certain instances, whereas an approximation algorithm is judged by its worst-case guarantees. In addition, this algorithm would be useless if the number of alternatives is not small compared with the number of agents.[1]

Betzler et al. [4] have investigated the parameterized computational complexity of the Dodgson and Young rules. The authors have devised a fixed parameter algorithm for exact computation of the Dodgson score, where the fixed parameter is the "edit distance," i.e., the number of exchanges. Specifically, if k is an upper bound on the Dodgson score of a given alternative, n is the number of agents, and m the number of alternatives, the algorithm runs in time $\mathcal{O}(2^k \cdot nk + nm)$. Notice that in general it may hold that $k = \Omega(nm)$. In contrast, computing the Young score is $W[2]$-complete; this implies that there is no algorithm that computes the Young score exactly, and whose running time is polynomial in nm and only exponential in k, where the parameter k is the number of remaining votes. These results complement ours nicely, as we shall also demonstrate that computing the Dodgson score is in a sense easier than computing the Young score, albeit in the context of approximation.

Putting computational complexity aside, several works by social choice theorists have considered comparing the ranking produced by Dodgson, i.e., the ordering of the alternatives by nondecreasing Dodgson score, with elections based on simpler voting rules. Such comparisons have always revealed sharp discrepancies. For example, the Dodgson winner can appear in any position in the Kemeny ranking [29] and in the ranking of any positional scoring rule [30] (e.g., Borda or Plurality), Dodgson rankings can be exactly the opposite of Borda [21] and Copeland rankings [19], while the winner of Kemeny of Slater elections can appear in any position of the Dodgson ranking [20].

More distantly related to our work is research that is concerned with exactly resolving hard-to-compute voting rules by heuristic methods. Typical examples include works regarding the Kemeny rule [7] and the Slater rule [6]. Another more remotely related field of research is concerned with finding approximate, efficient representations of voting rules, by eliciting as little

[1] This would normally not happen in political elections, but can certainly be the case in many other settings. For instance, consider a group of agents trying to reach an agreement on a joint plan, when multiple alternative plans are available. Specifically, think of a group of investors deciding which company to invest in.

information as possible; this line of research employs techniques from learning theory [26, 27].

Our results. In the context of approximating the Dodgson score, we devise an $\mathcal{O}(\log m)$ randomized approximation algorithm, where m is the number of alternatives. Our algorithm is based on solving the linear program proposed by Bartholdi et al. [2] and using randomized rounding. We then propose a second, deterministic and greedy, algorithm for the Dodgson score, with the same asymptotic approximation ratio. Although the latter algorithm is computationally superior in every way, we show that the former has the advantage of satisfying a flavor of monotonicity, which is a desirable property from a social choice point of view. We further observe that it follows from the work of McCabe-Dansted [23] that the Dodgson score cannot be approximated within sublogarithmic factors by polynomial-time algorithms unless $\mathcal{P} = \mathcal{NP}$. We prove a more explicit inapproximability result of $(1/2 - \epsilon)\ln m$, under the assumption that problems in $\mathcal{NP}$ do not have algorithms running in quasi-polynomial time; this implies that the approximation ratio achieved by our greedy algorithm is optimal up to a factor of 2.

Some of the results mentioned above [29, 30, 19, 20, 21] establish that there are sharp discrepancies between the Dodgson ranking and the rankings produced by other rank aggregation rules. Some of these rules (e.g., Borda and Copeland) are polynomial-time computable, so the corresponding results can be viewed as negative results regarding the approximability of the Dodgson ranking by polynomial-time algorithms. We show that the problem of distinguishing between whether a given alternative is the unique Dodgson winner or in the last $O(\sqrt{m})$ positions in any Dodgson ranking is $\mathcal{NP}$-hard. This theorem provides a complexity-theoretic explanation for some of the observed discrepancies, but in fact is much wider in scope as it applies to any efficiently computable rank aggregation rule.

The problem of calculating the Young score seems at first glance simple compared with the Dodgson score (we discuss in Section 4 why this seems so). Therefore, we found the following result quite surprising: it is $\mathcal{NP}$-hard to approximate the Young score within any factor. Specifically, we show that it is $\mathcal{NP}$-hard to distinguish between the case where the Young score of a given alternative is 0, and the case where the score is greater than 0. As a corollary we obtain an inapproximability result for the Young ranking.

Structure of the paper. In Section 2, we introduce some notations and definitions. In Section 3, we present our upper and lower bounds for approximating Dodgson elections. In Section 4, we prove that the Young score and ranking are inapproximable.

2 Preliminaries

Let $N = \{1, \ldots, n\}$ be a set of agents, and let A be the set of *alternatives*. We denote $|A| = m$, and denote the alternatives themselves by letters, such as $a \in A$. Indices referring to agents appear in superscript. Each agent $i \in N$ holds a binary relation R^i over A that satisfies irreflexivity, asymmetry, transitivity and totality. Informally, R^i is a ranking of the alternatives. Let $L = L(A)$ be the set of all rankings over A; we have that each $R^i \in L$. We denote $R^N = \langle R^1, \ldots, R^n \rangle \in L^N$, and refer to this vector as a *preference profile*. We may also use Q^i to denote the preferences of agent i, in cases where we want to distinguish between two different rankings R^i and Q^i. For sets of alternatives $B_1, B_2 \subseteq A$, we write $B_1 R^i B_2$ if for all $a \in B_1$ and $b \in B_2$, aR^ib.

Let $a, b \in A$. Denote $\{i \in N : \ aR^ib\}$ as $P(a, b)$. We say that *a beats b in a pairwise election* if $|P(a,b)| > n/2$, that is, a is preferred to b by the majority of agents. A *Condorcet winner* is an alternative that beats every other alternative in a pairwise election.

The *Dodgson score* of a given alternative a^*, with respect to a given preference profile R^N, is the least number of exchanges between adjacent alternatives in R^N needed to make a^* a Condorcet winner. For instance, let $N = \{1, 2, 3\}$, $A = \{a, b, c\}$, and let R^N be given by:

R^1	R^2	R^3
a	b	a
b	a	c
c	c	b

In this example, the Dodgson score of a is 0 (a is a Condorcet winner), the score of b is 1, and the score of c is 3. Bartholdi et al. [2] have shown that computing the Dodgson score is an $\mathcal{NP}$-complete problem.

The *Young score* of a^* with respect to R^N is the size of the largest subset of agents for whom a^* is a Condorcet winner. This is the definition given by Young himself [36], and used in subsequent works [32]. If for every nonempty subset of agents a^* is not a Condorcet winner, its Young score is 0. In the above example, the Young score of a is 3, the score of b is 1, and the score of c is 0.

Notice that, equivalently, a Young winner is an alternative such that one has to remove the least number of agents in order to make it a Condorcet winner. However, these two definitions are not equivalent in the context of approximation; we employ the former (original, prevalent) definition, but touch on the latter as well.

As the Young winner problem is known to be intractable [32], the Young score problem must also be hard; otherwise, we would be able to calculate the scores of all the alternatives efficiently, and identify the alternatives with maximum score.

3 Approximability of Dodgson

We begin by presenting our approximation algorithms for the Dodgson score. Let us first introduce some common notations.

Let $a^* \in A$ be a distinguished alternative, whose Dodgson score we wish to compute. Define the *deficit* of a^* with respect to $a \in A$, simply denoted $\text{def}(a)$ when the identity of a^* is clear, as the number of additional agents that must rank a^* above a in order for a^* to beat a in a pairwise election. For instance, if 4 agents prefer a to a^* and only one agent prefers a^* to a, then $\text{def}(a) = 2$. If a^* beats a in a pairwise election (namely a^* is preferred by the majority of agents) then $\text{def}(a) = 0$. We say that alternatives $a \in A$ with $\text{def}(a) > 0$ are *alive*. Alternatives that are not alive, i.e., $\text{def}(a) = 0$, are *dead*.

3.1 A Randomized Rounding algorithm

Bartholdi et al. [2] provide an integer linear programming (ILP) formulation for the Dodgson score. The number of constraints and variables in their program depends solely on the number of alternatives. Therefore, if the number of alternatives is constant, the program is solvable in polynomial time using the algorithm of Lenstra [22]. However, if the number of alternatives is not constant, the LP is of gargantuan size.[2]

Fortunately, it is easy to modify the abovementioned ILP to obtain a program of polynomial size. As before, let $a^* \in A$ be the alternative whose score we wish to compute. Let the variables of the program be $x^i_j \in \{0,1\}$ for all $i \in N$ and $j \in \{0, \ldots, m-1\}$; $x^i_j = 1$ if and only if a^* is moved upward, or *pushed*, by j positions in the ranking of agent i. Define constants $e^i_{ja} \in \{0,1\}$, for all $i \in N$, $j \in \{0, \ldots, m-1\}$, and $a \in A \setminus \{a^*\}$, which depend on the given preference profile; $e^i_{ja} = 1$ iff pushing a^* by j positions in the ranking of agent i makes a^* gain an *additional* vote against a (note that $e^i_{ja} = 0$ for all j if $a^* R^i a$). The ILP that computes the Dodgson score of a^* is given by:

$$\begin{aligned} \text{minimize} \quad & \sum_{i,j} j \cdot x^i_j \\ \text{subject to} \quad & \forall i \in N,\ \sum_j x^i_j = 1 \\ & \forall a \in A \setminus \{a^*\},\ \sum_{i,j} x^i_j e^i_{ja} \geq \text{def}(a) \\ & \forall i \in N, \forall j \in \{0, \ldots, m-1\},\ x^i_j \in \{0,1\} \end{aligned} \tag{1}$$

This ILP can be relaxed by requiring merely that $0 \leq x^i_j \leq 1$ for all i and j. The resulting linear program (LP) can be solved efficiently.

We are now ready to present our randomized rounding algorithm.

