



HAL
open science

Simultaneous Elicitation of Scoring Rule and Agent Preferences for Robust Winner Determination

Beatrice Napolitano, Olivier Cailloux, Paolo Viappiani

► **To cite this version:**

Beatrice Napolitano, Olivier Cailloux, Paolo Viappiani. Simultaneous Elicitation of Scoring Rule and Agent Preferences for Robust Winner Determination. *Algorithmic Decision Theory (ADT 2021)*, Nov 2021, Toulouse, France. pp.51-67, 10.1007/978-3-030-87756-9_4. hal-03384433

HAL Id: hal-03384433

<https://univ-paris-dauphine.hal.science/hal-03384433v1>

Submitted on 18 Oct 2021

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.

Simultaneous Elicitation of Scoring Rule and Agent Preferences for Robust Winner Determination

Beatrice Napolitano¹[0000–0002–3887–6206], Olivier Cailloux¹, and Paolo Viappiani²

¹ Université Paris-Dauphine, Université PSL, CNRS, LAMSADE

² LIP6, UMR 7606, CNRS and Sorbonne Université
beatrice.napolitano@dauphine.fr

Abstract. Social choice deals with the problem of determining a consensus choice from the preferences of different agents. In the classical setting, the voting rule is fixed beforehand and full information concerning the preferences of the agents is provided. This assumption of full preference information has recently been questioned by a number of researchers and several methods for eliciting the preferences of the agents have been proposed. In this paper we argue that in many situations one should consider as well the voting rule to be partially specified. Focusing on positional scoring rules, we assume that the chair, while not able to give a precise definition of the rule, is capable of answering simple questions requiring to pick a winner from a concrete profile. In addition, we assume that the agent preferences also have to be elicited. We propose a method for robust approximate winner determination and interactive elicitation based on minimax regret; we develop several strategies for choosing the questions to ask to the chair and the agents in order to converge quickly to a near-optimal alternative. Finally, we analyze these strategies in experiments where the rule and the preferences are simultaneously elicited.

Keywords: Uncertainty in AI · Computational Social Choice · Preference Elicitation.

1 Introduction

Aggregation of preference information is a central task in many computer systems (recommender systems, search engines, etc). In many situations, such as in group recommender systems, the goal is to find a consensus choice; social choice theory can provide foundations for such applications. The traditional approach to social choice assumes that 1) the full preference orderings of the agents and 2) the social choice function are expressed beforehand. These represent two strong hypotheses. Requiring agents to express full preference orderings can be prohibitively costly (in terms of cognitive and communication cost). This observation has motivated several works assuming partial preference orders: one early work is by Conitzer and Sandholm [7] who studied the complexity of communication when using different voting rules; Konczak and Lang [15] studied the computation of possible and necessary winners for various voting rules; Xia and Conitzer [34] then showed that, while the identification of a necessary co-winner in scoring rules is polynomial, the determination of possible co-winners is NP-hard; additional complexity results were given by Walsh [32] and Pini et al. [25].

Since in many practical situations there would be too many possible winners but no necessary winners, several works addressed the problem of agent preferences elicitation using a variety of approaches (minimax regret, Bayesian methods, etc.) with the goal of converging to a necessary winner [24,14,21,26,2,9]. Among those, Walsh [33] and Conitzer [6] analyzed when to stop the elicitation process.

A second concern is the ability of the chair (the person or organization supervising the voting process) to provide a precise definition of the voting rule, suggesting the relaxation of the second hypothesis. Indeed, it is often difficult for non-experts to formalize a voting rule on the basis of some generic preferences over a desired aggregation method. Here we provide two examples of such situations.

Consider, as a first example, a chair that is about to hire a new employee whose performances are evaluated by several experts. The members of the chair may not have a voting rule in mind at the start of the process, and might not wish to agree on a specific voting rule. However, they might be willing to answer a few questions requiring to select who should be the winner out of specific profiles.

Consider, as a second example, the reviewing process of a conference where the best paper must be elected. The agents express their preferences on the papers they reviewed, but they are not aware of the voting rule the Program Chair will apply when aggregating them. Nonetheless, reviewers are still willing to participate in the process. Also, the PC may not have a specific voting rule in mind, and she will find it hard to provide a precise scoring vector if asked. Maybe she strongly believes that being ranked once in the first position is “much more” valuable than being ranked two times second, but does not know exactly how much more (though she can judge example cases).

In this paper, we focus on positional scoring rules with convex weights, that are a particularly common method used to aggregate rankings. We develop methods, based on the notion of minimax regret, for determining a robust “winner” under uncertainty of both the voting rule and the agent preferences. We provide incremental elicitation methods that at each step of the elicitation question either one of the agents or the chair, and we discuss several heuristics to choose questions that quickly reduce the regret. Answers to questions are encoded as constraints; questions to the agents are comparisons between pairs of alternatives while questions to the chair ask to select a winner out of a synthetic profile.

While some previous works have considered partially specified aggregation methods [30,20,31], we do not know of any work considering both sources of uncertainty at the same time. Actually, very few works altogether have considered the problem of eliciting a voting rule by asking questions to the chair. We mention the work of Cailloux and Endriss [5] that assumes a different representation for the rule. Additionally, some works address the manipulability of voting rules [11,10,8,1] and strategic behaviors [12,17,27].

Our approach is evaluated on simulations with synthetic and real datasets where both the voting rule and the agent preferences are initially unknown to the system and incrementally revealed through questioning. We assume the chair to be human, thus able to answer questions about a limited number of alternatives, so we focus on small scale social choice situations. We compare the effectiveness of several questioning strategies based on the current knowledge of the rule and preferences. To summarize our con-

tributions: 1) we provide a novel mechanism for eliciting a voting rule by translating abstract questions about weights to a choice of an alternative given a concrete profile; 2) we show that with our elicitation method it is possible to reach low regret with a reasonable number of questions; 3) we present elicitation strategies that achieve good results within reasonable computation time; 4) we show that for the class of rules considered, asking a few questions to the chair suffice to reach low regret; 5) our experiments suggest that low degree of similarity among preferences (as in impartial culture) is a more challenging setting than less varied profiles.

