



HAL
open science

Possibilistic Preference Elicitation by Minimax Regret

Loïc Adam, Sébastien Destercke

► **To cite this version:**

Loïc Adam, Sébastien Destercke. Possibilistic Preference Elicitation by Minimax Regret. 37th Conference on Uncertainty in Artificial Intelligence (UAI 2021), Jul 2021, Online, United States. pp.718-727. hal-03521623

HAL Id: hal-03521623

<https://hal.science/hal-03521623>

Submitted on 24 Jan 2022

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.

Possibilistic Preference Elicitation by Minimax Regret

Loïc Adam¹

Sébastien Destercke¹

¹UMR CNRS 7253 Heudiasyc, Sorbonne Université, Université de Technologie de Compiègne, CS 60319 - 60203 Compiègne cedex, France.

Abstract

Identifying the preferences of a given user through elicitation is a central part of multi-criteria decision aid (MCDA) or preference learning tasks. Two classical ways to perform this elicitation is to use either a robust or a Bayesian approach. However, both have their shortcoming: the robust approach has strong guarantees through very strong hypotheses, but cannot integrate uncertain information. While the Bayesian approach can integrate uncertainties, but sacrifices the previous guarantees and asks for stronger model assumptions. In this paper, we propose and test a method based on possibility theory, which keeps the guarantees of the robust approach without needing its strong hypotheses. Among other things, we show that it can detect user errors as well as model misspecification.

1 INTRODUCTION

Interacting with an agent through an elicitation process is an essential task to identify her preferences. Among the possible ways to interact, incremental elicitation [Benabbou et al., 2017] is particularly interesting, as each new question can account for the preferential information provided so far. In such settings, there are two main approaches:

- The robust approach, based on (regret) minimax optimisation [Boutilier, 2013, Bourdache and Perny, 2017] and provides strong performance guarantees, provided that two hypotheses are satisfied: (1) that the user is an oracle (never commits mistakes) and (2) that we are searching in the right family of preference models.
- The Bayesian approach, based on the use of probabilities [Bourdache et al., 2019], that can include user noise in the elicitation process through relatively complex models and approximate updating procedures

(e.g., MCMC, Gradient descent [Vendrov et al., 2020]). They also only offer guarantees in terms of expectation, rather than the strong certain ones provided by the robust approaches.

Our goal is to use a possibilistic approach rather than a probabilistic one to handle uncertainty in the elicitation. As we will argue, such an approach has two main advantages:

- It formally extends the robust approach, in the sense that we retrieve this latter one as a specific instance of our framework. We keep the same guarantees as the robust approach, without needing its strong assumption.
- It comes with a natural way to encode and measure inconsistency resulting from inconsistent assessments [Dubois and Prade, 1994]. This way we can detect inconsistencies between the user answers and the postulated model, offering an interesting tool to detect user mistakes or model misspecifications.

This contrasts with the Bayesian approaches [Bourdache et al., 2019], which do not have these two properties. Moreover, exact computations with our proposal remain tractable as long as the number of questions remains limited (which is typically the case in incremental elicitation).

In order to perform the incremental elicitation and choose the preferential information to ask the user, we adapt the commonly used Current Solution Strategy (CSS) method in robust approaches.

We introduce the necessary elements to present our method and its interest in Section 2. We then discuss our possibilistic extension of the robust approach and its interest in Section 3. Lastly, Section 4 shows some simulated experiments that confirm the potential interest of our approach.

2 SETTING UP THE STAGE

In this section, we recall both the necessary basics of standard robust approaches, as well as how we propose to use

possibility theory to model preferential information.

2.1 NOTATIONS

We consider preference problems where we assume that alternatives are evaluated by value functions. The ground-truth order for preferences is assumed to be complete. We also consider the choice problem of providing a recommendation, and not to retrieve the complete ranking.

Alternatives We denote by \mathbb{X} a finite set of available alternatives. Alternatives within \mathbb{X} are denoted $\mathbb{X} = \{x_1, x_2, \dots, x_k\}$. We assume that alternatives are summarised by q real values, the criteria, such that $x \in \mathbb{R}^q$. The i th criterion of an alternative $x \in \mathbb{X}$ is denoted x^i .

Given two alternatives $x, y \in \mathbb{X}$, we denote:

- $x \succ_p y$ if and only if x is strictly preferred to y ,
- $x \succeq_p y$ if and only if x is preferred or equally preferred to y .

Example 1 (Choosing the best sandwich (running example)). Imagine a user that wants to choose the best sandwich possible among multiple ones (the alternatives). Each sandwich is characterised by two criteria: flavour and price. Each criterion has a value between 0 (worst) and 10 (best). For example, 0 in price means an overpriced sandwich, and 10 a very cheap one. Table 1 lists some sandwiches: for example duck sandwich may have a lot of flavour, but is more expensive than the other sandwiches.