Randomized Rounding Algorithm
Input: An alternative a^* whose Dodgson score we wish to estimate, and a preference profile $R^N \in L^N$.
Output: An approximation of the Dodgson score of a^*.
The algorithm:

1. Solve the relaxed LP given by (1) to obtain a solution $\vec{x}$.
2. For $k = 1, \ldots, \alpha \cdot \log m$ (where $\alpha > 0$ is a constant to be chosen later)
   - For all $i \in N$, randomly and independently (from other agents and other iterations) choose a value X^i_k, such that $X^i_k = j$ with probability x^i_j.
3. For all $i \in N$, set $X^i_{max} = \max_k X^i_k$.
4. Let $\mathcal{X}'$ be the solution which moves a^* upward in the ranking of i by X^i_{max} positions; return $\text{cost}(\mathcal{X}') = \sum_{i \in N} X^i_{max}$.

We remark that if a^* is a Condorcet winner from the outset, clearly the algorithm will calculate a score of 0 (with probability 1). Therefore, if we defined a new (randomized) voting rule, which elects the alternative with minimal score according to the algorithm, this voting rule would satisfy the Condorcet criterion.

[2]Note that there is also an efficient solution if the number of agents n is constant; indeed, brute force search requires checking $\mathcal{O}(m^n)$ possibilities.

Theorem 3.1. *For any input a^* and R^N with m alternatives, the randomized rounding algorithm returns a $4\alpha \cdot \log m$-approximation of the Dodgson score of a^* with probability at least 1/2.*

The proof of the theorem is quite similar to the analysis of the randomized rounding algorithm for Set Cover [35, pp. 120-122], with one prominent additional argument, namely the application of Lemma A.1. The details of the proof are given in Appendix A.

Note that it is possible to verify in polynomial time whether the output of the algorithm is, at the same time, a valid solution (i.e., a^* is a Condorcet winner) and a $4\alpha \cdot \log m$-approximation (by comparing with OPT_f). Therefore, it is possible to repeat the algorithm from scratch to improve the probability of success. The expected number of repetitions required to achieve both the foregoing properties is at most 2.

3.2 A Deterministic Combinatorial Algorithm

In this section, we present a deterministic, combinatorial, greedy algorithm for approximating the Dodgson score of a given alternative. Consider, once again, a special alternative a^*, and recall that a live alternative is one with a positive deficit. In each step, the algorithm selects the most cost-effective push of alternative a^* in the preference of some agent. The *cost-effectiveness* of pushing a^* in the preference of an agent $i \in N$ is the ratio between the total number of positions a^* is moved upwards in the preference of i *compared with the original profile* R^N, and the number of *currently* live alternatives that a^* overtakes as a result of this push. For example, if for some agent the algorithm raises a^* by one position where the alternative over which a^* is raised is dead, and later by a second position that causes a^* to overtake a live alternative, then the cost-effectiveness of the push is two and not one, since a^* ends up being two positions higher than its original position and only overtakes one live alternative.

After selecting the most cost-effective push, the algorithm decreases $\mathrm{def}(a)$ by one for each live alternative a that a^* overtakes. Alternatives $a \in A$ with $\mathrm{def}(a) = 0$ become dead. The algorithm terminates when no live alternatives remain. The input and output of the algorithm are as before.

Greedy Algorithm:

1. Let A' be the set of live alternatives, namely those alternatives $a \in A$ with $\mathrm{def}(a) > 0$.
2. While $A' \neq \emptyset$:
   - Perform the most cost-effective push, namely push a^* in the preferences of agent $i \in N$ in a way that minimizes the ratio between the *total* number of positions moved upwards in the preferences of i and the number of currently live alternatives overtaken by a^*.
   - Recalculate A'.
3. Return the number of exchanges performed.

By the definition of the algorithm, it is clear that it produces a profile where a^* is a Condorcet winner. It is important to notice that, as is the case with the randomized rounding algorithm, if a^* is initially a Condorcet winner then the algorithm calculates a Dodgson score of zero, so as a voting rule the algorithm satisfies the Condorcet criterion.

Theorem 3.2. *For any input a^* and R^N with m alternatives, the greedy algorithm returns an H_{m-1}-approximation of the Dodgson score of a^*, where H_k is the k-th harmonic number.*

The proof of Theorem 3.2 (Appendix B) uses the dual fitting technique, and is based on the connection between our problem and the Constrained Set Multicover problem [28].

3.3 Interlude: On the Desirability of Approximation Algorithms as Voting Rules

In Section 1 we stated that an approximation algorithm for the Dodgson score should be considered as a new voting rule. This implies that our approximation algorithms should be compared according to two conceptually different, but not orthogonal, dimensions: their algorithmic properties and their social choice properties. Our greedy algorithm is clearly superior to the randomized rounding algorithm in terms of algorithmic properties: the former is combinatorial whereas the latter is LP-based; the former is deterministic whereas the latter is randomized. In the sequel we suggest, however, that the latter has some desirable properties from a social choice point of view. It is important to note at this point that randomized voting rules are considered legitimate in the social choice literature (see, e.g., [13, 9]), hence our randomized rounding algorithm may be considered a valid voting rule.

In most algorithmic mechanism design settings [25], such as combinatorial auctions or scheduling, one usually seeks approximation algorithms that are truthful, i.e., the agents cannot benefit by lying. However, the well-known Gibbard-Satterthwaite Theorem [12, 33] precludes voting rules that are both truthful and reasonable, in a sense. Therefore, other desiderata are looked for in voting rules.

We have been careful to emphasize that both the randomized rounding algorithm and the greedy algorithm satisfy the Condorcet property. Let us now consider the *monotonicity* property, one of the major desiderata on the basis of which voting rules are compared. Many different notions of monotonicity can be found in the literature; for our purposes, a (score-based) voting rule is *weakly monotonic* if and only if pushing an alternative in the preferences of the agents cannot worsen the score of the alternative, that is, increase it when a lower score is desirable (as in Dodgson), or decrease it when a higher score is desirable. All prominent score-based voting rules (positional scoring rules, Copeland, Maximin) are weakly monotonic; it is straightforward to see that the Dodgson and Young rules are weakly monotonic as well.

We first claim that our randomized rounding algorithm, or, more accurately, a slight variant thereof, is weakly monotonic. Indeed, consider the variant of the algorithm where $\mathcal{X}'$ is the solution that moves a^* upward in the ranking of i by $\sum_k X_k^i$ positions rather than $\max_k X_k^i$; the cost of this solution is

$$\text{cost}(\mathcal{X}') = \sum_k \sum_{i \in N} X_k^i \quad .$$

It is easy to verify (see (5)) that the exact same worst-case approximation bound holds for this variant as well (although in practice its approximation ratio would usually be significantly worse).

Now, consider a situation where a^* is moved upwards in the preferences of the agents. It is obvious that this decreases the value of OPT_f. In addition, for every k, we have $\mathbb{E}\left[\sum_i X_k^i\right] = \text{OPT}_f$. Therefore, by the linearity of expectation, the expected cost of the solution produced by the algorithm $\mathbb{E}\left[\sum_k \sum_{i \in N} X_k^i\right]$ decreases as well.

In contrast, let us now consider the greedy algorithm. We design a preference profile and a push of a^* that demonstrate that the algorithm is not weakly monotonic. Agents 1 through 6 vote according to the profile R^N given in Figure 1(a). The positions marked by "." are placeholders for the rest of the alternatives, in some arbitrary order. Let $A' = \{a_1, \ldots, a_4\}$, $A'' = \{b_1, \ldots, b_{17}\}$.

Notice that def$(a) = 1$ for all $a \in A'$, def$(b) = 0$ for all $b \in A''$. The optimal sequence of exchanges moves a^* all the way to the top of the preferences of agent 2, with a cost of seven. The greedy algorithm, given this preference profile, indeed chooses this sequence.

R^1	R^2	R^3	R^4	R^5	R^6
a_4	a_4				
a_3	a_3	a_4			
a_2	a_2	b_4	a_3		
a_1	a_1	b_5	b_9	a_2	
.	b_1	b_6	b_{10}	b_{13}	a_1
.	b_2	b_7	b_{11}	b_{14}	b_{16}
.	b_3	b_8	b_{12}	b_{15}	b_{17}
.	a^*	a^*	a^*	a^*	a^*
.	.	.	.	.	.
a^*	.	.	.	.	.

(a) Original Profile.

R^1	R^2	Q^3	Q^4	Q^5	Q^6
a_4	a_4				
a_3	a_3	a_4			
a_2	a_2	b_4	a_3		
a_1	a_1	b_5	b_9	a_2	
.	b_1	b_6	b_{10}	b_{13}	a_1
.	b_2	a^*	a^*	a^*	a^*
.	b_3	b_7	b_{11}	b_{14}	b_{16}
.	a^*	b_8	b_{12}	b_{15}	b_{17}
.	.	.	.	.	.
a^*	.	.	.	.	.

(b) Improvement of a^*.

Figure 1: The greedy algorithm is nonmonotonic: an example.

On the other hand, consider the profile $(R^1, R^2, Q^3, Q^4, Q^5, Q^6)$ given in Figure 1(b) (where the position of a^* was improved by two positions in the preferences of agents 3 through 6). First notice that the deficits have not changed compared to the profile R^N. The greedy algorithm would in fact push a^* to the top of the preferences of agents 6, 5, 4, and 3 (in this order), with a total cost of ten. Note that the optimal solution still has a cost of seven.

The following stronger notion of monotonicity is often considered in the literature: pushing a winning alternative in the preferences of the agents cannot harm it, that is, cannot make it lose the election. We say that a voting rule that satisfies this property is *strongly monotonic*.[3] Interestingly, Dodgson itself is not strongly monotonic [34], a fact that is considered by many to be a serious flaw. However, this does not preclude the existence of an approximation algorithm for the Dodgson score that is strongly monotonic as a voting rule. An intriguing open question is the existence of such algorithms with a good approximation ratio.

Additionally, there are other prominent social choice properties that are often considered, e.g., *homogeneity*: a voting rule is said to be homogeneous if duplicating the electorate does not change the outcome of the election. We leave the comparison of our two algorithms on the basis of additional social choice desiderata, as well as more general questions regarding the design of socially desirable approximation algorithms, for future work.

