



Social Ranking under Incomplete Knowledge: Elicitation of the Lex-cel Necessary Winners

Sébastien Konieczny, Stefano Moretti, Ariane Ravier, Paolo Viappiani

► To cite this version:

Sébastien Konieczny, Stefano Moretti, Ariane Ravier, Paolo Viappiani. Social Ranking under Incomplete Knowledge: Elicitation of the Lex-cel Necessary Winners. 16th International Conference on Scalable Uncertainty Management. SUM 2024., Nov 2024, Palermo, Italy. hal-04750114

HAL Id: hal-04750114

<https://hal.science/hal-04750114v1>

Submitted on 23 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Social Ranking under Incomplete Knowledge: Elicitation of the Lex-cel Necessary Winners

Ariane Ravier¹, Sébastien Konieczny², Stefano Moretti¹, and Paolo Viappiani¹

¹ LAMSADE, CNRS, Université Paris-Dauphine, Université PSL, 75016, Paris, France.
{ariane.ravier, stefano.moretti, paolo.viappiani}@lamsade.dauphine.fr

² CRIL - CNRS - Université d'Artois - France. konieczny@cril.fr

Abstract. In this paper, we introduce two elicitation approaches for the determination of *lex-cel necessary winners* (i.e., the elements that are ranked highest according to *lex-cel*) in a social ranking problem, when the knowledge about preferences over the coalitions is incomplete, and the initially accessible input is a subset of an existing total preorder. The first approach is preorder-driven, and elicits enough of the underlying total preorder to determine the *lex-cel* necessary winners. The second approach is element-driven and guides comparisons based on strategically-located coalitions. Finally, we present experimental results and discuss the performance of each approach depending on various parameters and scenarios.

1 Introduction

The concept of social ranking has been recently introduced in the literature to address the problem of ranking individual elements based on their contribution in establishing the position of groups or coalitions within a society [1]. To generate a social ranking, a ranking over sets of objects (such as groups, team's squads, winning and losing coalitions within a voting body, etc.) is taken as input. This generates a ranking over individual elements (such as group members, team's players, voters, etc.) that reflects their overall influence within the ranking over sets.

In [2], a social ranking solution called *lexicographic excellence* (*lex-cel*) is identified as the only one satisfying a set of appealing properties. According to this rule, individual elements are ranked based on their occurrences in the highest positions of the ranking over coalitions (note that this ranking may contain equivalences) as follows. First, one counts the number of occurrences of each element (i.e., the number of sets containing the element) in every equivalence class of the ranking over sets. When comparing two elements, one must focus to the number of occurrences in the best equivalence class of sets, the greater number of occurrences being the better element. In case the number of occurrences in the best equivalence class is the same, attention shifts to the number of occurrences in the second-best equivalence class of sets, with the higher number determining the better element. If parity persists, the third-best equivalence class is considered, and so on, until all the equivalence classes are examined (of course, in case of a tie in each equivalence class of sets, the two elements are declared indifferent). For instance, in a situation with only two individual elements 1 and 2 and a ranking over

sets such that the set $\{1, 2\}$ and the singleton $\{2\}$ are in the best equivalence class and both are strictly better than the singleton $\{1\}$, the lex-cel places 2 above 1, as 2 occurs twice in the best equivalence class, while 1 occurs just once.

Generalizations of lex-cel have been proposed in [3,4], considering the size of coalitions in addition to the occurrence of elements in the sets, and in [5], under variable domains of coalitions. Other social ranking solutions, based on different criteria, have been proposed, for instance, in [6,7,8,9]. As noticed in [10], however, generating a complete ranking over individual elements may not be necessary. In such cases, the objective may be to select only the most important ones. As a motivating example, the authors of [10] consider the problem of identifying the most influential scientists within an academic association based on the number and quality of their publications and taking into account the contributions of scientists to multi-authored publications. They first propose to establish a ranking over groups of scientists according to a predefined bibliometric criterion, then they adopt a property-driven approach to argue that most influential scientists across different groups are the elements with the highest position according to lex-cel computed on the ranking of groups.

In this paper we apply the lex-cel method to identify the most significant elements (also called “winners”) in the presence of partial information. Indeed, we observe that the process of acquiring the necessary information to compute the lex-cel is often highly time-consuming, as one can argue that collecting and processing data to compare any possible pair of groups of a finite set becomes quickly unaffordable in practice. So, the main contribution of this paper is to provide elicitation procedures aiming at identifying winners (according to lex-cel) using a subset of the ranking relation over coalitions. In other words, our goal is to find a procedure to determine the winner(s) of the lex-cel by asking a minimum number of questions to an expert (*e.g.*, the director of the academic association) about pairwise comparisons among groups.

Our approach bears similarity with works in computational social choice that deal with determining possible and necessary winners [11,12,13] when only a partial preference profile, composed of partial orders, is given in input. Adaptive elicitation approaches have been proposed that refine the partial profile in an iterative manner until a necessary winner is identified or regret drops below a threshold [14,15,16].

We start by presenting some preliminary notions and concepts, then introduce the problem at hand. We describe, first, a preorder-driven elicitation method to determine the lex-cel necessary winners by reconstituting the underlying total preorder, and, second, an element-driven method, based on a strategic choice of coalitions to compare. Finally, we present experimental results and discuss the performance of each method depending on various scenarios.

2 Preliminaries

Given a finite set $\mathcal{E} = \{1, \dots, n\}$, a **binary relation** $\succsim_{\mathcal{E}}$ over $\mathcal{E}$ is a subset of the Cartesian product $\mathcal{E} \times \mathcal{E}$. In the following, given two elements $i, j \in \mathcal{E}$, the fact that $(i, j) \in \succsim_{\mathcal{E}}$ is also denoted by $i \succsim_{\mathcal{E}} j$. A binary relation $\succsim_{\mathcal{E}}$ over $\mathcal{E}$ that is reflexive and transitive is called **preorder** over $\mathcal{E}$, and the pair $(\mathcal{E}, \succsim_{\mathcal{E}})$ is a **preordered set**. The set of all total preorders on $\mathcal{E}$ is denoted by $\mathcal{R}(\mathcal{E})$. In the remaining of the paper a preorder may also be referred to as a **preference relation**. A total preorder that is

also antisymmetric is called a **linear order**. An **extension** of a preorder $\succsim_{\mathcal{E}}$ is a total preorder $\succsim_{\mathcal{E}}^* \in \mathcal{R}(\mathcal{E})$ such that $\succsim_{\mathcal{E}} \subseteq \succsim_{\mathcal{E}}^*$. If an extension is a linear order it is also called **linear extension**.