2 Social choice with partial information

We now introduce some basic concepts. We consider a set A of m alternatives (products, restaurants, public projects, job candidates, etc.) and an infinite set \mathbb{N} of potential agents.

A *profile* $(\succ_j, j \in N)$ considers a finite subset of agents $N \subset \mathbb{N}$ and associates to each agent a preference order $\succ_j \in \mathcal{L}(A)$, a linear order over the alternatives. A profile is equivalently represented by $\mathbf{v} = (v_j, j \in N)$ where $v_j(x) \in \{1, \dots, m\}$ denotes the rank of alternative x in the preference order \succ_j . A social choice function $f : \cup_{\emptyset \neq N \subset \mathbb{N}, N \text{ finite}} \mathcal{L}(A)^N \rightarrow \mathcal{P}^*(A)$ associates to each profile a set of (tied) winners, where $\mathcal{P}^*(A)$ is the powerset of A excluding the empty set. Among the many possible social choice functions, we consider convex *positional scoring rules* (PSRs). A PSR $f^{\mathbf{w}}$ is parameterized by a *scoring vector* \mathbf{w} associating weights $w_r \in [0, 1]$ to positions, with $1 = w_1 \geq w_2 \geq \dots \geq w_m = 0$. Let α_r^x be the number of times that alternative x was ranked in the r -th position. Given \mathbf{v} and \mathbf{w} , an alternative $x \in A$ obtains the score

$$s(x; \mathbf{v}, \mathbf{w}) = \sum_{j \in N} w_{v_j(x)} = \sum_{r=1}^m \alpha_r^x w_r. \quad (1)$$

The winners $f^{\mathbf{w}}(\mathbf{v})$ are the alternatives with highest score.

An important class of PSRs is the one using convex weights [30,19], meaning that the difference between the weight of the first position and the weight of the second position is at least as large as the difference between the weights of the second and third positions, etc.

$$\forall r \in \{1, \dots, m-2\} : w_r - w_{r+1} \geq w_{r+1} - w_{r+2}. \quad (2)$$

The constraint above is a natural and common assumption, often used when aggregating rankings in sport competitions (such as F1 racing, alpine skiing world cup): losing ranks at the top is more damaging than losing ranks at the bottom. Let \mathcal{W} denote the set of such convex weight vectors.

We consider a specific finite set of agents $N^* \subset \mathbb{N}$ and let $\mathbf{v}^* = (\succ_j^*, j \in N^*)$ and \mathbf{w}^* denote the profile and weight vector, unknown to us, that represent the preferences of the agents in N^* and of the chair.

At a given time, our knowledge of agent j 's preference is encoded by a partial order $\succ_j^p \subseteq \succ_j^*$ over the alternatives, a transitive and asymmetric relation (we assume that preference information is truthful). An incomplete profile $\mathbf{p} = (\succ_j^p, j \in N^*)$ maps

each agent to a partial preference. Let $C(\succ_j^p) = \{\succ \in \mathcal{L}(A) \mid \succ_j^p \subseteq \succ\}$ denote the set of possible completions of \succ_j^p and $C(\mathbf{p}) = \prod_{j \in N} C(\succ_j^p)$ the set of complete profiles extending \mathbf{p} . Note that $\mathbf{v}^* \in C(\mathbf{p})$.

The vector \mathbf{w}^* is also unknown but we assume that the chair is able to specify additional preference information taking the form of linear constraints about \mathbf{w}^* . Let $W \subseteq \mathcal{W}$ denote the set of weight vectors compatible with the preferences expressed by the chair about the scoring vector. We will show in Section 4 that the additional preferences we use can be elicited by showing a complete profile of a synthetic election and asking who should be elected in this case.

3 Robust winner determination

It is desirable in an elicitation protocol such as ours to be able to stop before reaching full knowledge of the agent preferences or of the preferences of the chair about the voting rule. As, often, there are no necessary winners and too many possible winners, it is useful to declare a winner given partial information. As a decision criterion to determine a winner, we propose to use minimax regret [29]. This decision criterion has been used for robust optimization under data uncertainty [16] as well as in decision-making with uncertain utility values [28,3]. In particular, Lu and Boutilier [21] have adopted minimax regret for winner determination in social choice where the preferences of agents are partially known, while the social choice function is known.

We consider the simultaneous presence of incomplete knowledge in agent preferences and in the weights of the PSR. We use *maximum regret* to quantify the worst-case error, and let the alternatives that minimize this quantity win, giving some robustness in face of ignorance. Intuitively, the quality of a proposed alternative a is how far a is from the optimal one in the worst case, given the current knowledge.

Given \mathbf{p} and W (that represent the current knowledge about agent preferences and the PSR), the maximum regret is considered by assuming that an adversary can both 1) extend the partial profile \mathbf{p} into a complete profile, and 2) instantiate the weights choosing among any weight vector in W . We formalize the notion of minimax regret in multiple steps. First of all, $\text{Regret}(x, \mathbf{v}, \mathbf{w})$ is the “regret” of selecting x as a winner instead of the optimal alternative under \mathbf{v} and \mathbf{w} :

$$\text{Regret}(x, \mathbf{v}, \mathbf{w}) = \max_{y \in A} s(y; \mathbf{v}, \mathbf{w}) - s(x; \mathbf{v}, \mathbf{w}).$$

The *pairwise maximum regret* of x relative to y given the partial profile \mathbf{p} and the set of weights W is the worst-case loss of choosing x instead of y under all possible realizations of the full profile *and* all possible instantiations of the weights:

$$\text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} \max_{\mathbf{v} \in C(\mathbf{p})} s(y; \mathbf{v}, \mathbf{w}) - s(x; \mathbf{v}, \mathbf{w}).$$

The maximum regret is the worst-case loss of x :

$$\text{MR}(x; \mathbf{p}, W) = \max_{y \in A} \text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} \max_{\mathbf{v} \in C(\mathbf{p})} \text{Regret}(x; \mathbf{v}, \mathbf{w}). \quad (3)$$