3.4 Lower Bounds

McCabe-Dansted [23] gives a polynomial-time reduction from the Minimum Dominating Set problem to the Dodgson score problem with the following property: given a graph G with k vertices, the reduction creates a preference profile with $n = \Theta(k)$ agents and $m = \Theta(k^4)$ alternatives, such that the size of the minimum dominating set of G is $\lfloor k^{-2}\mathrm{sc}_D(a^*)\rfloor$, where $\mathrm{sc}_D(a^*)$ is the Dodgson score of a distinguished alternative $a^* \in A$. Since the Minimum Dominating Set problem is known to be

[3] This is often simply referred to as *monotonic* in the literature

$\mathcal{NP}$-hard to approximate to within logarithmic factors [31], this implies that the Dodgson score problem is also hard to approximate to a factor of $\Omega(\log m)$. Due to the relation of Minimum Dominating Set to Minimum Set Cover, using an inapproximability result due to Feige [10], the explicit inapproximability bound can become $\left(\frac{1}{4} - \epsilon\right) \ln m$ under the assumption that problems in $\mathcal{NP}$ do not have quasi-polynomial-time algorithms.[4] This means that our algorithms are asymptotically optimal.

In the following, we present an alternative and more natural reduction directly from Minimum Set Cover that allows us to obtain a better explicit inapproximability bound. This bound implies that our greedy algorithm is optimal up to a factor of 2. The proof of the following theorem is given in Appendix C.

Theorem 3.3. *There exists $\beta > 0$ such that it is $\mathcal{NP}$-hard to approximate the Dodgson score of a given alternative in an election with m alternatives to within a factor of $\beta \ln m$. Furthermore, for any $\epsilon > 0$, there is no polynomial-time $\left(\frac{1}{2} - \epsilon\right) \ln m$-approximation for the Dodgson score of a given alternative unless problems in $\mathcal{NP}$ of input size k have algorithms running in time $k^{\mathcal{O}(\log \log k)}$.*

A related question is the approximability of the Dodgson ranking, that is, the ranking of alternatives given by ordering them by nondecreasing Dodgson score. To the best of our knowledge, no rank aggregation function, which maps preferences profiles to rankings of the alternatives, is known to provably produce rankings that are close to the Dodgson ranking [29, 30, 19, 20, 21] (see the survey of related work in Section 1).

Our next result establishes that efficient approximation algorithms are unlikely to exist unless $\mathcal{P} = \mathcal{NP}$, by proving that the problem of distinguishing between whether a given alternative is the unique Dodgson winner or in the last $O(\sqrt{m})$ positions is $\mathcal{NP}$-hard.

Theorem 3.4. *Given a preference profile with m alternatives and an alternative a^*, it is $\mathcal{NP}$-hard to decide whether a^* is a Dodgson winner or has rank at least $m - 6\sqrt{m}$ in any Dodgson ranking.*

Our proof, given in Appendix D, uses a reduction from Minimum Vertex Cover in 3-regular graphs and exploits a very weak statement concerning its inapproximability (marginally stronger than its $\mathcal{NP}$-hardness) that follows from the work of Berman and Karpinski [3]. The approach is similar to the proof of Theorem 3.3, albeit considerably more involved. This result provides a complexity-theoretic explanation for the sharp discrepancies observed in the Social Choice Theory literature when comparing Dodgson elections with simpler, efficiently computable, voting rules.

4 Approximability of Young

Recall that the Young score of a given alternative $a^* \in A$ is the size of the largest subset of agents for which a^* is a Condorcet winner.

It is straightforward to obtain a simple ILP for the Young score problem. As before, let $a^* \in A$ be the alternative whose Young score we wish to compute. Let the variables of the program be $x^i \in \{0, 1\}$ for all $i \in N$; $x^i = 1$ iff agent i is included in the subset of agents for a^*. Define constants $e^i_a \in \{-1, 1\}$ for all $i \in N$ and $a \in A \setminus \{a^*\}$, which depend on the given preference profile; $e^i_a = 1$ iff agent i ranks a^* higher than a. The ILP that computes the Young score of a^* is given by:

[4] Both inapproximability bounds have not been explicitly observed by McCabe-Dansted.

$$\begin{aligned} \text{maximize} \quad & \sum_{i \in N} x^i \\ \text{subject to} \quad & \forall a \in A \setminus \{a^*\},\ \sum_{i \in N} x^i e_a^i \geq 1 \\ & \forall i \in N,\ x^i \in \{0,1\} \end{aligned} \tag{2}$$

The ILP (2) for the Young score is seemingly simpler than the one for the Dodgson score, given as (1). This might seem to indicate that the problem can be easily approximated by similar techniques. Therefore, the following result is quite surprising.

Theorem 4.1. *It is $\mathcal{NP}$-hard to approximate the Young score by any factor.*

This result becomes more self-evident when we notice that the Young score has the rare property of being nonmonotonic as an optimization problem, in the following sense: given a subset of agents that make a^* a Condorcet winner, it is not necessarily the case that a smaller subset of the agents would satisfy the same property. This stands in contrast to many approximable optimization problems, in which a solution which is worse than a valid solution is also a valid solution. Consider the Set Cover problem, for instance: if one adds more subsets to a valid cover, one obtains a valid cover. The same goes for the Dodgson score problem: if a sequence of exchanges makes a^* a Condorcet winner, introducing more exchanges (that push a^* upwards) on top of the existing ones would not undo this fact.

In order to prove the inapproximability of the Young score, we define the following problem.

NonEmptySubset
Instance: An alternative a^*, and a preference profile $R^N \in L^N$.
Question: Is there a nonempty subset of agents $C \subseteq N$, $C \neq \emptyset$, for which a^* is a Condorcet winner?

To prove Theorem 4.1, it is sufficient to prove that NonEmptySubset is $\mathcal{NP}$-hard. Indeed, this implies that it is $\mathcal{NP}$-hard to distinguish whether the Young score of a given alternative is zero or greater than zero, which directly entails that the score cannot be approximated.

Lemma 4.2. NonEmptySubset *is $\mathcal{NP}$-complete.*

The proof of the lemma appears in Appendix E. A short discussion is in order. Theorem 4.1 states that the Young score cannot be efficiently approximated to any factor. The proof shows that, in fact, it is impossible to efficiently distinguish between a zero and a nonzero score. However, the proof actually shows more: it constructs a family of instances, where it is hard to distinguish between a score of zero and almost $2m/3$. Now, if one looks at an alternative formulation of the Young score problem where all the scores are scaled by an additive constant, it is no longer true that it is hard to approximate the score to *any* factor; however, the proof still shows that it is hard to approximate the Young score, even under this alternative formulation, to a factor of $\Omega(m)$.

As noted in Section 2, one can imagine another alternative formulation of the Young score. Indeed, one might ask: given a preference profile, what is the *smallest* number of agents that must be *removed* in order to make a^* a Condorcet winner? This minimization problem, where the score is the number of agents that are removed, is referred to as the *Dual Young score* by Betzler et al. [4].

Of course, a Young winner according to the primal formulation is always a winner according to the dual formulation, and vice versa. Notice that it is easy to obtain an ϵn-approximation under the dual formulation for any constant $\epsilon > 0$ by enumerating all subsets of agents of size at least $n - 1/\epsilon$ and checking whether a^* is the Condorcet winner in the preferences of these agents. However, we conjecture that the dual Young score is hard to approximate significantly better; we leave this issue for future work.

Finally, the strong inapproximability result for the Young score intuitively implies that the Young ranking cannot be approximated. The following corollary, whose proof (given in Appendix F) is a straightforward variation on the proof of Lemma 4.2, shows that this is indeed the case. It can be viewed as an analog of Theorem 3.4 for Young.

Corollary 4.3. *For any constant $\epsilon > 0$, given a preference profile with m alternatives and an alternative a^*, it is $\mathcal{NP}$-hard to decide whether a^* has rank $\mathcal{O}(m^\epsilon)$ or is ranked in place m (that is, ranked last) in any Young ranking.*

5 Acknowledgments

The authors wish to thank Nati Linial for pointing their attention to Jogdeo and Samuels [16], as well as Felix Fischer for helpful discussions. Ariel Procaccia is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

References

[1] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: ranking and clustering. In *Proceedings of the 37th annual ACM symposium on Theory of computing (STOC)*, pages 684–693, 2005.

[2] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6:157–165, 1989.

[3] P. Berman and M. Karpinski. On some tighter inapproximability results. In *Proceedings of the 26th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 1644 of *Lecture Notes in Computer Science*, pages 200–209. Springer, 1999.

[4] N. Betzler, J. Guo, and R. Niedermeier. Parameterized computational complexity of Dodgson and Young elections. In *Proceedings of the 11th Scandinavian Workshop on Algorithm Theory (SWAT)*, 2008.

[5] D. Black. *Theory of Committees and Elections*. Cambridge University Press, 1958.

[6] V. Conitzer. Computing Slater rankings using similarities among candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 613–619, 2006.

[7] V. Conitzer, A. Davenport, and J. Kalagnanam. Improved bounds for computing Kemeny rankings. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 620–626, 2006.

[8] D. Coppersmith, L. Fleischer, and A. Rudra. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proceedings of the 17th annual ACM-SIAM symposium on Discrete algorithm (SODA)*, pages 776–782, 2006.

[9] C. d'Aspremont and B. Peleg. Ordinal Bayesian incentive compatible representations of committees. *Social Choice and Welfare*, 5:261–279, 1988.

[10] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):643–652, 1998.

[11] M. R. Garey and D. S. Johnson. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[12] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

[13] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.

[14] E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.

[15] C. Homan and L. A. Hemaspaandra. Guarantees for the success frequency of an algorithm for finding Dodgson election winners. In *Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 528–539, 2006.

[16] K. Jogdeo and S. Samuels. Monotone convergence of binomial probabilities and a generalization of Ramanujan's equation. *Annals of Mathematical Statistics*, 39:1191–1195, 1968.

[17] M. Karpinski. Approximating bounded degree instances of NP-hard problems. In *Proceedings of the 13th International Symposium on Fundamentals of Computation Theory (FCT)*, volume 2138 of *Lecture Notes in Computer Science*, pages 24–34. Springer, 2001.

[18] C. Kenyon-Mathieu and W. Schudy. How to rank with few errors. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 95–103, 2007.

[19] C. Klamler. A comparison of the Dodgson method and the Copeland rule. *Economics Bulletin*, 4(8):1–7, 2003.

[20] C. Klamler. The Dodgson ranking and its relation to Kemeny's method and Slater's rule. *Social Choice and Welfare*, 23(1):91–102, 2004.