Given a total preorder $\succsim_{\mathcal{E}}$ over $\mathcal{E}$, we denote by $\sim_{\mathcal{E}}$ the symmetric part of $\succsim_{\mathcal{E}}$ (i.e. $i \sim_{\mathcal{E}} j$ if $i \succsim_{\mathcal{E}} j$ and $j \succsim_{\mathcal{E}} i$) and by $\succ_{\mathcal{E}}$ its asymmetric part (i.e. $i \succ_{\mathcal{E}} j$ and not $j \succ_{\mathcal{E}} i$). A **maximal element** $i \in \mathcal{E}$ of a preordered set $(\mathcal{E}, \succsim_{\mathcal{E}})$ is such that there exists no $j \in \mathcal{E}$ verifying $j \succ_{\mathcal{E}} i$; we denote by $\mathcal{M}^{\succsim_{\mathcal{E}}}$ the set of all maximal elements of $(\mathcal{E}, \succsim_{\mathcal{E}})$. Given a preordered set $(\mathcal{E}, \succsim_{\mathcal{E}})$, the preordered set restricted to a subset $C \subseteq \mathcal{E}$ is defined as the pair $(C, \succsim_C^C)$, where $\succsim_C^C \subseteq C \times C$ is such that $i \succsim_C^C j \Leftrightarrow i \succsim_{\mathcal{E}} j$ for all $i, j \in C$; if $\succsim_C^C$ is total, the preordered set $(C, \succsim_C^C)$ is called a **chain**. A **maximal chain** is a chain $(C, \succsim_C^C)$ such that there exists no $D \subseteq \mathcal{E}$ with $C \subseteq D$ and such that $(D, \succsim_D^D)$ is a chain.

Let X be a finite set. We denote by $\mathcal{P}(X)$ the set of non-empty subsets of X , i.e. $\mathcal{P}(X) = \{A \subseteq X : A \neq \emptyset\}$. The elements of $\mathcal{P}(X)$ are called **coalitions**. A total preorder $\succsim_{\mathcal{P}(X)} \in \mathcal{R}(\mathcal{P}(X))$ is called a **power relation** (when the set $\mathcal{P}(X)$ is clear from the context, a power relation on $\mathcal{P}(X)$ is simply denoted by $\succsim$). Consider a power relation $\succsim \in \mathcal{R}(\mathcal{P}(X))$ and two coalitions $S, T \in \mathcal{P}(X)$ such that $S \succsim T$ (interpreted as “ S is at least as powerful as T w.r.t. the power relation $\succsim$ ”). The **quotient order** of $\succsim$ is denoted by $\Sigma_1 \succ \Sigma_2 \succ \dots \succ \Sigma_l$ where each **equivalence class** $\Sigma_k, k \in \{1, \dots, l\}$ is generated by the symmetric part of $\succsim$. This means that all coalitions in Σ_1 are “equally powerful” and are “strictly more powerful” than the sets in Σ_2 and so on. We denote by $x_k = |\{S \in \Sigma_k : x \in S\}|$ the number of occurrences of element x in Σ_k , for any $k = 1, \dots, l$. Let us denote by $\theta^{\succsim}(x)$, for any $x \in X$, the l -dimensional vector $\theta^{\succsim}(x) = (x_1, \dots, x_l)$ associated with $\succsim$. This vector is called the **occurrence vector** of x , as it indicates the number of occurrences of x in each equivalence class of $\succsim$.


A **social ranking** $R : \mathcal{R}(\mathcal{P}(X)) \rightarrow \mathcal{R}(X)$ is a mapping that associates a power relation $\succsim \in \mathcal{R}(\mathcal{P}(X))$ over coalitions with a total preorder $R^{\succsim} \in \mathcal{R}(X)$ over single elements. Let $\mathbf{i}, \mathbf{j}$ be two vectors of the same size, the lexicographic total preorder $\geq_L$ is such that $\mathbf{i} \geq_L \mathbf{j}$ if either $\mathbf{i} = \mathbf{j}$ or there exists t such that $i_t > j_t$ and $i_r = j_r$ for all $r \in \{1, \dots, t-1\}$. The *lexicographic excellence* (lex-cel) [2] is a social ranking $R_{le}^{\succsim} : \mathcal{R}(\mathcal{P}(X)) \rightarrow \mathcal{R}(X)$ such that for all $\succsim \in \mathcal{R}(\mathcal{P}(X))$ and all $x, y \in X$: $x R_{le}^{\succsim} y \Leftrightarrow \theta^{\succsim}(x) \geq_L \theta^{\succsim}(y)$.

The symmetric and the asymmetric part of $R_{le}^{\succsim}$ are denoted by $I_{le}^{\succsim}$ and $P_{le}^{\succsim}$, respectively. An element $x \in X$ such that $x R_{le}^{\succsim} y$ for all $y \in X$ is a *winner* in $R_{le}^{\succsim}$. There may be several winners; an axiomatic approach to the problem of selecting the winners provided by the lex-cel has been studied in [10].

Remark 1. It is easy to show that if $\succsim$ is a linear order over $\mathcal{P}(X)$, then $\theta^{\succsim}(x)$ is a 0, 1-vector and $\theta^{\succsim}(x) \neq \theta^{\succsim}(y)$ for all $x, y \in X$ with $x \neq y$; so, $R_{le}^{\succsim}$ is a linear order over X and there is a unique winner in $R_{le}^{\succsim}$.

Example 1. Let $X = \{1, 2, 3, 4\}$, and $\succsim_{\mathcal{P}(X)}$ (or, simply, $\succsim$) be a preorder represented by the preference graph of Figure 1 (left-side) and its maximal chains (right-side).