	Flavour	Price
Cheese	5	9
Duck	10	0
Fish	8	4
Ham	7	7

Table 1: Grades of sandwiches

Aggregation models We assume that an alternative x is valued by its utility, measured by a function $f_\omega(x)$ parameterised by $\omega \in \Omega$. We will speak of the model ω in the rest of the paper. Two alternatives $x, y \in \mathbb{X}$ can be compared given the evaluation function:

$$x \succeq_\omega y \iff f_\omega(x) \geq f_\omega(y). \quad (1)$$

A first widely used model that we will consider here is the weighted sum (WS) model. Given a vector of weights $\omega = \{\omega^1, \dots, \omega^q\} \in \mathbb{R}^q$:

$$f_\omega(x) = \sum_{i=1}^q \omega^i x^i, \quad (2)$$

with $\omega^i \geq 0$ and $\sum_i \omega^i = 1$.

Ordered weighted averaging (OWA) model is another famous model [Yager, 1988] we will investigate further on. OWA is different from WS, since the criteria values are ordered in increasing order. Given a vector of weights $\omega \in \mathbb{R}^q$ and the ordered criteria values $x^{(1)} \leq \dots \leq x^{(q)}$, we have:

$$f_\omega(x) = \sum_{i=1}^q \omega^i x^{(i)}, \quad (3)$$

with $\omega^i \geq 0$, $\sum_i \omega^i = 1$.

More complex models can be considered within the framework presented here, like Choquet integrals. As long as the models we pick are linear in ω , Equation (1) is equivalent to a linear constraint and we can use linear programming techniques at our advantage, as in [Benabbou et al., 2017].

Example 2 (Application of models). Given the sandwiches of Table 1, first suppose the user uses a WS model such that $\omega = (0.9, 0.1)$. This means she really values the flavour over the price. Between a duck and a fish sandwich, the user choose the duck, as it scores 9 against 7.6 for fish.

If we suppose the user evaluates each sandwich with an OWA model such that $\omega = (0.7, 0.3)$, meaning she penalises a sandwich with a low score in any category, she will prefer the fish sandwich: $f_\omega(\text{fish}) = 4 \times 0.7 + 8 \times 0.3 = 5.2$ while $f_\omega(\text{duck}) = 0 \times 0.7 + 10 \times 0.3 = 3$.

Example 2 assumes we know the model ω . In reality, we do not know such parameters, and directly asking the agent about them makes little sense. To find these parameters, we need a preference elicitation method, that we now present.

2.2 ROBUST ELICITATION

Motivation Finding a unique model ω from pairwise comparisons is a difficult task. However, it is often possible to draw inferences without complete information. In robust approaches, a subset Ω' of possible models ω is identified from preferential information, from which are inferred those preferences that hold for every model. This results in a partial preorder over \mathbb{X} where:

$$x \succeq^{\Omega'} y \iff \forall \omega \in \Omega' f_\omega(x) \geq f_\omega(y). \quad (4)$$

In order to make recommendations, an elicitation strategy that reduces Ω' as quickly as possible is needed. Such a strategy should also aim at making good recommendations even if $\succeq^{\Omega'}$ does not have a single maximal element, as in practice one may stop collecting information before that.

Regret based elicitation Minmax regret is a well-known notion when considering decision problems under uncertainty and set-valued information [Savage, 1951]. It still provides strong guarantees on the recommendation quality, while being less conservative than standard minmax.

We now introduce the elements of regret-based elicitation strategies, that we will extend in Section 3. The regret of choosing an alternative x over the alternative y according to a model ω is defined by:

$$R_\omega(x, y) = f_\omega(y) - f_\omega(x). \quad (5)$$

Given a set Ω' of models, the pairwise max regret is:

$$\text{PMR}(x, y, \Omega') = \max_{\omega \in \Omega'} R_\omega(x, y), \quad (6)$$

corresponding to the maximum regret of choosing x over y for any model $\omega \in \Omega'$. The max regret of choosing x is:

$$\text{MR}(x, \Omega') = \max_{y \in \mathbb{X}} \text{PMR}(x, y, \Omega'), \quad (7)$$

corresponding to the regret of choosing x in the worst case scenario. Lastly, the min max regret of a set \mathbb{X} of alternatives given a set Ω' of possible models is:

$$\text{mMR}(\Omega') = \min_{x \in \mathbb{X}} \text{MR}(x, \Omega'), \quad (8)$$

and $x^* = \arg \text{mMR}(\Omega')$ is the alternative that gives the minimal regret in a worst-case scenario, which would correspond to the current recommendation if no further information can be collected.