[21] C. Klamler. The Dodgson ranking and the Borda count: a binary comparison. *Mathematical Social Sciences*, 48(1):103–108, 2004.

[22] H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8:538–548, 1983.

[23] J. C. McCabe-Dansted. Approximability and computational feasibility of Dodgson's rule. Master's thesis, University of Auckland, 2006.

[24] J. C. McCabe-Dansted, G. Pritchard, and A. M. Slinko. Approximability of Dodgson's rule. In *Proceedings of the 1st International Workshop on Computational Social Choice (COMSOC)*, pages 331–344, 2006.

[25] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.

[26] A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. Learning voting trees. In *Proceedings of the 22nd National Conference on Artificial Intelligence (AAAI)*, pages 110–115, 2007.

[27] A. D. Procaccia, A. Zohar, and J. S. Rosenschein. Automated design of scoring rules by learning from examples. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 951–958, 2008.

[28] S. Rajagopalan and V. V. Vazirani. Primal-dual RNC approximation algorithms for set cover and covering integer programs. *SIAM Journal on Computing*, 28:526–541, 1999.

[29] T. C. Ratliff. A comparison of Dodgson's method and Kemeny's rule. *Social Choice and Welfare*, 18(1):79–89, 2001.

[30] T. C. Ratliff. A comparison of Dodgson's method and the Borda count. *Economic Theory*, 20(2):357–372, 2002.

[31] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th ACM Symposium on Theory of Computing (STOC)*, pages 475–484, 1997.

[32] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.

[33] M. 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.

[34] N. Tideman. *Collective Decisions and Voting: The Potential for Public Choice*. Ashgate, 2006.

[35] V. V. Vazirani. *Approximation Algorithms*. Springer, 2001.

[36] H. P. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16:335–353, 1977.

A Proof of Theorem 3.1

Fix some iteration k of the algorithm's for loop. Let $X^i = X^i_k$, $i \in N$, be independent discrete random variables such that $X^i = j$ with probability x^i_j. Consider the sequence of exchanges induced by the variables X^i, i.e., each agent $i \in N$ moves a^* upward by j places with probability x^i_j. As a result of the constraint $\forall i \in N$, $\sum_j x^i_j = 1$, these are legal random variables. Moreover, let $\mathcal{X}$ be the chosen sequence of exchanges, and denote the optimal fractional solution of the LP by $\text{OPT}_f = \sum_{i,j} j \cdot x^i_j$; it holds that

$$\mathbb{E}[\text{cost}(\mathcal{X})] = \mathbb{E}\left[\sum_{i \in N} X^i\right] = \text{OPT}_f \quad . \tag{3}$$

Now, fix some alternative $a \neq a^*$. We wish to bound the probability that a^* does not beat a after the exchanges given by $\mathcal{X}$ are made in R^N.

Let Y^i, $i \in N$, be independent Bernoulli trials, such that $Y^i = 1$ iff aR^ia^*, and a^* is moved above a in the preferences of agent i. In other words, $Y^i = 1$ if agent i becomes an additional agent that ranks a^* above a as a result of the exchanges. We want to provide an upper bound on $\Pr[\sum_{i \in N} Y^i < \text{def}(a)]$. Denote

$$p^i = \sum_{j:\ e^i_{ja}=1} x^i_j \quad .$$

Notice that $Y^i = 1$ with probability p^i, so $\mathbb{E}[\sum_i Y^i] = \sum_i p^i$. Moreover, by the constraint $\forall a \in A \setminus \{a^*\}$, $\sum_{i,j} x^i_j e^i_{ja} \geq \text{def}(a)$, we have that $\sum_i p^i \geq \text{def}(a)$. We now employ a deceivingly intuitive but nontrivial result:

Lemma A.1 (Jogdeo and Samuels [16]). *Let $Y^1, \ldots, Y^n$ be independent heterogeneous Bernoulli trials. Suppose that $\mathbb{E}[\sum_i Y^i]$ is an integer. Then*

$$\Pr\left[\sum_i Y^i < \mathbb{E}\left[\sum_i Y^i\right]\right] < 1/2 \quad .$$

Since def(a) is an integer, and $\mathbb{E}[\sum_i Y^i] = \sum_i p^i \geq \text{def}(a)$, it follows from the lemma that:

$$\Pr[a \text{ not beaten in } \mathcal{X}] = \Pr\left[\sum_i Y^i < \text{def}(a)\right] < 1/2 \quad .$$

At this point, we choose the value of the constant α to be such that $2^{\alpha \log m} \geq 4m$. Note that if $m \geq 4$, we can choose $\alpha \leq 2$. As in the algorithm, set $X^i_{max} = \max_k X^i_k$. Denote by $\mathcal{X}'$ the induced sequence of exchanges. It holds that a is not beaten in a pairwise election under $\mathcal{X}'$ only if a is not beaten under the exchanges obtained in each one of the $\alpha \cdot \log m$ individual iterations. Therefore,

$$\Pr[a \text{ not beaten in } \mathcal{X}'] < \left(\frac{1}{2}\right)^{\alpha \cdot \log m} \leq \frac{1}{4m} \quad .$$

By the union bound we get:[5]

$$\Pr[a^* \text{ is not a Condorcet winner in } \mathcal{X}'] \leq m \cdot \frac{1}{4m} = 1/4 \quad . \tag{4}$$

$X^i_1, \cdots, X^i_{\alpha \log m}$ are i.i.d. random variables; it holds that

$$X^i_{max} = \max_k X^i_k \leq \sum_k X^i_k \quad ,$$

and thus

$$\mathbb{E}\left[X^i_{max}\right] \leq \mathbb{E}\left[\sum_k X^i_k\right] = \alpha \cdot \log m \cdot \mathbb{E}[X^i_1] \quad . \tag{5}$$

Therefore, by the linearity of expectation,

$$\begin{aligned}
\mathbb{E}[\text{cost}(\mathcal{X}')] &= \mathbb{E}\left[\sum_i X^i_{max}\right] \\
&\leq \alpha \cdot \log m \cdot \mathbb{E}\left[\sum_i X^i_1\right] \\
&= \alpha \cdot \log m \cdot \mathbb{E}[\text{cost}(\mathcal{X})] \\
&= \alpha \cdot \log m \cdot \text{OPT}_f \\
&\leq \alpha \cdot \log m \cdot \text{OPT} \quad ,
\end{aligned}$$

where OPT is the Dodgson score of a^*, i.e., the optimal integral solution to the ILP (1).

By Markov's inequality we have that

$$\Pr[\text{cost}(\mathcal{X}') > \text{OPT} \cdot 4\alpha \cdot \log m] \leq 1/4 \quad . \tag{6}$$

We now apply the union bound once again on (4) and (6), and obtain that with probability at least 1/2, a^* is a Condorcet winner under $\mathcal{X}'$ and, at the same time, $\text{cost}(\mathcal{X}') \leq \text{OPT} \cdot 4 \cdot \alpha \cdot \log m$. This completes the proof of Theorem 3.1. □

[5]Strictly speaking, we can use $m-1$ instead of m.

B Proof of Theorem 3.2

We may view the problem of approximating the Dodgson score as the following covering problem with different covering requirements and constraints. The ground set is the set of live alternatives. For each live alternative $a \in A \setminus \{a^*\}$, its deficit def($a$) is in fact its covering requirement, i.e., the number of different sets it has to belong to in the final cover. For each agent $i \in N$ that ranks a^* in place r^i, we have a subcollection $\mathcal{S}^i$ consisting of the sets S_k^i for $k = 1, \ldots, r^i - 1$, where the set S_k^i contains the (initially) live alternatives that appear in positions $r^i - k$ to $r^i - 1$ in the preference of agent i. The set S_k^i has cost k. Now, the covering problem to be solved is the following. We wish to select at most one set from each of the different subcollections so that each alternative $a \in A \setminus \{a^*\}$ appears in at least def(a) sets and the total cost of the selected sets is minimized. The optimal cost is the Dodgson score of a^* and, hence, the cost of any approximate cover that satisfies the covering requirements and the constraints is an upper bound on the Dodgson score.

In terms of this covering problem, the greedy algorithm mentioned above can be thought of as working as follows. In each step, it selects the most cost-effective set where the cost-effectiveness of a set is defined as the ratio between the cost of the set and the number of live alternatives it covers that have not been previously covered by sets belonging to the same subcollection. For these live alternatives, the algorithm decreases their covering requirements at the end of the step. The algorithm terminates when all alternatives have died (i.e., their covering requirement has become zero). The output of the algorithm consists of the maximum-cost sets that were picked from each subcollection.

We remark that the covering problem we use is closely related to the Constrained Set Multicover problem considered in Rajagopalan and Vazirani [28] (see also [35, pp. 112–116]), with the additional constraint that at most one set has to be selected from each subcollection.

We now turn to the formal part of the proof. We find it convenient to formulate the Dodgson score problem as the following integer linear program, which is very similar but not identical to the LP (1) that we have used for the randomized rounding algorithm.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{i \in N} \sum_{k=1}^{r^i - 1} k \cdot x_{S_k^i} \\
\text{subject to} \quad & \forall a \in A \setminus \{a^*\}, \ \sum_{i \in N} \sum_{S \in \mathcal{S}^i : a \in S} x_S \geq \text{def}(a) \\
& \forall i \in N, \ \sum_{S \in \mathcal{S}^i} x_S \leq 1 \\
& x \in \{0, 1\}
\end{aligned}
$$

The variable x_S associated with a set S denotes whether S is included in the solution ($x_S = 1$) or not ($x_S = 0$). We relax the integrality constraint in order to obtain a linear programming relaxation and we compute its dual linear program.

$$
\begin{aligned}
\text{maximize} \quad & \sum_{a\in A\setminus\{a^*\}} \mathrm{def}(a)\cdot y_a - \sum_{i\in N} z^i \\
\text{subject to} \quad & \forall i\in N, k=1,\ldots,r^i-1, \ \sum_{a\in S_k^i} y_a - z^i \le k \\
& \forall i \in N,\ z^i \ge 0 \\
& \forall a \in A\setminus\{a^*\},\ y_a \ge 0
\end{aligned}
$$

For a set S that is picked by the algorithm to cover alternative $a \in A \setminus \{a^*\}$ for the j-th time (the j-th copy of a), we set $p(a,j)$ to be equal to the cost-effectiveness of S when it is picked. Informally, p distributes equally the cost of S among the copies of the live alternatives it covers. For each $i \in N$ and $k = 1,\ldots,r^i-1$, we define a set T_k^i that, if S_k^i was picked by the algorithm, contains the alternatives in S_k^i that were alive when S_k^i was picked (namely those whose covering requirement decreased), and is empty otherwise. In the case where $a \in T_k^i$, we use $j(a,S_k^i)$ to denote the index of the copy of a that S_k^i covers when it is picked.