The following power relation $\succsim^*$ is an extension of $\succsim$: $1234 \succ^* 134 \succ^* 123 \succ^* 24 \succ^* 3 \succ^* 34 \succ^* 2 \succ^* 1 \succ^* 13 \succ^* 23 \succ^* 124 \succ^* 12 \succ^* 234 \sim^* 4 \succ^* 14$, from which we determine the following occurrence vectors: $\theta^{\succsim^*}(1) = (1, 1, 1, 0, 0, 0, 0)$,

Fig. 1: Preference graph and maximal chains of $\succ$

$1, 1, 0, 1, 1, 0, 1, \theta^{\succ^*}(2) = (1, 0, 1, 1, 0, 0, 1, 0, 0, 1, 1, 1, 1, 0), \theta^{\succ^*}(3) = (1, 1, 1, 0, 1, 1, 0, 0, 1, 1, 0, 0, 1, 0), \theta^{\succ^*}(4) = (1, 1, 0, 1, 0, 1, 0, 0, 0, 0, 1, 0, 2, 1)$.

According to lex-cel, we have $3 \ P_{le}^{\succ^*} 1 \ P_{le}^{\succ^*} 4 \ P_{le}^{\succ^*} 2$, and the winner in $R_{le}^{\succ^*}$ is 3.

3 Problem description

We consider social rankings situations where only partial information about the coalitions is available; instead of a power relation in $\mathcal{R}(\mathcal{P}(X))$, we are given a preorder $\succsim_{\mathcal{P}(X)}$. In a way similar to what is done in social choice [11], we define possible and necessary winners, by considering winners in extensions of $\succsim_{\mathcal{P}(X)}$.

Definition 1. An element $x \in X$ is a **lex-cel possible winner** on $\succsim_{\mathcal{P}(X)}$ if there exists an extension $\succsim_{\mathcal{P}(X)}^*$ of $\succsim_{\mathcal{P}(X)}$ such that x is a winner in $R_{le}^{\succsim_{\mathcal{P}(X)}^*}$.

Definition 2. An element $x \in X$ is a **lex-cel necessary winner** on $\succsim_{\mathcal{P}(X)}$ if x is a winner in $R_{le}^{\succsim_{\mathcal{P}(X)}^*}$ for every extension $\succsim_{\mathcal{P}(X)}^*$ of $\succsim_{\mathcal{P}(X)}$.

Example 2. Let $\succsim_{\mathcal{P}(X)}$ be the preorder introduced in Example 1. Possible extensions of $\succsim_{\mathcal{P}(X)}$ (other than $\succeq^*$) are $\succeq^a$ and $\succeq^b$ such that $134 \succ^a 34 \succ^a 2 \succ^a 24 \succ^a 1234 \succ^a 3 \succ^a 13 \succ^a 123 \succ^a 23 \succ^a 124 \succ^a 1 \succ^a 12 \succ^a 234 \succ^a 4 \succ^a 14$ and $123 \succ^b 134 \succ^b 1 \succ^b 34 \succ^b 12 \succ^b 1234 \succ^b 2 \succ^b 3 \succ^b 24 \succ^b 23 \succ^b 13 \succ^b 4 \succ^b 14 \succ^b 124 \succ^b 234$.

One can check that 4 is winning in $R_{le}^{\succeq^a}$, 1 is winning in $R_{le}^{\succeq^b}$, but there is no extension of $\succeq$ such that 2 is winning. Actually, only 1, 3 and 4 are lex-cel possible winners of $\succeq$.

Remark 2. Since an extension of $\succsim_{\mathcal{P}(X)}$ must exist, the set of lex-cel possible winners on a preorder $\succsim_{\mathcal{P}(X)}$ is non-empty.

The problem of generating linear extensions of a partially ordered set (poset) has been widely studied in the literature [17,18]. However, as it will be better clarified later, we are not interested in generating the set of lex-cel possible winners. Our goal is to find a “parsimonious” procedure to elicit the winners in $R_{le}^{\succsim_{\mathcal{P}(X)}^*}$, where $\succsim_{\mathcal{P}(X)}^*$ is the (unique) true extension of $\succsim_{\mathcal{P}(X)}$, and the concept of lex-cel possible winner is only used to guide our elicitation procedures.

Proposition 1 Consider³ the quotient order $\Sigma_1 \succ_{\mathcal{P}(X)} \dots \succ_{\mathcal{P}(X)} \Sigma_l$. Each winner in $R_{le}^{\succ_{\mathcal{P}(X)}}$ belongs to Σ_1 .

Now, let $\succ_{\mathcal{P}(X)}$ be a preorder and let $\succ_{\mathcal{P}(X)}^* \in \mathcal{R}(\mathcal{P}(X))$ be an extension of $\succ_{\mathcal{P}(X)}$. Notice that the maximal coalitions in the (totally) preordered set $(\mathcal{P}(X), \succ_{\mathcal{P}(X)}^*)$ are also maximal in the preordered set $(\mathcal{P}(X), \succ_{\mathcal{P}(X)})$, i.e. $\mathcal{M}^{\succ_{\mathcal{P}(X)}} \subseteq \mathcal{M}^{\succ_{\mathcal{P}(X)}^*}$.

Let $\succ_{\mathcal{P}(X)}^* \in \mathcal{R}(\mathcal{P}(X))$ be a power relation, and suppose that it is only partially known, i.e. only a subset $\succ_{\mathcal{P}(X)} \subseteq \succ_{\mathcal{P}(X)}^*$ is initially revealed. Our problem is to gather, through an elicitation process, sufficient information over $\succ_{\mathcal{P}(X)}^* \setminus \succ_{\mathcal{P}(X)}$ to determine the set of winners in $R_{le}^{\succ_{\mathcal{P}(X)}^*}$. As we are mainly concerned with the cognitive effort provided by the user, we base our evaluation of an elicitation method's performance on the number of queries it submits. Starting from the (input) preorder $\succ_{\mathcal{P}(X)} \subseteq \succ_{\mathcal{P}(X)}^*$, we investigate different elicitation procedures aimed at submitting queries to a user (expert) about unknown relations between pairs of coalitions. We assume that each query asked to the user over an arbitrary pair of coalitions (let's say, coalitions C_1 and C_2 in $\mathcal{P}(X)$) receives in answer a set of pairs expressing the true value of the relation between the two coalitions, i.e., $C_1 \succ C_2$ or $C_2 \succ C_1$ (or both, if $C_1 \sim C_2$). A tuple (C_i, C_j) indicates that $C_i \succ_{\mathcal{P}(X)}^* C_j$.