$\text{MR}(x; \mathbf{p}, W)$ is the result of an adversarial selection of the complete profile $\mathbf{v} \in C(\mathbf{p})$ and of the scoring vector $\mathbf{w} \in W$ that jointly maximize the loss between x and the true winner under \mathbf{v} and \mathbf{w} . Finally, $\text{MMR}(\mathbf{p}, W) = \min_{x \in A} \text{MR}(x; \mathbf{p}, W)$ is the value of *minimax regret* under \mathbf{p} and W , obtained when recommending a *minimax optimal* alternative $x_{\mathbf{p}, W}^* \in A_{\mathbf{p}, W}^* = \text{argmin}_{x \in A} \text{MR}(x; \mathbf{p}, W)$. Picking as consensus choice an alternative associated with minimax regret provides a recommendation that gives worst-case guarantees. In cases of ties, we can return all minimax alternatives $A_{\mathbf{p}, W}^*$ as winners or pick one of them using some tie-breaking strategy.

Observe that if $\text{MMR}(\mathbf{p}, W) = 0$, then any $x_{\mathbf{p}, W}^* \in A_{\mathbf{p}, W}^*$ is a necessary winner: any valid completion of the profile and choice of $\mathbf{w} \in W$ gives to $x_{\mathbf{p}, W}^*$ the highest score.

We note that our notion of regret gives some cardinal meaning to the scores: instead of just being used to select winners under the corresponding PSR, their differences are considered as representing the regret of the chair.

Computation of minimax regret Given a voting rule and a partially specified profile, Xia and Conitzer [34] determine necessary winners by showing constructions that attempt to maximize the score difference between a proposed winner and a chosen alternative. This reasoning was later adopted by Lu and Boutilier [21] who used the considerations on the worst-case completions for computing the minimax regret.

In order to compute pairwise maximum regret, and therefore minimax regret, we decompose the PMR into the contributions associated to each agent by adapting this same reasoning to our setting. The context is however more challenging due to the presence of uncertainty in the weights.

Recall that, in the computation of $s(x; \mathbf{v}, \mathbf{w})$, $w_{v_j(x)}$ represents the score that x obtains in the ranking v_j (see Eq. (1)). Since scoring rules are additively decomposable, we can consider separately the contribution of each agent to the total score. Thus, we can write the actual regret of choosing x instead of y as $s(y; \mathbf{v}, \mathbf{w}) - s(x; \mathbf{v}, \mathbf{w}) = \sum_{j \in N} w_{v_j(y)} - w_{v_j(x)}$, and we obtain

$$\text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} \sum_{j \in N} \max_{v_j \in C(\succ_j^{\mathbf{p}})} [w_{v_j(y)} - w_{v_j(x)}].$$

The following propositions show that the procedure for completing a partial profile, proposed by Lu and Boutilier [21] when considering a fixed weight vector, also applies in our setting. We write $a \succeq_j^{\mathbf{p}} b$ iff $a \succ_j^{\mathbf{p}} b \vee a = b$ and adopt the canonical notation when considering a relation as a function, writing $\succeq_j^{\mathbf{p}}(x)$ for $\{y \mid x \succeq_j^{\mathbf{p}} y\}$.

Proposition 1. *There exists a completion $\hat{\mathbf{v}} \in C(\mathbf{p})$ of the partial profile \mathbf{p} such that $\text{PMR}(x, y; \mathbf{p}, W) = \max_{\mathbf{w} \in W} [s(y; \hat{\mathbf{v}}, \mathbf{w}) - s(x; \hat{\mathbf{v}}, \mathbf{w})]$ and such that the linear order \hat{v}_j of each agent j satisfies:*

$$a \succ_j x \Leftrightarrow \neg(x \succeq_j^{\mathbf{p}} a); \quad (4)$$

$$y \succ_j a \Leftrightarrow \neg(a \succeq_j^{\mathbf{p}} y) \wedge \neg((x \succeq_j^{\mathbf{p}} y) \wedge \neg(x \succeq_j^{\mathbf{p}} a)). \quad (5)$$

Proof Sketch. Consider our knowledge $\succeq_j^{\mathbf{p}}$ about the preference of the agent j . The adversary's goal is to make the score of y as high as possible and the score of x as

low as possible. To do this, he should complete \succ_j^p to \succ_j by placing above x as many alternatives as possible; that is, all the alternatives except those that are known to be worse than x (those a such that $x \succeq_j^p a$); and similarly, he should put below y all the alternatives he can. Two conditions must be excluded for a to go below y . The alternatives such that $a \succeq_j^p y$ can't be put below y . Furthermore, the first objective must take priority over the second one: when an alternative should go above x according to the first objective (because $\neg(x \succeq_j^p a)$), and x is known to be better than y (thus $x \succeq_j^p y$), then a should be put above x (irrespective of whether $a \succeq_j^p y$), which will move both x and y one rank lower than if a had been put below y . This maximizes the adversary's interests: because the weight vector is convex, the score difference will be lower when both alternatives are ranked lower (Equation 2), and that difference of scores is in favor of x when $x \succ_j^p y$, thus to be minimized from the the adversary's point of view. \square

Proposition 2. *The rank of x in the PMR-maximizing linear orders of agent j is $\hat{v}_j(x) = 1 + |A| - |\succeq_j^p(x)|$, and the rank of y is $\hat{v}_j(y) = 1 + |\prec_j^p(y)| + |\beta|$, where $|\beta| = |A \setminus (\succeq_j^p(x) \cup \prec_j^p(y))|$ if $(x \succeq_j^p y)$ and $|\beta| = 0$ otherwise.*

Proof. The rank of x is directly obtained from Eq.(4). The rank of y is obtained by complementing Eq.(5), obtaining $a \succeq_j y \Leftrightarrow (a \succeq_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a))$, and, observing that $a \succ_j y \Leftrightarrow a \neq y \wedge a \succeq_j y$, obtaining that $a \succ_j y$ if and only if

$$(a \neq y) \wedge [(a \succeq_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a))], \quad (6)$$

or equivalently, if and only if

$$(a \succ_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a)). \quad (7)$$

Indeed, (6) \Rightarrow (7), and (7) \Rightarrow (6) because $(x \succeq_j^p y) \wedge \neg(x \succeq_j^p a) \Rightarrow a \neq y$ (as when $a = y$, $(x \succeq_j^p y)$ and $\neg(x \succeq_j^p a)$ are opposite claims). Suffices now to rewrite Eq. (7) to let the two disjuncts designate disjoint sets:

$$a \succ_j y \Leftrightarrow (a \succ_j^p y) \vee ((x \succeq_j^p y) \wedge \neg(x \succeq_j^p a) \wedge \neg(a \succ_j^p y)). \quad (8)$$

\square

Note that in Proposition 2, in the case $(x \succeq_j^p y)$, β is the number of alternatives incomparable with both x and y .