Now, we shall show that by setting

$$
y_a \quad = \quad \frac{p(a,\mathrm{def}(a))}{H_{m-1}}
$$

for each alternative $a \in A \setminus \{a^*\}$, where H_{m-1} is the $m-1$ harmonic number, and

$$
z^i = \frac{1}{H_{m-1}} \sum_{k=1}^{r^i-1} \sum_{a\in T_k^i} (p(a,\mathrm{def}(a)) - p(a,j(a,S_k^i)))
$$

for each agent $i \in N$, the constraints of the dual linear program are satisfied. The variables y_a are clearly non-negative. Since the algorithm always selects the most cost-effective set, it holds that $p(a,\mathrm{def}(a)) \ge p(a,j)$ for every alternative a with $\mathrm{def}(a) > 0$ and $j \le \mathrm{def}(a)$ and, hence, z^i is non-negative.

In order to show that the first constraint of the dual linear program is also satisfied, consider an agent $i \in N$ and integer λ such that $1 \le \lambda \le r^i - 1$. We have

$$
\begin{aligned}
\sum_{a\in S_\lambda^i} y_a - z^i \quad &= \quad \frac{1}{H_{m-1}} \left[\sum_{a\in S_\lambda^i} p(a,\mathrm{def}(a)) - \sum_{k=1}^{r^i-1}\sum_{a\in T_k^i} (p(a,\mathrm{def}(a)) - p(a,j(a,S_k^i))) \right] \\
&\le \quad \frac{1}{H_{m-1}} \left[\sum_{a\in S_\lambda^i} p(a,\mathrm{def}(a)) - \sum_{a\in T_\lambda^i} (p(a,\mathrm{def}(a)) - p(a,j(a,S_\lambda^i))) \right] \\
&= \quad \frac{1}{H_{m-1}} \left[\sum_{a\in S_\lambda^i\setminus T_\lambda^i} p(a,\mathrm{def}(a)) + \sum_{a\in T_\lambda^i} p(a,j(a,S_\lambda^i)) \right] \qquad (7)
\end{aligned}
$$

Let $s = |S_\lambda^i|$ and $s' = |S_\lambda^i \setminus T_\lambda^i|$. We number the alternatives of $S_\lambda^i \setminus T_\lambda^i$ in the order in which they die during the execution of the algorithm. Let this order be $a_1, a_2, \ldots, a_{s'}$. When a_t dies, we have that

$$
p(a_t,\mathrm{def}(a_t)) \le \frac{\lambda}{s-t+1} \qquad (8)
$$

since otherwise the set S^i_λ would have been used to cover a_t. We distinguish between the following two cases:
Case 1: If $T^i_\lambda = \emptyset$, using inequalities (7) and (8), we obtain

$$\sum_{a \in S^i_\lambda} y_a - z^i \quad \leq \quad \frac{1}{H_{m-1}} \sum_{t=1}^{s} p(a_t, \mathrm{def}(a_t)) \leq \frac{1}{H_{m-1}} \sum_{t=1}^{s} \frac{\lambda}{s-t+1} \leq \lambda \quad .$$

Case 2: If $T^i_\lambda \neq \emptyset$ then $s' \leq s-1$. We have

$$\sum_{a \in T^i_\lambda} p(a, j(a, S^i_\lambda)) \quad = \quad \lambda \quad ,$$

since the cost of the set S^i_λ is equally distributed among the copies it covers. Also, using (8),

$$\sum_{a \in S^i_\lambda \setminus T^i_\lambda} p(a, \mathrm{def}(a)) \quad = \quad \sum_{t=1}^{s'} p(a_t, \mathrm{def}(a_t)) \leq \sum_{t=1}^{s-1} \frac{\lambda}{s-t+1} = \lambda(H_{s-1} - 1) \quad .$$

So, inequality (7) again yields

$$\sum_{a \in S^i_\lambda} y_a - z^i \quad \leq \quad \lambda \quad ,$$

implying that the constraints of the dual linear program are always satisfied.

Now, denote by OPT the optimal objective value of the integer linear program. By duality, we have that any feasible solution to the dual of its linear programming relaxation has objective value at most OPT. Hence,

$$\begin{aligned}
H_{m-1} \cdot \mathrm{OPT} \quad &\geq \quad H_{m-1} \left(\sum_{a \in A \setminus \{a^*\}} \mathrm{def}(a) \cdot y_a - \sum_{i \in N} z^i \right) \\
&= \quad \sum_{a \in A \setminus \{a^*\}} \mathrm{def}(a) \cdot p(a, \mathrm{def}(a)) - \sum_{i \in N} \sum_{k=1}^{r^i - 1} \sum_{a \in T^i_k} \left(p(a, \mathrm{def}(a)) - p(a, j(a, S^i_k)) \right) \\
&= \quad \sum_{i \in N} \sum_{k=1}^{r^i - 1} \sum_{a \in T^i_k} p(a, j(a, S^i_k)) \\
&= \quad \sum_{i \in N} \sum_{k \in \{1, \ldots, r^i - 1\} : T^i_k \neq \emptyset} k \quad .
\end{aligned}$$

The theorem follows since the last expression clearly upper-bounds the cost of the algorithm. □

C Proof of Theorem 3.3

Our inapproximability result for Dodgson score uses a reduction from Minimum Set Cover and the following well-known statements of its inapproximability.

Theorem C.1 (Raz and Safra [31]). *There exists a constant $\alpha > 0$ such that, given an instance $(U, \mathcal{S})$ of Minimum Set Cover with $|U| = n$ and an integer $K \le n$, it is $\mathcal{NP}$-hard to distinguish between the following two cases:*

- *$(U, \mathcal{S})$ has a cover of size at most K.*
- *Any cover of $(U, \mathcal{S})$ has size at least $\alpha K \ln n$.*

Theorem C.2 (Feige [10]). *For any constant $\epsilon > 0$, given an instance $(U, \mathcal{S})$ of Minimum Set Cover with $|U| = n$ and an integer $K \le n$, there is no polynomial-time algorithm that distinguishes between the following two cases:*

- *$(U, \mathcal{S})$ has a cover of size at most K, and*
- *Any cover of $(U, \mathcal{S})$ has size at least $(1 - \epsilon)K \ln n$,*

unless $\mathcal{NP} \subseteq DTIME(n^{O(\log \log n)})$.

Given an instance of Minimum Set Cover consisting of a set of n elements, a collection of sets over these elements and an integer $K \le n$, we construct a preference profile with $m = (1 + \zeta)n + \lceil \alpha\zeta K n \ln n \rceil + 1$ alternatives and a specific alternative a^* in which we show that if we could distinguish in polynomial time between the following two cases:

- a^* has Dodgson score at most $(1 + \zeta)Kn$, and
- a^* has Dodgson score at least $\alpha\zeta K n \ln n$,

then we could have distinguished between the two cases of Theorems C.1 and C.2 for the original Minimum Set Cover instance, contradicting the above inapproximability statements. Here, α is the inapproximability constant in Theorem C.1 or C.2 (in the latter $\alpha = 1 - \epsilon$), and ζ is an arbitrarily large positive constant. In this way, we obtain an inapproximability bound of $\frac{\alpha\zeta}{1+\zeta} \ln n$. Since $m = (1 + \zeta)n + \lceil \alpha\zeta K n \ln n \rceil + 1$, it holds that $\ln n \ge \frac{1}{2} \ln m - \mathcal{O}(\ln \ln m)$, and hence the inapproximability bound for Dodgson score can be expressed in terms of the number of alternatives m as stated in Theorem 3.3.

We now present our reduction. Given an instance $(U, \mathcal{S})$ of Minimum Set Cover consisting of a set U of n elements, a collection $\mathcal{S}$ of sets $S_1, S_2, \ldots, S_{|\mathcal{S}|}$ and an integer $K \le n$, we construct the following preference profile. There are the following alternatives:

- A set of n basic alternatives each corresponding to an element of U.
- A set Z of ζn alternatives where ζ is a positive constant.
- A set F of $\lceil \alpha\zeta K n \ln n \rceil$ alternatives, where α is the constant from Theorem C.1.
- A specific alternative a^*.

There are the following $2|\mathcal{S}| + 1$ agents:

- A critical agent ℓ^i for each set $S_i \in \mathcal{S}$.
- An indifferent agent r^i for each set $S_i \in \mathcal{S}$.

- A special agent v^*.

The preferences of the agents are defined as follows:

- The special agent v^* ranks a^* in the first position of its preferences and the rest of the alternatives occupy the remaining positions in arbitrary order (i.e., $a^* \, R^{v^*} U \cup Z \cup F$).
- The critical agent ℓ^i ranks the basic alternatives corresponding to the elements of S_i in the first positions of its preference (in arbitrary order), next the alternatives of Z, next a^*, next the alternatives of F, and, in the last positions of its preference, the basic alternatives corresponding to the elements in $U \backslash S_i$ (i.e., $S_i \, R^{\ell^i} Z \, R^{\ell^i} a^* \, R^{\ell^i} F \, R^{\ell^i} U \backslash S_i$).
- We construct the ranking of the indifferent agents as follows: For each element u of U, we remove u from one of the sets of $\mathcal{S}$ (selecting arbitrarily among the sets of $\mathcal{S}$ containing u). Let $S'_1, S'_2, \ldots, S'_{|\mathcal{S}|}$ be the resulting sets and denote by $\mathcal{S}'$ their collection. The indifferent agent r^i ranks the basic alternatives corresponding to the elements in $U \setminus S'_i$ in the first positions of its preference, next the alternatives of F, next a^*, next the alternatives of Z and, in the last positions of its preference, the basic alternatives corresponding to elements in S'_i (if any). I.e., $U \setminus S'_i \, R^{r^i} F \, R^{r^i} a^* \, R^{r^i} Z \, R^{r^i} S'_i$.