Our objective is to minimize the number of queries submitted to the user to determine a non-empty set of lex-cel necessary winners on the set of revealed relations, that is finding a set $Q \subseteq (\succ_{\mathcal{P}(X)}^* \setminus \succ_{\mathcal{P}(X)})$ such that:

$$Q \in \arg \min_{R \subseteq (\succ_{\mathcal{P}(X)}^* \setminus \succ_{\mathcal{P}(X)})} \{ |R| : \mathcal{NW}^{\succ_{\mathcal{P}(X)} \cup R} \neq \emptyset \} \quad (1)$$

where $\mathcal{NW}^{\succ_{\mathcal{P}(X)} \cup R}$ is the set of lex-cel necessary winners on $\succ_{\mathcal{P}(X)} \cup R$. Problem (1) is called the NECESSARY WINNER IN SOCIAL RANKING (NWSR) problem. It is an extension of the social ranking problem, in a context where the user's preferences are only partially known. This can be because the user has not yet proceeded to the comparison of all possible coalitions, as this can prove cognitively taxing. However, it is assumed that there exists an underlying total preorder over $\mathcal{P}(X)$, which is obviously compatible with the preorder given as input. As the total preorder represents the user's preferences, we assume it may be discovered through queries. From the above definitions and propositions, we can now present the different steps of our two elicitation approaches for NWSR.

4 Elicitation by reconstitution: RECO

A first preorder-based approach aims at reconstituting sufficiently many of the most preferred coalitions in the underlying total preorder so as to determine the lex-cel necessary winners with certainty. The process is iterative: at each step, we present the user with two coalitions over which they must express their preference.

Recall that the restriction of a preorder $\succ_{\mathcal{P}(X)}$ over $\mathcal{P}(X)$ to a collection $\mathcal{S} \subseteq \mathcal{P}(X)$ is the preorder $\succ_{\mathcal{P}(X)}^{\mathcal{S}}$ over $\mathcal{S}$ such that $C_1 \succ_{\mathcal{P}(X)}^{\mathcal{S}} C_2 \Leftrightarrow C_1 \succ_{\mathcal{P}(X)} C_2$ for all

³ Proofs for all propositions are available in the Appendix, available at [this url](#)

$C_1, C_2 \in \mathcal{S}$. With a slight abuse of notation, given a collection of coalitions $\mathcal{C} \subseteq \mathcal{P}(X)$, we denote by $x_{\mathcal{C}} = |\{C \in \mathcal{C} : x \in C\}|$ the number of occurrences of element x in the collection $\mathcal{C}$.

Algorithm 1: Elicitation by reconstitution

Input: A preorder $\succsim_{\mathcal{P}(X)}$;
Output: A preorder $\succsim \supseteq \succsim_{\mathcal{P}(X)}$ and a set W of lex-cel necessary winners on $\succsim$;
 $\mathcal{C} \leftarrow \mathcal{P}(X)$; $W \leftarrow X$; $\succsim \leftarrow \succsim_{\mathcal{P}(X)}^c$;
while $|W| > 1$ **and** $\mathcal{C} \neq \emptyset$ **do**
 if $S \sim T$ **for all** $S, T \in \mathcal{M}^{\succsim}$ **then**
 $W \leftarrow \{x \in W : x_{\mathcal{M}^{\succsim}} \geq y_{\mathcal{M}^{\succsim}} \text{ for each } y \in W\}$;
 $\mathcal{C} \leftarrow \mathcal{C} \setminus \mathcal{M}^{\succsim}$;
 $\succsim \leftarrow \succsim_{\mathcal{P}(X)}^c$;
 else
 $C_1, C_2 \leftarrow \text{SELECT}(\succsim_{\mathcal{P}(X)}^c)$;
 $\text{Newcomp} \leftarrow \text{QUERY}(C_1, C_2)$;
 $\succsim \leftarrow \text{Transitive_Closure}(\succsim \cup \text{Newcomp})$;
return $\succsim$;
return W .

Algorithm 1 takes as input a preorder $\succsim_{\mathcal{P}(X)}$ and returns an updated preorder $\succsim \supseteq \succsim_{\mathcal{P}(X)}$, as well as a set of winners W . The procedure starts by considering all elements in X as possible winners. As long as there are still several possible winners and not all the coalitions have been considered and ranked (this is to address the scenario in which there may be several necessary winners, due to equivalences in the total preorder), two coalitions are selected using $\text{SELECT}(\succsim)$, a subroutine which, given a preorder $\succsim$ over a set $\mathcal{S}$, selects two coalitions $C_1, C_2 \in \mathcal{S}$ such that $(C_1, C_2) \notin \succsim$ and $(C_2, C_1) \notin \succsim$. We study different approaches to coalition selection, *i.e.* different definitions of $\text{SELECT}(\succsim)$. Using the function $\text{QUERY}(C_1, C_2)$, the user is then asked to indicate the set of preference relations over C_1 and C_2 as a set of tuples. The function $\text{Transitive_Closure}$ returns the smallest (with respect to the number of related pairs) transitive superset of the updated $\succsim$. If all maximal elements from the maximal chains (*i.e.* all coalitions in $\mathcal{M}^{\succsim}$) have been evaluated as equivalent, they must belong to the same equivalence class Σ , which can no longer be extended: we therefore update our set of possible winners to contain only those who are most represented in Σ (in accordance with the lexcel), and remove from consideration the elements in $\mathcal{M}^{\succsim}$, as they are considered to have been added to the reconstructed order and their impact on the set of possible winners has been taken into account.

Proposition 2 *Algorithm 1 returns a set W of lex-cel necessary winners for the power relation $\succsim$.*

The performance of this approach depends heavily on the manner in which the two coalitions C_1 and C_2 are selected at each iteration. We therefore study several different manners of selecting coalitions, which we will refer to as *variants*.⁴

⁴ The pseudo-code for each variant is available in the Appendix.