Proposition 3. *The PMR can be written as:*

$$\text{PMR}(x, y; \mathbf{p}, W) = \max_{w \in W} \sum_{j \in N} w_{\hat{v}_j(y)} - w_{\hat{v}_j(x)} = \max_{w \in W} \sum_{r=1}^m (\hat{\alpha}_r^y - \hat{\alpha}_r^x) w_r, \quad (9)$$

where $\hat{\alpha}_r^y$ (resp. $\hat{\alpha}_r^x$) is the number of times y (resp. x) has rank r in the complete profile \hat{v} defined in Proposition 2.

Proposition 3 shows that PMR is linear in the weights. The pairwise max regret $\text{PMR}(x, y; \mathbf{p}, W)$ can thus be obtained by solving the following linear program defined on the variables w_1, \dots, w_m :

$$\max_{\mathbf{w}} \sum_{r=1}^m (\hat{\alpha}_r^y - \hat{\alpha}_r^x) w_r \quad \text{s.t. } w_1 = 1 \geq \dots \geq w_m = 0, \text{ Eq. (2) and } \mathbf{w} \in W. \quad (10)$$

The max regret $\text{MR}(x; \mathbf{p}, W)$ is determined by computing the pairwise regret of x with all other alternatives in A , and the recommended alternatives are the ones with least max regret. Observe that when the PMR of an alternative x (against some other alternative y) exceeds the best MR value found so far, we do not need to further evaluate x . This idea can be exploited using a minimax-search tree [4].

4 Interactive Elicitation

We propose an incremental elicitation method based on minimax regret. At each step, the system may ask a question either to one of the agents about her preferences or to the chair about the voting rule. The goal is to obtain relevant information to reduce minimax regret as quickly as possible. The elicitation can be terminated either after a given number of questions, or when the minimax regret is lower than a threshold (or when it drops to zero if we wish optimality).

Question types We distinguish between questions asked to the agents and questions asked to the chair. As *questions asked to the agents* we consider comparison queries relating two alternatives. The effect of a response to a question asked to an agent is the increase in our knowledge about the agent rankings, thus augmenting the partial profile \mathbf{p} . If agent j answers a comparison query stating that alternative a is preferred to b , then the partial order $\succ_j^{\mathbf{p}}$ is augmented with $a \succ_j^{\mathbf{p}} b$ and by transitive closure.

A bit more discussion is needed about *questions asked to the chair*. Such questions aim at refining our knowledge about the scoring rule; a response gives us a constraint on the weight vector \mathbf{w} . In particular, we want to obtain constraints of the type $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$ for $r \in \{1, \dots, m-2\}$, relating the difference between the importance of ranks r and $r+1$ with the difference between ranks $r+1$ and $r+2$.

Building concrete questions for the chair Even if the chair might be considered able to answer directly such abstract questions, we want to ensure that these questions can also, in principle, be asked in a more concrete way: in terms of winners of example profiles. Such questions have clear semantics whose understanding can be assumed to be shared by the chair, contrary to abstract questions about weights. Moreover, this way of questioning the chair is independent of the voting rule that is being elicited; whereas questions about weights only make sense when considering PSRs. Asking who should win in specific profiles has been used in experimental settings investigating the feeling of justice of individuals [13], but, to the best of our knowledge, the use of such questions to systematically guide an elicitation process about voting rules is novel. This is similar to favor, in decision theory, direct choice questions ("please choose either a or

b") compared to, say, questioning the decision maker about the shape of her utility function. The former are considered “observable”: acts of choice are translated to preference statements [22, Ch. 1].

Although questioning in terms of profiles and in terms of weights is logically equivalent in our setting, there is no a priori certainty that questioning the chair using different phrasing would yield logically equivalent answers: research in experimental psychology shows that participants’ answers differ widely when changing the phrasing of preference-related questions [18]. To get out of such conundrums, we need a language considered “fundamental”. Questions of the form “In this profile, who should win?” arguably provides such a natural language.

Thus, our task is to build a profile, given λ and $r \leq m - 2$, in such a way that the set of (tied) winners picked by the chair reveals whether $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$.

Proposition 4. *Given a rational $\lambda = p/q > 1$ and a rank r between 1 and $m - 2$, there exists a profile P such that, for any weight vector $w \in \mathcal{W}$, $a \in f(P)$ iff $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$ and $b \in f(P)$ iff $w_r - w_{r+1} \leq \lambda(w_{r+1} - w_{r+2})$, where f is the PSR parameterized with w .*

Proof. Define a linear order $>_1$ over A as placing a at rank r , b at rank $r + 1$, and the remaining alternatives arbitrarily. Define $>_2$ over A as placing a at rank $r + 2$, b at rank $r + 1$, and the remaining alternatives arbitrarily. Define an arbitrary linear ordering $>$ over $A \setminus \{a, b\}$. Define a linear order $>_3$ as placing a first, b second, and following the order of $>$ for the remaining positions. Finally, define a linear order $>_4$ as placing b first, a second, and following the *inverse* order of $>$ for the remaining positions.