Clearly, a^* is preferred to any alternative in Z by the special agent and by the $|\mathcal{S}|$ indifferent agents, namely by a majority of agents. Similarly, a^* is preferred to any alternative in F by the special agent and by the $|\mathcal{S}|$ critical agents. Now, for each element of U, denote by f_u the number of sets in $\mathcal{S}$ that contain u. Then, a^* is preferred to u by the special agent, by the $|\mathcal{S}| - f_u$ critical agents corresponding to sets in $\mathcal{S}$ that do not contain u, and by the $f_u - 1$ indifferent agents corresponding to sets in $\mathcal{S}'$ that contain u (i.e., by $|\mathcal{S}|$ agents in total). Hence, a^* has a deficit of exactly 1 with respect to each of the alternatives in U.

Theorem 3.3 follows by the next two lemmata that give bounds on the Dodgson score of alternative a^* in the two cases of interest: when $(U, \mathcal{S})$ has a cover of size at most K (Lemma C.3) and when any cover of $(U, \mathcal{S})$ has size at least $\alpha K \ln n$ (Lemma C.4).

Lemma C.3. *If $(U, \mathcal{S})$ has a cover of size K, then a^* has Dodgson score at most $(1+\zeta)Kn$.*

Proof. Let $H \subseteq \mathcal{S}$ be a cover for $(U, \mathcal{S})$ with $|H| = K$. By the definition of a cover, H covers all elements of U. Hence, by pushing a^* to the first position in the preference of the critical agent ℓ^i such that $S_i \in H$, a^* will decrease its deficit with respect to each of the basic alternatives by 1, and hence it will become a Condorcet winner. The total number of positions a^* rises is at most $|H| \cdot (|Z| + n) = (1+\zeta)nK$. □

Lemma C.4. *If any cover of $(U, \mathcal{S})$ has size at least $\alpha K \ln n$, then a^* has Dodgson score at least $\alpha \zeta K n \ln n$.*

Proof. We first assume that the minimum number of positions a^* has to rise in order to beat the basic alternatives and become a Condorcet winner includes raising a^* by at least $|F|$ positions in the ranking of some indifferent agent r^i. Hence, a^* rises $|F|$ positions in the preference of r^i in order to reach position $|U \backslash S'_i| + 1$ and at least n additional positions in order to beat the basic alternatives. Hence, its Dodgson score is at least $|F| + n \geq \alpha \zeta K n \ln n$.

Now, assume that the minimum number of positions a^* has to rise in order to beat the basic alternatives does not include raising a^* by at least $|F|$ positions in the ranking of some indifferent agent. We will show that if the Dodgson score of a^* is less than $\alpha\zeta K n \ln n$, then there exists a cover of $(U, \mathcal{S})$ of size less than $\alpha K \ln n$, contradicting the assumption of the lemma.

Let H be the set of critical agents such that a^* is pushed $|Z|$ positions in their preferences. In total, a^* rises $|H| \cdot |Z|$ positions in order to reach position $|S_i| + 1$ in the preferences of each critical agent ℓ^i belonging to H, plus at least n additional positions in order to decrease its deficit with respect to each of the alternatives in U by 1, that is, at least $\zeta|H|n + n$ positions in total. Hence, by denoting the Dodgson score of a^* by $sc_D(a^*)$, we have $|H| \leq \frac{1}{\zeta n} sc_D(a^*) - \frac{1}{\zeta} < \alpha K \ln n$. The proof is completed by observing that the union of the sets S_i for each critical agent ℓ^i belonging to H contains all the basic alternatives, i.e., H corresponds to a cover for $(U, \mathcal{S})$ of size less than $\alpha K \ln n$. □

D Proof of Theorem 3.4

Our proof relies on the inapproximability of Minimum Vertex Cover in 3-regular graphs [3]. In particular, our reduction uses the following very weak variant of an inapproximability result presented in [3].

Theorem D.1 (Berman and Karpinski [3], see also [17]). *Given a 3-regular graph G with $n = 22t$ nodes for some integer $t > 0$ and an integer K in $[n/2, n-6]$, it is $\mathcal{NP}$-hard to distinguish between the following two cases:*

- *G has a vertex cover of size at most K.*
- *Any vertex cover of G has size at least $K + 6$.*

Given an instance of Minimum Vertex Cover consisting of a 3-regular graph G with $n = 22t$ nodes $v_0, v_1, \ldots, v_{n-1}$ and an integer $K \in [n/2, n-6]$, we construct in polynomial time a preference profile in which if we could distinguish whether a particular alternative is a Dodgson winner or not very far from the last position in any Dodgson ranking, then we could also distinguish between the two cases mentioned in Theorem D.1 for the original Minimum Vertex Cover instance. The Dodgson election has the following sets of alternatives:

- A special alternative a^*.
- A set F of $4Kn/11 + 3n/2$ alternatives. These alternatives are partitioned into n disjoint blocks F_0, F_1, ..., F_{n-1} so that each block contains either $\lceil 4K/11 + 3/2 \rceil$ or $\lfloor 4K/11 + 3/2 \rfloor$ alternatives.
- A set A of n alternatives $a_0, a_1, \ldots, a_{n-1}$.
- An alternative u_j for each edge e_j of G. Let U be the set of these $3n/2$ alternatives. We denote by S_i the set of the three alternatives of U that correspond to the edges of G which are incident to node v_i.

For each node v_i of G, there are two agents: one left agent ℓ^i and one right agent r^i. The preferences of the left agent ℓ^i are as follows:

- The three alternatives of S_i are ranked by agent ℓ^i in the first three positions of its preference (in arbitrary order).

- From position 4 to position $4n/11 + 3$, agent ℓ^i ranks the alternatives $a_i, a_{(i+1) \bmod n}, \dots, a_{(i+4n/11-1) \bmod n}$ in this order.
- In position $4n/11 + 4$, agent ℓ^i ranks a^*.
- From position $4n/11 + 5$ to position $4Kn/11 + 41n/22 + 4$, agent ℓ^i ranks the alternatives of F in the following order: The alternatives of F_i are ranked in positions from $4n/11 + 5$ to $4n/11 + 4 + |F_i|$ (in arbitrary order). Next, agent ℓ^i ranks the alternatives of sets $F_0, \dots, F_{i-1}, F_{i+1}, \dots, F_{n-1}$ in this order (the relative order of the alternatives of the same block is arbitrary).
- From position $4Kn/11 + 41n/22 + 5$ to position $4Kn/11 + 5n/2 + 4$, agent ℓ^i ranks the alternatives $a_{(i+4n/11) \bmod n}, a_{(i+4n/11+1) \bmod n}, \dots, a_{(i-1) \bmod n}$ in this order.
- In the last $3n/2 - 3$ positions, agent ℓ^i ranks the alternatives of $U \setminus S_i$ (in arbitrary order).

The preferences of the right agent r^i are as follows:

- In the first $3n/2 - 3$ positions, agent r^i ranks the alternatives of $U \setminus S_i$ in reverse relative order to the order ℓ^i ranks them.
- From position $3n/2 - 2$ to position $4Kn/11 + 3n - |F_i| - 3$, agent r^i ranks the alternatives of the blocks $F_{n-1}, F_{n-2}, \dots, F_{i+1}, F_{i-1}, \dots, F_0$ in this order so that the alternatives of block F_j are ranked in reverse relative order to the order ℓ^i ranks them.
- From position $4Kn/11 + 3n - |F_i| - 2$ to position $4Kn/11 + 40n/11 - |F_i| - 5$, agent r^i ranks the alternatives $a_{(n-i-1) \bmod n}, a_{(n-i-2) \bmod n}, \dots, a_{(4n/11-i+2) \bmod n}$ in this order.
- In position $4Kn/11 + 40n/11 - |F_i| - 4$, agent r^i ranks a^*.
- From position $4Kn/11 + 40n/11 - |F_i| - 3$ to position $4Kn/11 + 40n/11 - 4$, agent r^i ranks the alternatives of F_i in reverse relative order to the order ℓ^i ranks them.
- From position $4Kn/11+40n/11-3$ to position $4Kn/11+4n-2$, agent r^i ranks the alternatives $a_{(4n/11-i+1) \bmod n}$, $a_{(4n/11-i) \bmod n}$, ..., $a_{(n-i) \bmod n}$ in this order.
- The three alternatives of S_i are ranked in the last three positions in the preference of agent r^i, in reverse relative order to the order ℓ^i ranks them.

We observe that a^* beats all alternatives but the alternatives of U. In particular, a^* is preferred to each alternative of F by $n + 1$ agents. Specifically, a^* is ranked above an alternative belonging to the block F_i by the n left agents and by the right agent r^i. Also, the alternative a_i is ranked below a^* by the $7n/11$ left agents $\ell^{(i+1) \bmod n}, \ell^{(i+2) \bmod n}, \dots, \ell^{(i+7n/11-1) \bmod n}$ and by the $4n/11+2$ right agents $r^{(i+7n/11-1) \bmod n}, r^{(i+7n/11) \bmod n}, \dots, r^i$. Hence, a^* beats all alternatives in set A as well since it is ranked above each of them by $n+2$ agents. Also, a^* is ranked above the alternative u_j corresponding to the edge e_j of G by the left agents ℓ^i and $\ell^{i'}$ and by all right agents besides r^i and $r^{i'}$ so that nodes v_i and $v_{i'}$ are the endpoints of edge e_j in G. Hence, a^* has a deficit of 1 with respect to each of the alternatives in U.

We also observe that the alternatives in F beat each alternative in A. Note that in each agent where a^* is preferred to an alternative in A besides by the right agent r^i, an alternative of block

F_i is also preferred to the alternative in A. Hence, each alternative of F beats each alternative of A since it is ranked above it by $n+1$ agents. Furthermore, similarly to a^*, each alternative in F is preferred to each alternative of U by n agents. Also, when an alternative f of F is ranked above another alternative f' of F by agent ℓ^i, f' is ranked above f by agent r^i. Hence, an alternative of F has a deficit of 1 with respect to U and each other alternative in F, and a deficit of 2 with respect to a^*.