The *Complete randomness variant* (C-RAND) is a first naïve approach consists in simply selecting two uncomparing coalitions randomly within $\mathcal{P}(X)$. However, it quickly appears that this approach is far from efficient, as we aim at reconstituting the preorder over the top coalitions, and the complete randomness variant can compare coalitions regardless of their likelihood to be among the most preferred. As the top coalition in $\succsim_{\mathcal{P}(X)}^*$ is necessarily a maximal element of $\succsim_{\mathcal{P}(X)}$, restricting the pool of candidate coalitions to $\mathcal{M}^{\succsim_{\mathcal{P}(X)}}$ significantly improves the relevance of the queries.

The *Restricted randomness variant* (R-RAND) consists in selecting at random two uncomparing coalitions within $\mathcal{M}^{\succsim_{\mathcal{P}(X)}}$.

We wish to reduce the impact of randomness on the selection process, and therefore introduce new strategies, the first of which shares the lexicographic approach of lex-cel. The *Restricted maximal length variant* (R-MAX_LEN) makes the hypothesis that longer maximal chains are more likely to be informative. It uses a vector $\Gamma = (\gamma_1, \dots, \gamma_s)$, where s is the number of maximal chains in $\succsim_{\mathcal{P}(X)}$ and γ_i is the i -th longest maximal chain in $\succsim_{\mathcal{P}(X)}$. We define $M^\Gamma = (m_1, \dots, m_s)$ an s -sized vector such that m_i is one of the maximal elements of γ_i .

Remark 3. Note that randomness may still play a part in this variant. Indeed, let $\succsim_{\mathcal{P}(X)}$ be the preorder presented in Example 1, the associated set of maximal elements is then such that $\mathcal{M}^{\succsim_{\mathcal{P}(X)}} = \{1234, 134, 34, 123\}$. As all maximal chains are of identical length, the tie-breaker is the order in which they have been given as input. Therefore, in this particular instance, the first comparison will be between 1234 and 134.

Finally, we try to focus on a maximization of a query's informativity by exploiting properties specific to lex-cel; this gives rise to the *Restricted minimal nonempty intersection* (R-MIN_INT) variant. We can define a category of queries we call *informative* with regards to a set of lex-cel possible winners. Given a preorder $\succsim_{\mathcal{P}(X)}$ over which we can determine a set of lex-cel possible winners W , a query over a pair of coalitions $C_1, C_2 \in \mathcal{P}(X)$ is said to be informative if, provided C_1 and C_2 are the top coalitions among $\mathcal{C}$, the recovered preference over them allows for the reduction of the set W . A query is maximally informative when the preference it uncovers can reduce W to a single element: the lex-cel necessary winner has then been uncovered.

For an element x to be removed from W in favour of another element $y \in W$, simply by comparing C_1 and C_2 as top coalitions among $\mathcal{C}$, implies that $\exists C \in \{C_1, C_2\}$ such that $x \notin C$ yet $y \in C$. This means the intersection of C_1 and C_2 must be smaller than the largest of the two coalitions. The result of the query must then lead to the placement of the coalition containing y before any containing x . We therefore seek to minimize the intersection of pairs of candidate coalitions, in order to minimize the number of elements that may remain in W after submission of the query. Note that if this intersection set is empty, then the query is considered non-informative, for when none of the candidates are present in the k -th equivalence class, the k -th component of both their occurrence vectors will be 0, and will not allow the lexicographic relation to discriminate between the candidates.

Example 3. Let $\succsim_{\mathcal{P}(X)}$ be the preorder presented in Example 1. The set $\mathcal{M}^{\succsim_{\mathcal{P}(X)}}$ is therefore such that $\mathcal{M}^{\succsim_{\mathcal{P}(X)}} = \{1234, 134, 34, 123\}$. We have that $|1234 \cap 134| = |1234 \cap 123| = 3$, $|1234 \cap 34| = |134 \cap 34| = |134 \cap 123| = 2$ and $|123 \cap 34| = 1$. The

minimal non-empty intersection is of size 1, and the relative comparison will therefore be between 123 and 34.

The existence of informative queries highlighted in the restricted lexicographic variant leads us to observe that eliciting by reconstituting the total preorder may not always be an efficient approach. Indeed, there exist coalitions which will bring no information to lex-cel, and any query about their placement in the total preorder is therefore superfluous. A second approach to elicitation suggests to focus directly on the elements themselves, and only submit to the user targeted queries over key informative coalitions which will help determining the preference over elements without needing to discover the entire preorder.

5 Elicitation through prospection: PROSP

A second element-based approach aims at directly discerning the preference over the individual elements themselves. To do so, it considers the likelihood of certain elements within the population to be a lex-cel necessary winner, then determines key coalitions, the preference between which will be decisive during the application of lex-cel to assess the dominance of an element over another.⁵

Given a total preorder $\succsim_{\mathcal{P}(X)}^*$, and $x, y \in X$, we denote by $C_{x,-y}$ the most highly ranked coalition in $\succsim_{\mathcal{P}(X)}^*$ verifying that $x \in C_{x,-y}$ and $y \notin C_{x,-y}$.

Proposition 3 *Given a total preorder $\succsim_{\mathcal{P}(X)}^*$, let $x, y \in X$ be present in a best coalition $C^* \in \succsim_{\mathcal{P}(X)}^*$. It holds that $x R_{le}^{\succsim_{\mathcal{P}(X)}^*} y$ if $C_{x,-y} \succ_{\mathcal{P}(X)}^* C_{y,-x}$.*

Remark 4. Given a total preorder $\succsim_{\mathcal{P}(X)}^*$, let $x, y \in X$ be present in a best coalition $C^* \in \succsim_{\mathcal{P}(X)}^*$. If $C_{x,-y} \sim_{\mathcal{P}(X)}^* C_{y,-x}$, then neither $C_{x,-y}$ nor $C_{y,-x}$ will be decisive in the comparison of x and y . As such, $C_{x,-y}$ and $C_{y,-x}$ may be removed from consideration, and their values updated.