Define P as the profile of $3(p+q)$ agents containing q times $>_1$, p times $>_2$, and $>_3$ and $>_4$ each $p + q$ times. As a result, a obtains the following ranks: q times r , p times $r + 2$, $p + q$ times first, and $p + q$ times second. The alternative b obtains the ranks $r + 1$, 2 and 1, each $p + q$ times. Consider any alternative $c \in A \setminus \{a, b\}$. Its score is maximal when it comes first in $>_1$, first in $>_2$ and first in $>$, by convexity of the weights. In that case, c is positioned at the ranks 1, 3 and m , each $p + q$ times.

Letting $s(x)$ denote the score of x at P , we obtain $s(a) = qw_r + pw_{r+2} + (p+q)w_1 + (p+q)w_2$, thus, $s(a) \geq (p+q)w_m + (p+q)w_1 + (p+q)w_2$; $s(b) = (p+q)w_{r+1} + (p+q)w_2 + (p+q)w_1$; and, $\forall c \in A \setminus \{a, b\}$, $s(c) \leq (p+q)w_1 + (p+q)w_3 + (p+q)w_m$. It follows that a or b maximize s (as $s(a) \geq s(c)$). We conclude by observing that $a \in f(P) \Leftrightarrow s(a) \geq s(b) \Leftrightarrow qw_r + pw_{r+2} \geq (p+q)w_{r+1} \Leftrightarrow w_r - w_{r+1} \geq (p/q)(w_{r+1} - w_{r+2})$, and similarly for $b \in f(P)$. \square

Example. Suppose we want to ask the following question to the chair: $w_2 - w_3 \geq 2(w_3 - w_4)$. We show the profile in Figure 1a to the chair and ask who should win (each column is the preference of one agent). Both a and b have scores higher than c and d for all convex weights, thus either a or b will be picked under our hypothesis; and $s(a) \geq s(b) \Leftrightarrow w_2 + 2w_4 \geq 3w_3$. Figure 1b represents the same profile using a compressed view, the numbers in bold indicating the number of agents having the preference in the corresponding column. As the proof shows, constructed profiles require only four different linear orders.

Fig. 1: Profile representing a question to the chair in extended (a) and compact (b) form.

$c d d a a a b b b$ $a c c b b b a a a$ $b b b c c c d d d$ $d a a d d d c c c$	1 2 3 3 $c d a b$ $a c b a$ $b b c d$ $d a d c$
(a)	(b)

Elicitation strategies We develop several strategies for simultaneous elicitation of agent preferences and of the PSR. While it is of course possible to first fully elicit the agent preferences and afterwards elicit weights, we want to investigate approaches that are able to recommend winning alternatives before obtaining complete knowledge of the profile or the rule. We define here various strategies; a strategy tells us, given the current partial knowledge (\mathbf{p}, W) , which question to ask next.

The *Random* strategy is used as a baseline. It first chooses equiprobably whether to question the chair or the agents. In the first case, it draws one rank in $1 \leq r \leq m - 2$ equiprobably, takes the middle of the interval of values for λ that are still possible considering our knowledge so far, and asks whether $w_r - w_{r+1} \geq \lambda(w_{r+1} - w_{r+2})$. In the second case, it draws equiprobably among the agents whose preference is not known entirely; it then draws an alternative a among those involved in some incomparabilities in \succ_j^p and an alternative b among those incomparable with a in \succ_j^p .

Let $(x^*, \bar{y}, \bar{v}, \bar{w})$ be the current solution of the minimax regret, where x^* is the minimax optimal alternative and $\bar{y}, \bar{v}, \bar{w}$ the corresponding adversarial choices. The *Pessimistic* strategy considers a set of $n + (m - 2)$ candidate questions: one per agent, and one per rank (excluding the first and the last one which are known).

The candidate questions to the agents are chosen by extending the idea of Lu and Boutilier [21], that privilege learning about the relationship of x^* and \bar{y} to the other alternatives if possible. Given $j \in N^*$, if x^* and \bar{y} are incomparable in \succ_j^p , the candidate question concerns the pair (x^*, \bar{y}) , otherwise, it concerns the pair (x^*, z) for some z incomparable to x^* (randomly chosen), or if none such z exist, the pair (\bar{y}, z) for some z incomparable to \bar{y} , or, if both x^* and \bar{y} are comparable to every alternatives in \succ_j^p , any incomparable pair is picked at random.

The candidate questions to the chair are determined as in the Random strategy.

Once having selected $n + m - 2$ candidate questions, the Pessimistic strategy selects the one that leads to minimal regret in the worst case. Assume that a question q_1 has type t_1 (being ‘‘chair’’ or ‘‘agent’’), and leads to the new knowledge states (\mathbf{p}_1, W_1) if answered positively and (\mathbf{p}'_1, W'_1) if the answer is negative. Define

$$R_1^{\max} = \max\{\text{MMR}(\mathbf{p}_1, W_1), \text{MMR}(\mathbf{p}'_1, W'_1)\}$$

and

$$R_1^{\min} = \min\{\text{MMR}(\mathbf{p}_1, W_1), \text{MMR}(\mathbf{p}'_1, W'_1)\}\epsilon_t + \epsilon'_t.$$

The terms ϵ_t and ϵ'_t are real numbers associated to the type t of question; these parameters are used to fine tune the choice of the question type. Define similarly t_2, R_2^{\max}

and R_2^{\min} for q_2 . Pessimistic considers question q_1 to be better than q_2 iff $R_1^{\max} < R_2^{\max}$ or $[R_1^{\max} = R_2^{\max} \text{ and } R_1^{\min} < R_2^{\min}]$. In other words if the maximal *a posteriori* MMR of two questions are (approximately) equal, then it considers the (penalized) minimal MMR values.

The *Extended pessimistic* strategy uses the same criterion as the pessimistic strategy, but extending it to a bigger set of candidate questions, the same as those considered by the Random strategy. These candidate questions are then evaluated using the same operator as for the Pessimistic strategy. Extended pessimistic is applicable only to very small problem instances: its complexity is in $O(n^2m^5)$, because we consider $O(m^2)$ questions for each agent and need for each question to compute MMR twice, whose complexity is $O(nm^3)$.