Furthermore, observe that each alternative in A is ranked above the alternative u_j corresponding to the edge e_j of G by the left agents ℓ^i and $\ell^{i'}$ and by all right agents besides r^i and $r^{i'}$ so that nodes v_i and $v_{i'}$ are the endpoints of edge e_j in G, i.e., by n agents. Also, when an alternative a of A is preferred to another alternative a' of A by agent ℓ^i, a' is preferred to a by agent r^i. Hence, an alternative in A has a deficit of 1 with respect to each of the alternatives in U and F and each other alternative in A, and a deficit of 3 with respect to a^*. This immediately yields that the Dodgson score of each alternative in U is at least $4Kn/11 + 4n + 2$.

Similarly, when an alternative u of U is preferred to another alternative u' of U by agent ℓ^i, u' is preferred to u by agent r^i. Hence, an alternative in U has a deficit of 1 with respect to each of the alternatives in A and F, each other alternative in U, and a^*. This immediately yields that the Dodgson score of each alternative in U is at least $4Kn/11 + 4n$.

The next lemma gives upper and lower bounds on the Dodgson score of the alternatives in F.

Lemma D.2. *Each alternative in F has Dodgson score between $4Kn/11 + 3n + 1$ and $4Kn/11 + 37n/11 + 2K/11 + 3/4$.*

Proof. Since each alternative in F has a deficit of 1 with respect to each alternative in U and each other alternative in F, and a deficit of 2 with respect to a^*, its Dodgson score is at least $|U| + |F| - 1 + 2 = 4Kn/11 + 3n + 1$.

Now, consider an alternative f belonging to block F_i. f is at distance at most

$$\lfloor \frac{|F_i| - 1}{2} + 1 \rfloor \leq \frac{|F_i| + 1}{2} \leq \frac{\lceil 4K/11 + 3/2 \rceil + 1}{2} \leq 2K/11 + 7/4$$

from a^* in the preferences of either the left agent ℓ^i or the right agent r^i. Hence, by raising f at most $2K/11 + 7/4$ positions in the preferences of either ℓ^i or r^i, its deficit with respect to a^* decreases by 1. Consider a left agent $\ell^{i'}$ with $i' \neq i$ and let F' be the subset of alternatives in F that are higher than f in the preferences of $\ell^{i'}$. By pushing f to the first position in the preferences of agent $\ell^{i'}$ (i.e., $4n/11 + 3 + |F'|$ additional positions), f decreases its deficit by 1 with respect to each alternative of F' and a^*, as well as with respect to the three alternatives of $S_{i'}$ in the first three positions in the preferences of $\ell^{i'}$. Now, consider the right agent $r^{i'}$. In the preferences of $r^{i'}$, f is ranked higher than the alternatives in F' and lower than the alternatives in $F \backslash F' - \{f\}$. Hence, by pushing f to the first position in the preferences of agent $r^{i'}$ (i.e., $|F \backslash F' - \{f\}| + 3n/2 - 3 = 4Kn/11 - |F'| + 3n - 4$ additional positions), f decreases its deficit by 1 with respect to each alternative of $F \backslash F' - \{f\}$ as well the alternatives of $U \backslash S_{i'}$ in the first $3n/2 - 3$ positions in the preferences of r^i. Hence, by pushing $4Kn/11 + 37n/11 + 2K/11 + 3/4$ positions, f becomes a Condorcet winner. □

The next two lemmata give bounds on the Dodgson score of alternative a^* in the two cases of interest: when G has a vertex cover of size at most K (Lemma D.3), and when any vertex cover of G has size at least $K + 6$ (Lemma D.4).

Lemma D.3. *If G has a vertex cover of size at most K, then the Dodgson score of a^* is less than $4Kn/11 + 3n$.*

Proof. Let $H \subseteq V$ be a vertex cover of G with $|H| = K$. By the definition of the vertex cover, H covers all edges of G and this implies that $\cup_{i:v_i \in H} S_i = U$. Hence, by pushing a^* to the first position in the preferences of each of the K left agents ℓ^i such that $v_i \in H$, a^* decreases it deficit with respect to each of the alternatives in U by 1, and becomes a Condorcet winner. The total number of positions a^* rises is $K(4n/11 + 3) < 4Kn/11 + 3n$. The last inequality is true since $K < n$. □

Lemma D.4. *If any vertex cover of G has size at least $K + 6$, then the Dodgson score of a^* is larger than $4Kn/11 + 37n/11 + 2K/11 + 3/4$.*

Proof. First assume that the minimum sequence of exchanges that makes a^* beat the alternatives of U and become a Condorcet winner includes pushing a^* to one of the first $3n/2 - 3$ positions in the preferences of some right agent r^i. Certainly, not all alternatives of U are beaten in this way since the three alternatives of S_i are ranked below a^* by agent r^i. So, in order to beat the remaining 3 alternatives of S_i, a^* has to either be pushed to one of the first three positions of a left agent or to one of the first $3n/2 - 3$ positions of another right agent $r^{i'}$ with $i' \neq i$. Hence, a^* must be first pushed to position $3n/2 - 2$ of agent r^i (i.e., $|F \backslash F_i| + 7n/11 - 2$ positions), to position 4 of a left agent (i.e., $4n/11$ additional positions) or to position $3n/2 - 2$ of agent $r^{i'}$ (i.e., $|F \backslash F_{i'}| + 7n/11 - 2$ additional positions), and then rise at least $3n/2$ additional positions in order to beat all alternatives of U. In total, a^* rises at least

$$
\begin{aligned}
& |F \backslash F_i| + 7n/11 - 2 + \min\{4n/11, |F \backslash F_{i'}| + 7n/11 - 2\} + 3n/2 \\
\geq \;& |F| - |F_i| + 5n/2 - 2 \\
\geq \;& 4Kn/11 + 4n - \lceil 4K/11 + 3/2 \rceil - 2 \\
\geq \;& 4Kn/11 + 4n - 4K/11 - 9/2 \\
= \;& 4Kn/11 + 37n/11 + n/11 + 6n/11 - 4K/11 - 9/2 \\
\geq \;& 4Kn/11 + 37n/11 + 22/11 + 6(K+6)/11 - 4K/11 - 9/2 \\
> \;& 4Kn/11 + 37n/11 + 2K/11 + 3/4
\end{aligned}
$$

positions. The fourth inequality holds since $n \geq 22$ and $n \geq K + 6$.

Now, assume that the minimum sequence of exchanges for making a^* a Condorcet winner does not include raising a^* to any of the first $3n/2 - 3$ positions of any right agent. We will show that if a^* has Dodgson score at most $4Kn/11 + 37n/11 + 2K/11 + 3/4$, then G has a vertex cover of size less than $K + 6$, contradicting the assumption of the lemma.

Let H be the set of left agents where a^* rises to one of the first three positions in order to beat all the alternatives of U. In total, a^* rises $4|H|n/11$ positions in order to reach position 4 in the preferences of each of the agents in H plus at least $3n/2$ additional positions in order to decrease its deficit with respect to the alternatives in U by at least 1, i.e., at least $4|H|n/11 + 3n/2$ positions in total. Hence, by denoting the Dodgson score of a^* by $sc_D(a^*)$, we have $|H| \leq \frac{11}{4n}(sc_D(a^*) - 3n/2)$.

Since $\cup_{i:\ell^i \in H} S_i = U$, the set of nodes of G consisting of nodes v_i such that agent ℓ^i belongs to H is a vertex cover of G of size $|H|$. Assuming that the Dodgson score of a^* is at most $4Kn/11 + 37n/11 + 2K/11 + 3/4$, we have

$$
\begin{aligned}
|H| &\leq \frac{11}{4n}(sc_D(a^*) - 3n/2) \\
&\leq \frac{11}{4n}(4Kn/11 + 41n/22 + 2K/11 + 3/4) \\
&< K + 6 \;\; ,
\end{aligned}
$$

where the last inequality follows since $K \le n-6$. □

By Lemmata D.2, D.3, and D.4, we obtain that if G has a vertex cover of size at most K, then a^* is the unique Dodgson winner, while if every vertex cover of G has size at least $K+6$, then a^* is below all alternatives in F in any Dodgson ranking. Denote by $m = |F|+|A|+|U|+1 = 4Kn/11+4n+1$ the total number of alternatives. Then, the rank of a^* in the second case is at least

$$\begin{aligned} |F|+1 &= 4Kn/11+3n/2+1 = m-5n/2 = m-\sqrt{25n^2/4} \\ &\ge m-\sqrt{25nK/2} \ge m-6\sqrt{4Kn/11+4n+1} = m-6\sqrt{m} \quad , \end{aligned}$$

where the first inequality follows since $K \ge n/2$. By Theorem D.1, we obtain the desired result. □

D.1 An example

We present an example of the construction in the proof of Theorem 3.4. Consider an instance of Minimum Vertex Cover with the 22-node 3-regular graph of Figure 2, and $K = 12$.

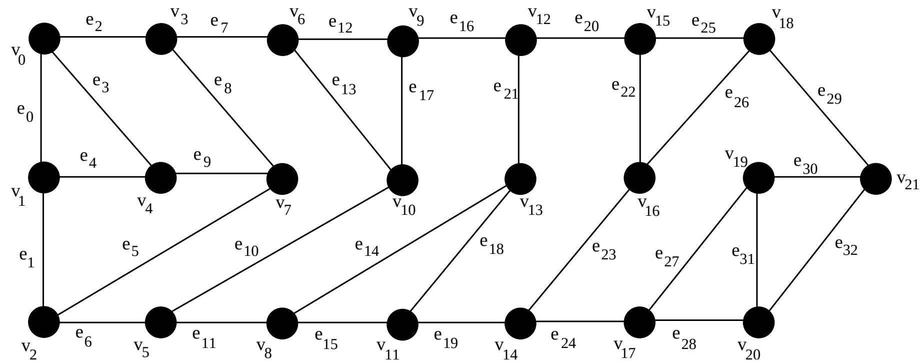

Figure 2: A 3-regular graph with 22 nodes.