For any $x \in X$, let $\mathcal{P}_x = \{C \in \mathcal{P}(X) : x \in C\}$. For any pair $x, y \in X$, $x \neq y$, let $\mathcal{P}_{x,y} = \{\mathcal{P}_x \cup \mathcal{P}_y\} \setminus \{\mathcal{P}_x \cap \mathcal{P}_y\}$ be the set of coalitions in $\mathcal{P}(X)$ containing either x or y , but not both elements. $\mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}^*}$ denotes the set of maximal elements in $\mathcal{P}_{x,y}^{\succsim_{\mathcal{P}(X)}^*}$, that can then be partitioned into two disjoint subsets $\mathcal{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}^*} = \{C \in \mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}^*} : x \in C\}$ and $\mathcal{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}^*} = \{C \in \mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}^*} : y \in C\}$.

At each step, the procedure determines the set W of lex-cel possible winners in the current preorder, then selects the two deemed most likely to be more preferred amongst them, and determines coalitions which would play a key role in establishing a preference between the two elements. The user is only presented queries relative to these key coalitions.

From Proposition 1, we know that, in order to be a winner, an element $x \in X$ must be present in a top coalition of the total preference relation. While the exact set of the top coalitions is unknown, we proceed using $\mathcal{M}_{\mathcal{P}(X)}^{\succsim_{\mathcal{P}(X)}^*}$, as we know all top coalitions must belong to that set.

⁵ The pseudo-code for PROSP is presented in the Appendix.

Determining the most likely lex-cel possible winners: the global lex-cel. In order to minimize the number of queries to the user, we try and determine the possible winners most likely to perform better. As a lex-cel necessary winner is one of the most present elements in the top coalitions of the total preorder, and since the higher a coalition is placed within a maximal chain, the most likely it is to be among the top coalitions, we study the presence of each lex-cel possible winner in each position. To do so, we present a generalisation of lex-cel, applicable to any set of maximal chains.

Let $\mathcal{A}^{\succsim_{\mathcal{P}(X)}}$ be the set of all maximal chains in $\succsim_{\mathcal{P}(X)}$. We denote by l the size of a longest maximal chain in $\mathcal{A}^{\succsim_{\mathcal{P}(X)}}$, and by A_k the k -th ranked equivalence class in a maximal chain $A \in \mathcal{A}^{\succsim_{\mathcal{P}(X)}}$, $\forall k \in \{1, \dots, |A|\}$. For any maximal chain $A \in \mathcal{A}^{\succsim_{\mathcal{P}(X)}}$ and for any element $i \in X$, we denote by $i_k^A = |\{C \in A_k \mid i \in C\}|$ the number of occurrences of i in the k -th equivalence class of the maximal chain A . We then denote by $\vartheta^{\succsim_{\mathcal{P}(X)}}(i)$ the l -dimensional vector

$$\vartheta^{\succsim_{\mathcal{P}(X)}}(i) = \left(\sum_{A \in \mathcal{A}^{\succsim_{\mathcal{P}(X)}}} i_1^A, \dots, \sum_{A \in \mathcal{A}^{\succsim_{\mathcal{P}(X)}}} i_l^A \right).$$

Definition 3. The global lexicographic-excellence (global lex-cel) is the binary relation $R_{gle}^{\succsim_{\mathcal{P}(X)}}$ such that, for any preorder $\succsim_{\mathcal{P}(X)}$, and for any $i, j \in X$:

$$i R_{gle}^{\succsim_{\mathcal{P}(X)}} j \iff \vartheta^{\succsim_{\mathcal{P}(X)}}(i) \geq_L \vartheta^{\succsim_{\mathcal{P}(X)}}(j)$$

Example 4. Let $\succsim_{\mathcal{P}(X)}$ be the preorder presented in Example 1. All maximal chains in $\succsim_{\mathcal{P}(X)}$ are of length at most $l = 5$, therefore we determine the l -sized vectors associated with every element in the population: $\vartheta^{\succsim_{\mathcal{P}(X)}}(1) = (4, 1, 5, 3, 3)$, $\vartheta^{\succsim_{\mathcal{P}(X)}}(2) = (2, 4, 2, 3, 3)$, $\vartheta^{\succsim_{\mathcal{P}(X)}}(3) = (6, 1, 5, 0, 3)$, and $\vartheta^{\succsim_{\mathcal{P}(X)}}(4) = (5, 2, 0, 6, 6)$, from which we can determine the following total preorder over X :

$$3 R_{gle}^{\succsim_{\mathcal{P}(X)}} 4 R_{gle}^{\succsim_{\mathcal{P}(X)}} 1 R_{gle}^{\succsim_{\mathcal{P}(X)}} 2.$$

Using the global lex-cel, we determine a preliminary total preorder over X , from which we select the two best elements as the pair driving the current iteration of the procedure.

Querying the user. We submit queries to determine the top coalition $C \in \mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}}$. Depending on whether C is in $\mathcal{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ or $\mathcal{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$, we are then able to determine the preference relation between x and y .

Proposition 4 Given a total preorder $\succsim_{\mathcal{P}(X)}^*$. Let $\succsim_{\mathcal{P}(X)}$ be a preorder such that $\succsim_{\mathcal{P}(X)} \subseteq \succsim_{\mathcal{P}(X)}^*$. Let $x, y \in X$ be the two best elements in X according to the global lex-cel. It holds that $x R_{le}^{\succsim_{\mathcal{P}(X)}} y$ iff there exists a coalition in $\mathcal{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ preferred to every coalition in $\mathcal{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$ or if $\mathcal{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}} = \emptyset$ and $\mathcal{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}} \neq \emptyset$.

We consider that a coalition in $\mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}}$ is more likely to be the most preferred in the set when its worst position in any maximal chain is still a low position. For this reason, instead of the sets $\mathcal{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ and $\mathcal{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$, the procedure uses the vector of coalitions ordered in ascending worst position in a maximal chain of $\succsim_{\mathcal{P}(X)}$, denoted by $\hat{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ and $\hat{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$ respectively.