The *Two phases* strategy is developed in order to investigate the effect of varying the proportion of questions of the two types, when asking all questions to the chair at the beginning or at the end. It is parameterized by q_c , the number of questions to be asked to the chair. The *Two phases-ca* variant first asks q_c questions to the chair, then $k - q_c$ questions to the agents, using in both cases Pessimistic to select the specific questions; whereas the *Two phases-ac* variant starts with $k - q_c$ questions to the agents, then questions the chair.

Finally, the *Elitist* strategy aims at uncovering as quickly as possible the top alternatives of all agents. For any agent j , it asks to compare an alternative currently undominated in \succ_j^p with one that is currently incomparable. Thus, the top alternative for j will be known after having asked exactly $m - 1$ questions to j . After having asked $n(m - 1)$ questions to the agents, it questions the chair only, using the same approach as Pessimistic. This strategy can be expected to perform well when the chair assigns a large weight to the first rank, as compared to the other ranks. It is used to further challenge Pessimistic, which is not specifically tailored to such a situation.

5 Empirical Evaluation

We performed several numerical experiments using both real data and randomly generated profiles in order to validate our approach and test the performance of our elicitation strategies.

Given a problem size (m, n) , a number of questions k and a strategy to test, we first create an “oracle”, representing the true preferences of the agents (randomly generated or coming from real data) and the weights associated with the chair’s scoring rule (randomly generated). We start with empty knowledge ($\mathbf{p} = \emptyset, W = \mathcal{W}$) about the preference orderings of the agents and the weights of the chair. We obtain the first question to be asked using the strategy under test. We then use the oracle to answer the question and update the system’s knowledge, which is thus used to obtain the next question. This is repeated until k answers have been obtained, computing the resulting MMR values along the way for various values of k . We repeat this whole experiment a variable number of times, for a given (m, n, k) , and report the average resulting MMR and standard deviation *sd*. The sizes of the considered scenarios are comparable to the ones used by Cailloux and Endriss [5].

The oracle is built as follows. For the real preferences, we used three datasets from PrefLib [23]: *T Shirt* (researchers voted on tee shirt designs; $m = 11, n = 30$), *Courses* (students voted on courses; $m = 9, n = 146$; referred to as AGH on PrefLib) and *Skate* (judges voted on skaters at the Euros Pairs Short Program; $m = 14, n = 9$). For the synthetic datasets, we follow an Impartial Culture (IC) assumption: the linear order of each agent is drawn i.i.d. uniformly. We believe IC to be a challenging situation and expect the number of questions to ask, in order to reach a certain level of regret, to decrease with less varied profiles. To generate the scoring rule weights, we first draw $m - 1$ numbers uniformly at random (in the interval $[0, 1]$ representing weight “differences”), normalize and sort them; a sequence of convex decreasing weights is then obtained by a decumulative sum. The penalty parameters for the Pessimistic and Extended pessimistic strategies are $\epsilon_{\text{chair}} = 1.1$, $\epsilon'_{\text{chair}} = 10^{-6}$, $\epsilon_{\text{agent}} = 1.0$ and $\epsilon'_{\text{agent}} = 0$.

Fig. 2: Average MMR in problems of size (5, 10) after k questions.

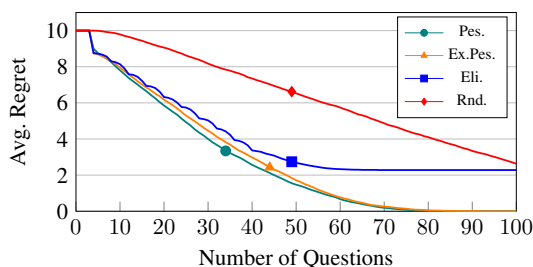


Table 1: Average MMR in problems of size (10, 20) after k questions assuming geometric weights.

k	Pes. \pm sd	Eli. \pm sd
0	20.0 \pm 0.0	20.0 \pm 0.0
50	16.0 \pm 0.5	17.3 \pm 0.4
100	12.5 \pm 0.9	15.6 \pm 0.4
150	9.6 \pm 1.4	13.9 \pm 0.8
200	7.4 \pm 1.3	11.0 \pm 1.1
250	5.3 \pm 1.5	6.6 \pm 0.8
300	3.5 \pm 1.3	6.6 \pm 0.8

Comparison of strategies Our first experiment concerns small size situations. Figure 2 compares some of our strategies in the case $m = 5, n = 10$ (variations around this size yield similar conclusions), where the results are averaged over 200 runs. We see that asking random questions does not allow to reach a low regret level even after having asked 100 questions, whereas a low regret level (MMR = 1) is reached by Pessimistic before having asked 60 questions. This also holds for other problem sizes. For instance, for $m = 10, n = 20$ and 500 questions, Random strategy reaches an average regret (over 20 runs) of 9.3 (± 0.7) and Pessimistic 0.5 (± 0.5). We notice that Pessimistic performs slightly better than Extended pessimistic, showing that Pessimistic chooses candidate questions wisely; this is good news since Pessimistic is much faster: it takes

on average only 16s for a complete elicitation session (for $m = 5$, $n = 10$ and 100 questions), while Extended pessimistic takes 50s. Although their performance is close, Pessimistic performs systematically better in multiple runs of the experiment.

We also compared the Pessimistic strategy against Elitist in a situation specifically tailored to advantage Elitist. For that experiment specifically, instead of drawing the weights of the oracle randomly, we fix it to a “geometric” weight vector, such that $w_r - w_{r+1} = 2(w_{r+1} - w_{r+2})$, for all $r \leq m - 2$, so as to dramatically increase the importance of the weights associated to the top ranks. Even in that case, we see in Table 1 that Pessimistic performs better than Elitist.