The corresponding preference profile has 185 alternatives and 44 agents. In particular, the set F has 129 alternatives $f_0, f_1, \ldots, f_{128}$, which are partitioned in 22 blocks as follows. Block F_0 contains the six alternatives $f_0, f_1, \ldots, f_5$, block F_1 contains the six alternatives $f_6, \ldots, f_{11}$, …, block F_{18} contains the six alternatives $f_{108}, \ldots, f_{113}$, block F_{19} contains the five alternatives $f_{114}, \ldots, f_{118}$, …, and block F_{21} contains the alternatives $f_{124}, \ldots, f_{128}$. The set A has 22 alternatives $a_0, \ldots, a_{21}$. The set U has 33 alternatives $u_0, \ldots, u_{32}$, one alternative for each edge of the graph. The agents are partitioned into 22 left agents and 22 right agents. In order to compute the preferences of an agent, say agent ℓ^{17}, we first compute the set S_{17}, which contains the alternatives corresponding to the edges incident to node v_{17} of the graph, i.e., $S_{17} = \{u_{24}, u_{27}, u_{28}\}$. Now, the preferences of agent ℓ^{17} are:

$$\begin{aligned} &S_{17}\, R^{\ell^{17}} a_{17}\, R^{\ell^{17}} a_{18}\, R^{\ell^{17}} \ldots R^{\ell^{17}} a_1\, R^{\ell^{17}} a_2\, R^{\ell^{17}} a^*\, R^{\ell^{17}} F_{17}\, R^{\ell^{17}} F_0\, R^{\ell^{17}} \ldots F_{16} \\ &R^{\ell^{17}} F_{18}\, R^{\ell^{17}} \ldots R^{\ell^{17}} F_{21}\, R^{\ell^{17}} a_3\, R^{\ell^{17}} \ldots R^{\ell^{17}} a_{16}\, R^{\ell^{17}} U \setminus S_{17} \quad . \end{aligned}$$

Similarly, the preferences of agent r^{17} are:

$$\begin{aligned} &\overleftarrow{U \setminus S_{17}}\ R^{\ell^{17}}\ \overleftarrow{F_{21}}\ R^{\ell^{17}}\ \overleftarrow{F_{20}}\ R^{\ell^{17}} \ldots R^{\ell^{17}}\ \overleftarrow{F_{18}}\ R^{\ell^{17}}\ \overleftarrow{F_{16}}\ R^{\ell^{17}} \ldots R^{\ell^{17}}\ \overleftarrow{F_0}\ R^{\ell^{17}} \\ &a_4\, R^{\ell^{17}} a_3\, R^{\ell^{17}} \ldots R^{\ell^{17}} a_{15}\, R^{\ell^{17}} a^*\, R^{\ell^{17}}\ \overleftarrow{F_{17}}\ a_{16}\, R^{\ell^{17}} a_5\, R^{\ell^{17}}\ \overleftarrow{S_{17}} \quad , \end{aligned}$$

where the symbol $\leftarrow$ on top of a set of alternatives is used to denote that their order in the preferences of r^{17} is the reverse of the order ℓ^{17} ranks them.

E Proof of Lemma 4.2

The problem is clearly in $\mathcal{NP}$; a witness is given by a nonempty set of agents for which a^* is a Condorcet winner.

In order to show $\mathcal{NP}$-hardness, we present a polynomial-time reduction from the $\mathcal{NP}$-hard Exact Cover by 3-Sets (X3C) problem [11] to our problem. An instance of the X3C problem includes a finite set of elements U, $|U| = n$ (where n is divisible by 3), and a collection $\mathcal{S}$ of 3-element subsets of U, $\mathcal{S} = \{S_1, \ldots, S_k\}$, such that for every $1 \leq i \leq k$, $S_i \subseteq U$ and $|S_i| = 3$. The question is whether the collection $\mathcal{S}$ contains an *exact cover* for U, i.e., a subcollection $\mathcal{S}^* \subseteq \mathcal{S}$ of size $n/3$ such that every element of U occurs in exactly one subset in $\mathcal{S}$.

We next give the details of the reduction from X3C to NonEmptySubset. Given an instance of X3C, defined by the set U and a collection of 3-element sets $\mathcal{S}$, we construct the following instance of NonEmptySubset.

Define the set of alternatives as $A = U \cup \{a\} \cup \{a^*\}$. Let the set of agents be $N = N' \cup N''$, where N' and N'' are defined as follows. The set N' is composed of k agents, corresponding to the k subsets in $\mathcal{S}$, such that for all $i \in N'$, agent i prefers the alternatives in $U \setminus S_i$ to a^*, and prefers a^* to all the alternatives in $S_i \cup \{a\}$ (i.e., $U \setminus S_i\ R^i\ a^* R^i\ S_i \cup \{a\}$).

Subset N'' is composed of $\frac{n}{3} - 1$ agents who prefer a to a^* and a^* to U (i.e., for all $i \in N''$, $a\ R^i\ a^*\ R^i\ U$).

We next show that there is an exact cover in the given instance iff there is nonempty subset of agents for which a^* is a Condorcet winner in the constructed instance.

Sufficiency: Let $\mathcal{S}^*$ be an exact cover by 3-sets of U, and let $N^* \subseteq N'$ be the subset of agents corresponding to the $\frac{n}{3}$ subsets $S_i \in \mathcal{S}^*$. We show that a^* is a Condorcet winner for $C = N^* \cup N''$. Since $\mathcal{S}^*$ is an exact cover, for all $b \in U$ there exists exactly one agent in N^* that prefers a^* to b and $\frac{n}{3} - 1$ agents in N^* that prefer b to a^*. In addition, all $\frac{n}{3} - 1$ agents in N'' prefer a^* to b. Therefore, a^* beats b in a pairwise election.

It remains to show that a^* beats a in a pairwise election. This is true since all $\frac{n}{3}$ agents in N^* prefer a^* to a, and there are only $\frac{n}{3} - 1$ agents in N'' who prefer a to a^*. It follows that a^* is a Condorcet winner for $N^* \cup N''$.

Necessity: Assume the given instance of X3C has no exact cover. We have to show that there is no subset of agents for which a^* is a Condorcet winner. Let $C \subseteq N$, $C \neq \emptyset$, and let $N^* = C \cap N'$. We distinguish between three cases.

Case 1: $|N^*| = 0$. It must hold that $C \cap N'' \neq \emptyset$. In this case, a^* loses to a in a pairwise election, since all the agents in N'' prefer a to a^*.

Case 2: $0 < |N^*| \leq \frac{n}{3}$. Since there is no exact cover, the corresponding sets S_i cannot cover U. Thus there exists $b \in U$ that is ranked higher than a^* by all agents in N^*. In order for a^* to beat b in a pairwise election, C must include at least $|N^*| + 1$ agents from N''. However, this means that a beats a^* in a pairwise election (since a is ranked lower than a^* by $|N^*|$ agents, and higher than a^* by at least $|N^*| + 1$ agents). It follows that a^* is not a Condorcet winner for C.

Case 3: $|N^*| > \frac{n}{3}$. Let us award each alternative $b \in A \setminus \{a^*\}$ a point for each agent that ranks it above a^*, and subtract a point for each agent that ranks it below a^*. a^* is a Condorcet winner iff the score of every other alternative, counted this way, is negative. This implies that a^* is a

Condorcet winner only if for every subset $B \subseteq A$ of alternatives, the total score of the alternatives in B is at most $-|B|$.

We shall calculate the total score of the alternatives in U from the agents in N^*. Every agent in N^* prefers a^* to 3 alternatives in U and prefers $n-3$ alternatives in U to a^*. Thus, every agent in N^* contributes $(n-3)-3 = n-6$ points to the total score of U. Summing over all the agents in N^*, we have that the total score of U from N^* is $|N^*|(n-6)$. By $|N^*| > \frac{n}{3}$, we have that

$$|N^*|(n-6) \geq \left((\frac{n}{3}-1)+2\right)(n-6) = (\frac{n}{3}-1)n - 6 \quad .$$

Recall that every agent in N'' prefers a^* to all alternatives in U. However, since $|N''| = \frac{n}{3}-1$, agents from N'' can only subtract $(\frac{n}{3}-1)n$ from the total score of U. We conclude that the total score of U is at least -6. Since we can assume that $|U| = n > 6$,[6] a^* cannot beat all the alternatives in U in pairwise elections. This concludes the proof. □

F Proof of Corollary 4.3

Let $\epsilon > 0$ be a constant. We perform the same reduction as before, with the following differences. Let A' be the set of alternatives constructed in the reduction of Lemma 4.2, and $m' = |A'|$; we add a set B of $(m')^{1/\epsilon}$ additional alternatives, i.e., $A = A' \cup B$, $m = |A| = m' + (m')^{1/\epsilon}$. The set of agents is $N' \cup N'' \cup N^*$, the preferences of N' and N'' restricted to A' are as before, and all these agents rank B at the bottom. All the agents in N^* rank a^* last; for each $b \in A' \setminus \{a^*\}$, there is $i \in N^*$ that ranks b first and B just above a^*, i.e.,

$$b \ R^i \ A' \setminus \{a^*, b\} \ R^i \ B \ R^i \ a^* \quad .$$

For each $c \in B$, there is $i \in N^*$ that ranks c first and the rest of B just above a^*, namely

$$c \ R^i \ A' \setminus \{a^*\} \ R^i \ B \setminus \{c\} \ R^i \ a^* \quad .$$

Notice that the Young score of the alternatives in $A' \setminus \{a^*\}$ is at least one. The Young score of any alternative $c \in B$ is exactly one, since exactly one voter (in N^*) does not rank $A' \setminus \{a^*\}$ above c. Now, if there is an exact 3-cover, then the Young score of a^* is at least $2n/3 - 1$ (according to the proof of Lemma 4.2), so a^* is ranked above all the alternatives in B, that is, in the top $m'+1 = \mathcal{O}(m^\epsilon)$ places. On the other hand, if there is no exact 3-cover, then the Young score of a^* is zero by the same arguments as in Lemma 4.2, since the agents in N^* all rank a^* last. Hence a^* is placed last in any Young ranking. □

[6] X3C is obviously tractable for a constant n, as one can examine all the families $\mathcal{S}' \subseteq \mathcal{S}$ of constant size in polynomial time.