We submit to the user a query over the first coalition from each vector. Without loss of generality, when it is found that a coalition $C_2 \in \hat{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ is dominated by $C_1 \in \hat{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$, the value of C_2 is updated to the next coalition in $\hat{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$. In the event that C_1 and C_2 are equivalent, they are both replaced in $\mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}}$ by the first following coalition C' in their respective maximal chains such that $|\{x,y\} \cap C'| = 1$. If, without loss of generality, there is no such coalition placed after C_1 in any maximal chain, then C_1 is simply removed from $\mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}}$. The sets $\hat{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ and $\hat{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$ are then redetermined accordingly. This procedure is repeated until a coalition from one of the sets has been established to be preferred to all others in the complementary set, in which case, from Proposition 4, we can establish the preference relation between x and y . If $\mathcal{M}_{x,y}^{\succsim_{\mathcal{P}(X)}}$ becomes empty without having uncovered any preference between any coalition from $\hat{M}_{x,-y}^{\succsim_{\mathcal{P}(X)}}$ and any coalition from $\hat{M}_{y,-x}^{\succsim_{\mathcal{P}(X)}}$, we gather that x and y must be equivalent. In this case, we flag x as being equivalent to y , and artificially remove y from the pool of possible winners. If x turns out to be the necessary winner, its flag will serve as a reminder that it is not the only necessary winner: PROSP will return the set of elements having been found to be equivalent to x .

Example 5. Let $X = \{1, 2, 3\}$. Let the total preorder $\succsim_{\mathcal{P}(X)}^*$ be such that $1 \sim 2 \succ 12 \succ 13 \succ 23 \succ 123 \succ 3$, and let the available preorder $\succsim_{\mathcal{P}(X)} \subseteq \succsim_{\mathcal{P}(X)}^*$ contain two maximal chains $1 \succ 12 \succ 13 \succ 23 \succ 123 \succ 3$ and $2 \succ 12 \succ 13 \succ 23 \succ 123 \succ 3$, from which we obtain the set of lex-cel possible winners $W = \{1, 2\}$. It is unnecessary to apply the global lex-cel to W , as we can only compare the elements 1 and 2. We easily determine that $\mathcal{M}_{1,2}^{\succsim_{\mathcal{P}(X)}} = \{\{1\}, \{2\}\}$. The first query is between $\{1\}$ and $\{2\}$, to which the answer will be that $\{1\} \sim \{2\}$. Both elements are removed from consideration for the remainder of the comparison between elements 1 and 2. The set $\mathcal{M}_{1,2}^{\succsim_{\mathcal{P}(X)}}$ is then computed again using the preorder $\succsim_{\mathcal{P}(X)}$ from which all preferences concerning either $\{1\}$ or $\{2\}$ have been removed. The coalition $\{1\}$ is replaced by $\{1, 3\}$; the coalition $\{2\}$ is replaced by $\{1, 3\}$ as well. Consequently, the updated set is now $\mathcal{M}_{1,2}^{\succsim_{\mathcal{P}(X)}} = \{\{1, 3\}\}$ which, in turn, means that $\mathcal{M}_{1,-2}^{\succsim_{\mathcal{P}(X)}} = \{\{1, 3\}\}$ and $\mathcal{M}_{2,-1}^{\succsim_{\mathcal{P}(X)}} = \emptyset$. By construction of the algorithm, this means that 1 is preferred to 2, *i.e.* $1 R_{le} 2$, from which we update the set L containing elements found, in previous steps, to be dominated by another possible winner. In this case, we have $L = \{2\}$. Since $W \setminus L = \{1\}$, there is only one possible winner left, which, by definition, means that it is the necessary winner. The procedure then returns 1 as the necessary winner, which is indeed the winner for the preorder $\succsim_{\mathcal{P}(X)}$.

6 Experimental results

We compare the performance of both elicitation methods based on the number of queries submitted over preorders defined under different constraints. All presented results are from experiments over 100 to 1000 different preorders (referred to as "runs"). Unless specified otherwise, all preorders derived from the total preorder, which serve as input information for our problem, are generated uniformly at random. By this, we mean a total order over all coalitions is generated uniformly at random. This order is then turned




Fig. 2: Number of runs (out of 1000) for which each method submits at most x queries over a population of size 4

into a preorder by coursing through all coalitions from best to worst, and making each equivalent to it next preferred with a probability 0.2. Then, coursing through coalitions of the preorder from most to least preferred, each is added to one of m maximal chains, with m the number of maximal chains also determined uniformly at random within the set $\{1, \dots, \frac{2^n}{2}\}$, with n the size of the population. Once a coalition has been added to an i -th maximal chain, $\forall i \in \{1, \dots, m-1\}$, it may be added to an additional chain with probability 0.2.⁶

In any scenario, we find that the C-RAND variant of RECO is always inefficient, as it most often submits far too many queries to the user, most of which can be deemed uninformative. Furthermore, we find that its restricted version, R-RAND, systematically outperforms it. Indeed, Figure 2 shows that, even when the population contains only 4 elements, C-RAND is significantly outperformed by every other approach over all 1000 runs, as it always requires more queries than any other method: at the threshold of 20 questions, all other variants have solved NWSR on nearly all runs, yet C-RAND has only solved it for half of the runs. Note that, for a population of size $n = 4$, there are $2^n = 16$ coalitions to rank in total, and a worst-case scenario in terms of submitted queries would compare the worst coalition to all others ($2^n - 1 = 15$ queries), then the second worst to the rest ($2^n - 2 = 14$ queries), and so on, leading to a maximum of $\sum_{i \in \{0, 2^n-1\}} (2^n - 1) - i = 120$ queries.

The performance of C-RAND only worsens as the size of the population grows. As such, and for legibility purposes, we will not include the results of this particular variant in the following sections.

⁶ Additional results under particular parameters are available in the Appendix.

Uniformly at random. In the case where the underlying total preorder is generated uniformly at random, we observe that the performance of the R-MIN_INT and R-MAX_LEN variants of RECO never produce the least amount of queries. However, the performance of the remaining two approaches varies with the size of the population.

Indeed, we observe in Figure 3 that, for a population of size smaller or equal to 6, PROSP yields a smaller number of queries than the R-RAND variant of RECO. However, for a population of size greater than or equal to 7, the R-RAND variant outperforms PROSP more and more significantly.




Fig. 3: Number of queries submitted by each method based on the size of the population.

Additive preferences Preferences expressed by a ranking $\succsim^*$ are called **additive** if there exists a utility function $f : \mathcal{P}(X) \rightarrow \mathbb{N}$ such that $\forall C \subseteq X, f(C) = \sum_{x \in C} f(x)$ verifying that, $\forall C, D \subseteq X, C \succsim^* D \Leftrightarrow f(C) \geq f(D)$. Under the hypothesis of an underlying total preorder expressing additive preferences, Figure 4, testing over 100 runs for each size of X , shows that PROSP yields systematically better results than any other variant of RECO.