Evaluation of Pessimistic Strategy Our next set of experiments evaluate the Pessimistic strategy in absolute terms. We first wonder how many questions should be asked in order to achieve low regret, fixed at $n/10$: this is equivalent to the difference of score of an alternative x that results from switching from a profile P to a profile P' where a tenth of the agents rank x last instead of first. Table 2, first five columns, contains the result: it displays, for each dataset, the number of questions asked to the chair ($q_c^{\text{MMR} \leq n/10}$), and the quartiles of the number of questions asked to the agents ($q_a^{\text{MMR} \leq n/10}$), averaged over 20 runs. It is interesting to note that about twenty or thirty questions per agent on average suffice to reach a low regret in those instances. We find also noteworthy that the Pessimistic strategy chooses to ask zero questions to the chair but still achieves low regret, in most of those instances.

Another interesting measure is the average number of questions asked to the chair ($q_c^{\text{MMR}=0}$) and to the agents ($q_a^{\text{MMR}=0}$) before reaching zero regret. The results for various sizes are displayed in the last two columns of Table 2. Here, we see that the Pessimistic strategy does choose to question the chair when reaching low enough regret values. The m15n30 dataset did not reach zero regret in 1000 questions.

Table 2: Questions asked by Pessimistic strategy on several datasets to reach $\frac{n}{10}$ regret, columns 4 and 5, and zero regret, last two columns.

dataset	m	n	$q_c^{\text{MMR} \leq n/10}$	$q_a^{\text{MMR} \leq n/10}$	$q_c^{\text{MMR}=0}$	$q_a^{\text{MMR}=0}$
m5n20	5	20	0.0	[4.3 5.0 5.8]	5.3	[5.4 6.2 7.2]
m10n20	10	20	0.0	[13.9 16.1 18.4]	32.0	[19.7 21.8 24.7]
m11n30	11	30	0.0	[16.6 19.0 22.3]	45.2	[23.1 25.7 28.9]
tshirts	11	30	0.0	[13.1 16.6 19.6]	43.2	[28.2 32.0 35.6]
courses	9	146	0.0	[6.0 7.0 7.0]	0.0	[6.8 7.0 7.0]
m14n9	14	9	5.4	[30.3 33.5 36.7]	64.1	[37.6 40.5 44.3]
skate	14	9	0.0	[11.4 11.6 12.3]	0.0	[11.5 11.8 12.8]
m15n30	15	30	0.0	[25.0 29.5 33.7]		

Figure 3 shows the decrease in MMR according to the number of questions asked for various problem sizes. In particular, this shows important differences between some real datasets and the problems generated using IC. In the *Skate* problem, the value $\text{MMR} = 1$ is reached after less than 100 questions, while the IC case of the same size ($m = 14$, $n = 9$) requires more than 200 questions to reach that value. This reasoning also applies to

the *Courses* dataset but not to the *T Shirt* dataset. This can be explained by the high degree of similarity in the preference rankings of the *Skate* and the *Courses* problems, which helps reducing the regret faster. For example, in *Skate* the top-2 alternatives are the same for all agents, and 8 out of 9 agents rank the same alternative at position 3. By contrast, in *T Shirt*, the alternatives are evenly distributed in the preference rankings.

Fig. 3: Average MMR (normalized by n) after k questions with Pessimistic strategy for different datasets.




Table 3: Average MMR in problems of size $(10, 20)$ after 500 questions, among which q_c to the chair.

q_c	2 ph. ca \pm sd	2 ph. ac \pm sd
0	0.6 ± 0.6	0.6 ± 0.6
15	0.5 ± 0.5	0.6 ± 0.6
30	0.3 ± 0.5	0.4 ± 0.5
50	0.1 ± 0.2	0.1 ± 0.3
100	0.4 ± 0.6	0.2 ± 0.5
200	2.1 ± 1.6	2.4 ± 1.2
300	5.8 ± 1.7	6.7 ± 1.5
400	11.3 ± 1.1	11.9 ± 1.2
500	20.0 ± 0.0	20.0 ± 0.0

Comparison with Two Phases The experiments so far let the strategy free to question either the chair or an agent at each step. One may wonder what is lost in terms of regret by asking different proportions of questions to the chair and the agents. Such restrictions may be useful because of (partial) unavailability of the chair, or because the estimated cognitive costs may differ sensibly.

Table 3 shows the MMR value reached in problems of size $m = 10, n = 20$ after 500 questions, using the Two phases strategy, in the “ca” (chair then agents) and in the “ac” (agents then chair) variants. These numbers are to be compared with the MMR value reached after 500 questions with the Pessimistic strategy (displayed in Fig. 3),

which is 0.7; the Pessimistic strategy asks on average 13 (± 13) questions to the chair in this setting. The line $q_c = 0$, where no question is asked to the chair, suggest that it is possible to obtain a good-quality recommendation while knowing only that the voting rule is a scoring rule with convex weights, which is our basic hypothesis. However, we observe that asking no questions to the chair does not permit to reach $\text{MMR} = 0$. The strategy, indeed, obtains full knowledge of the profile after an average of 500 questions to the agents but never reaches zero.

6 Conclusions

In this paper we have considered a social choice setting with partial information about the agent preferences and voting rule. We have proposed the use of minimax regret both as a means of robust winner determination and as a guide to the process of simultaneous elicitation of preferences and voting rule. Our experimental results suggest that regret-based elicitation is effective and allows to quickly reduce worst-case regret significantly. They also show that, in our setting, good quality (low regret) recommendations can be achieved short of having full knowledge of weights or profile.

As part of our contribution, we provide an open-source library that can be found at <https://github.com/oliviercailloux/minimax>, to reproduce our experiments and perform many more.

Some directions for future works include developing new elicitation strategies, considering alternative heuristics; extending the elicitation to voting rules beyond scoring rules; eliciting preferences while restraining to concrete and easy questions.