Best-case scenario: a unique lex-cel possible winner. Under any restriction applied to the studied preorder, we find that PROSP maintains an enticing quality. Indeed, there exists what may be referred to as a “best-case scenario”, in which the preorder given as input is such that one could directly determine the lex-cel necessary winner without soliciting the user. This scenario arises when there is only one true lex-cel possible winner (an example is provided in the Appendix⁷).

By construction of RECO, it is essential to determine at least the first coalition of the total preorder to assess any preference over elements of X . This means that, with the exception of the very specific scenario in which all maximal chains in the initial preorder start with the same singleton, we know that at least one query will always be submitted to the user using any variant of RECO. On the other hand, because prospecting leads us to directly confront key coalitions, PROSP does not need any additional

⁷ The Appendix will be published as a technical report and is available at this url.




Fig. 4: Number of queries submitted by each method over additive total preorders

input from the user in order to determine the lex-cel necessary winner when there is only one lex-cel possible winner from the start. In simpler terms, PROSP never submits a query when there is no need for additional information from the user's expressed partial preferences, whereas the RECO will (nearly) always submit at least one query.

Discussion of the results. Under no restriction over the type of underlying total preorder, the R-RAND variant appears more the more efficient out of all proposed methods. However, PROSP remains enticing, as it submits fewer queries when the underlying preorder is additive, and maintains the notable advantage of not submitting any query when none is necessary.

Conclusions and perspectives

We introduced two elicitation approaches to determine the lex-cel necessary winners when provided with a preorder over the set of a population's coalitions. One approach aims at reconstituting enough of the underlying total preorder to be able to determine with certainty which elements are the true winners; the second determines key coalitions in the comparison of elements deemed most likely to be winners by the global lex-cel. Our experiments suggest that the performance of the different elicitation methods varies depending on the type of power relation being examined, as the random variant of RECO seems to fare better when there is no hypothesis on the total preorder's type, but PROSP yields better results on additive underlying total preorders. Furthermore, PROSP has the attractive quality to not submit queries when none is necessary.

Additionally, as both approaches aim to determine the lex-cel necessary winners with certainty, they require a great number of queries as the size of the population increases, which may still prove strenuous for the user. We believe it would also be perti-

ment to study approximated resolution methods for this problem, which could guarantee that the number of submitted queries would remain manageable to the user.

Finally, both approaches operate under the hypothesis that the underlying total pre-order is available, and any query will provide a fixed preference relation. This is a strong assumption, which could be made more flexible, and we believe it could be interesting to study scenarios in which not all queries may lead to a definite preference relation.

Acknowledgements

We acknowledge a financial support from the project THEMIS ANR-20-CE23-0018 of the French National Research Agency (ANR).

References

1. Stefano Moretti and Meltem Öztürk. Some axiomatic and algorithmic perspectives on the social ranking problem. In *Proceedings of ADT'17*, page 166–181, 2017.
2. Giulia Bernardi, Roberto Lucchetti, and Stefano Moretti. Ranking objects from a preference relation over their subsets. *Social Choice and Welfare*, 52:589–606, 2019.
3. Encarnación Algaba, Stefano Moretti, Eric Rémila, and Philippe Solal. Lexicographic solutions for coalitional rankings. *Social Choice and Welfare*, pages 1–33, 2021.
4. Sylvain Béal, Eric Rémila, and Philippe Solal. Lexicographic solutions for coalitional rankings based on individual and collective performances. *Journal of Mathematical Economics*, page 102738, 2022.
5. Takahiro Suzuki and Masahide Horita. Consistent social ranking solutions. *Social Choice and Welfare*, 62:549–569, 2024.
6. Tahar Allouche, Bruno Escoffier, Stefano Moretti, and Meltem Öztürk. Social ranking manipulability for the CP-majority, Banzhaf and lexicographic excellence solutions. In *Proceedings of IJCAI'20*, pages 17–23, 2020.
7. Adrian Haret, Khani Hossein, Stefano Moretti, and Meltem Öztürk. Ceteris paribus majority for social ranking. In *Proceedings of IJCAI'18*, pages 303–309, 2018.
8. Hossein Khani, Stefano Moretti, and Meltem Ozturk. An ordinal banzhaf index for social ranking. In *Proceedings of IJCAI'19*, pages 378–384, 2019.
9. Takahiro Suzuki and Masahide Horita. Which set of agents plays a key role? an impossibility in transforming binary relations. *Mathematical Social Sciences*, 129:12–19, 2024.
10. Sébastien Konieczny, Stefano Moretti, Ariane Ravier, and Paolo Viappiani. Selecting the most relevant elements from a ranking over sets. In *Proc. of SUM'22*, pages 172–185, 2022.
11. Kathrin Konczak and Jerome Lang. Voting procedures with incomplete preferences. *Proceedings of the IJCAI'05 Workshop on Advances in Preference Handling*, 01 2005.
12. Lirong Xia and Vincent Conitzer. Determining possible and necessary winners given partial orders. *J. Artif. Intell. Res.*, 41:25–67, 2011.
13. Jérôme Lang. Collective decision making under incomplete knowledge: Possible and necessary solutions. In *Proceedings of IJCAI'20*, pages 4885–4891, 2020.
14. Tyler Lu and Craig Boutilier. Preference elicitation and robust winner determination for single- and multi-winner social choice. *Artif. Intell.*, 279, 2020.
15. Nawal Benabbou, Serena Di Sabatino Di Diodoro, Patrice Perny, and Paolo Viappiani. Incremental preference elicitation in multi-attribute domains for choice and ranking with the borda count. In *Proceedings of SUM'16*, pages 81–95, 2016.

16. Beatrice Napolitano, Olivier Cailloux, and Paolo Viappiani. Simultaneous elicitation of scoring rule and agent preferences for robust winner determination. In *Proceedings of ADT'21*, pages 51–67, 2021.
17. Graham Brightwell and Peter Winkler. Counting linear extensions. *Order*, 8:225–242, 1991.
18. Gara Pruesse and Frank Ruskey. Generating linear extensions fast. *SIAM Journal on Computing*, 23(2):373–386, 1994.