References

1. Baumeister, D., Hogrebe, T.: Manipulative design of scoring systems. In: Proc. of AAMAS'19 (2019), <https://dl.acm.org/doi/10.5555/3306127.3331928>
2. Benabbou, N., Di Sabatino Di Diodoro, S., Perny, P., Viappiani, P.: Incremental preference elicitation in multi-attribute domains for choice and ranking with the borda count. In: Proc. of SUM 2016 (2016)
3. Boutilier, C., Patrascu, R., Poupart, P., Schuurmans, D.: Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence* (2006). <https://doi.org/10.1016/j.artint.2006.02.003>
4. Braziunas, D.: Decision-Theoretic Elicitation of Generalized Additive Utilities. Ph.D. thesis (2012), <https://dl.acm.org/doi/book/10.5555/2519022>
5. Cailloux, O., Endriss, U.: Eliciting a suitable voting rule via examples. In: Proc. of ECAI 2014 (2014). <https://doi.org/10.3233/978-1-61499-419-0-183>
6. Conitzer, V.: Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research* (2009). <https://doi.org/10.1613/jair.2606>
7. Conitzer, V., Sandholm, T.: Communication complexity of common voting rules. In: Proc. of EC'05 (2005). <https://doi.org/10.1145/1064009.1064018>
8. Conitzer, V., Walsh, T., Xia, L.: Dominating manipulations in voting with partial information. In: AAAI'11 (2011), <https://dl.acm.org/doi/10.5555/2900423.2900525>
9. Dey, P., Misra, N.: Preference elicitation for single crossing domain. In: Proc. of IJCAI'16 (2016), <https://dl.acm.org/doi/10.5555/3060621.3060653>

10. Dey, P., Misra, N., Narahari, Y.: Complexity of manipulation with partial information in voting. *Theoretical Computer Science* (2018). <https://doi.org/10.1016/j.tcs.2018.03.012>
11. Elkind, E., Erdélyi, G.: Manipulation under voting rule uncertainty. In: Proc. AAMAS'12 (2012), <https://dl.acm.org/doi/10.5555/2343776.2343786>
12. Endriss, U., Obraztsova, S., Polukarov, M., Rosenschein, J.S.: Strategic voting with incomplete information. In: IJCAI'16 (2016), <https://dl.acm.org/doi/10.5555/3060621.3060655>
13. Giritligil, A., Sertel, M., Kara, A.: Does majoritarian approval matter in selecting a social choice rule? an exploratory panel study. *Social Choice and Welfare* (2005). <https://doi.org/10.1007/s00355-005-0024-8>
14. Kalech, M., Kraus, S., Kaminka, G.A., Goldman, C.V.: Practical voting rules with partial information. *Autonomous Agents and Multi-Agent Systems* (2011). <https://doi.org/10.1007/s10458-010-9133-6>
15. Konczak, K., Lang, J.: Voting procedures with incomplete preferences. In: IJCAI'05 (2005)
16. Kouvelis, P., Yu, G.: *Robust Discrete Optimization and Its Applications*. Kluwer (1997)
17. Lev, O., Meir, R., Obraztsova, S., Polukarov, M.: Heuristic voting as ordinal dominance strategies. Proc. of the AAAI'19 (2019). <https://doi.org/10.1609/aaai.v33i01.33012077>
18. Lichtenstein, S., Slovic, P.: *The Construction of Preference*. Cambridge University Press (2006). <https://doi.org/10.1017/CBO9780511618031>
19. Llamazares, B.: Ranking candidates through convex sequences of variable weights. *Group Decision and Negotiation* (2016). <https://doi.org/10.1007/s10726-015-9452-8>
20. Llamazares, B., Peña, T.: Aggregating preferences rankings with variable weights. *European Journal of Operational Research* (2013). <https://doi.org/10.1016/j.ejor.2013.04.013>
21. Lu, T., Boutilier, C.: Robust approximation and incremental elicitation in voting protocols. In: Proc. of IJCAI'11 (2011). <https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-058>
22. Mas-Colell, A., Whinston, M.D., Green, J.R.: *Microeconomic Theory* (1995)
23. Mattei, N., Walsh, T.: Preflib: A library of preference data PREFLIB.ORG. In: Proc. of ADT 2013 (2013), <https://www.preflib.org/>
24. Naamani-Dery, L., Golan, I., Kalech, M., Rokach, L.: Preference elicitation for group decisions using the borda voting rule. *Group Decision and Negotiation* (2015). <https://doi.org/10.1007/s10726-015-9427-9>
25. Pini, M.S., Rossi, F., Venable, K.B., Walsh, T.: Incompleteness and incomparability in preference aggregation. In: Proc. of IJCAI'07 (2007). <https://doi.org/10.1016/j.artint.2010.11.009>
26. Pini, M.S., Rossi, F., Venable, K.B., Walsh, T.: Aggregating partially ordered preferences. *Journal of Logic and Computation* (2009). <https://doi.org/10.1093/logcom/exn012>
27. Reijngoud, A., Endriss, U.: Voter response to iterated poll information. In: Proc. of AAMAS'12 (2012), <https://dl.acm.org/doi/10.5555/2343776.2343787>
28. Salo, A.A., Hämäläinen, R.P.: Preference ratios in multiattribute evaluation (PRIME)-elicitation and decision procedures under incomplete information. *IEEE Trans. on Systems, Man and Cybernetics* (2001). <https://doi.org/10.1109/3468.983411>
29. Savage, L.J.: *The Foundations of Statistics*. Wiley (1954)
30. Stein, W.E., Mizzi, P.J., Pfaffenberger, R.C.: A stochastic dominance analysis of ranked voting systems with scoring. *EJOR* (1994). [https://doi.org/10.1016/0377-2217\(94\)90205-4](https://doi.org/10.1016/0377-2217(94)90205-4)
31. Viappiani, P.: Positional scoring rules with uncertain weights. In: *Scalable Uncertainty Management* (2018). https://doi.org/10.1007/978-3-030-00461-3_21
32. Walsh, T.: Uncertainty in preference elicitation and aggregation. In: Proc. of AAAI'07 (2007), <https://dl.acm.org/doi/10.5555/1619645.1619648>
33. Walsh, T.: Complexity of terminating preference elicitation. In: Proc. of AAMAS'08 (2008), <https://dl.acm.org/doi/abs/10.5555/1402298.1402357>
34. Xia, L., Conitzer, V.: Determining possible and necessary winners under common voting rules given partial orders. In: Proc. of AAAI'08 (2008). <https://doi.org/10.1613/jair.3186>