Example 3 (Initial choice). Supposing the user evaluates each sandwich with a WS model, we need to compute $\max_{\omega \in \Omega} R_\omega(x_i, y_j) \forall i, j \in \{\text{Cheese, Duck, Fish, Ham}\}$. We can compute the maximal regrets by optimising $\max_{\omega \in \Omega} (\omega \cdot (x - y))$. Since we optimise a linear function over a convex polytope Ω , the optimisation problem is solved easily and exactly using linear programming (LP).

We then obtain the PMR, as shown in Table 2:

x/y	y_C	y_D	y_F	y_H
x_C	0	5	3	2
x_D	9	0	4	7
x_F	5	2	0	3
x_H	2	3	1	0

Table 2: Initial PMR

For example, if a user chooses the duck sandwich over the fish sandwich, her maximal regret is $\text{PMR}(x_D, y_F) = 4$; for the model $\omega = (0, 1)$, we have $f_\omega(x_D) = 0$ and $f_\omega(x_F) = 4$.

x	x_C	x_D	x_F	x_H
MR	5	9	5	3

Table 3: Initial MR

The corresponding MR is given in Table 3. We obtain $\text{mMR} = 3$ and the best initial choice is $x^* = x_H$, which is the least regretted in the worst case scenario, when having no information on the DM preferences.

Elicitation sequence In our setting, we consider the quite common case where preferential information is collected through pairwise comparison: we present a pair (x, y) to the user, and she tells which one she prefers. We will denote by

$$\Omega_{x \succeq y} = \{\omega \in \Omega : f_\omega(x) \geq f_\omega(y)\}, \quad (9)$$

the subset of models consistent with the assessment $x \succeq y$, and $\Omega_{y \succeq x}$ the subset for $y \succeq x$.

An elicitation sequence corresponds to alternatively presenting a pair to the user, and updating the information with the answer. In the robust approach, if Ω^k is the possible subset of models at the k th step, the next step is to present a couple (x, y) to the user, and then compute $\Omega^{k+1} = \Omega^k \cap \Omega_{x \succeq y}$ if the user prefers x , $\Omega^{k+1} = \Omega^k \cap \Omega_{x \preceq y}$ otherwise.

Choosing a good pair (x, y) is therefore a critical step, and in this paper we consider the well-known CSS strategy [Boutilier et al., 2006], where given a subset Ω' , the user compares the current regret-based recommendation $x^* = \arg \text{mMR}(\Omega')$ (so our best option w.r.t this criterion) to its worst opponent $y^* = \arg \max_{y \in \mathbb{X}} \text{PMR}(x^*, y, \Omega')$. This heuristic strategy provides good results in general, and guarantees that the updated set will be non-empty.

Example 4 (Worst opponent of ham). In example 3, we already found that $x^* = x_H$. Given Table 2 and the CSS heuristic, $y^* = \arg \max_{y \in \mathbb{X}} \text{PMR}(x_H, y, \Omega') = y_D$. We create the query (x_H, y_D) and if the user does not commit mistakes, she will answer that she prefers duck to ham.

Decrease of regret A notable advantage of robust approaches combined with CSS is that, by construction, they guarantee that the elicitation sequence will converge, as we recall here.

Proposition 1. *Given $\Omega^{k+1} \subseteq \Omega^k$, the sets of possible model at steps k and $k + 1$, we have that:*

$$\text{PMR}(x, y, \Omega^k) \geq \text{PMR}(x, y, \Omega^{k+1}), \quad (10)$$

$$\text{MR}(x, \Omega^k) \geq \text{MR}(x, \Omega^{k+1}), \quad (11)$$

$$\text{mMR}(\Omega^k) \geq \text{mMR}(\Omega^{k+1}). \quad (12)$$

Proof. PMR. Suppose we have a function f and two sets Ω, Ω' such that $\Omega' \subseteq \Omega$. We have $\max_{x \in \Omega} f(x) \geq \max_{x \in \Omega'} f(x)$, the maximum of Ω being either in Ω' or in $\Omega \setminus \Omega'$. We can replace f by the PMR, Ω by Ω^k and Ω' by Ω^{k+1} since $\Omega^{k+1} \subseteq \Omega^k$. (10) is then proved.

Proof for MR and mMR directly follows, as they are maximum and minimum taken over decreasing values. \square

2.3 ELICITATION WITH UNCERTAINTY

Motivation The robust approach works quite well, provided two strong hypotheses are satisfied: the first one is that the user is an oracle (never makes mistakes), and the second one is that the chosen model family can perfectly describe the user preferences. However, such hypotheses are unrealistic in many applications, and it is desirable to account for possible mistakes through refined uncertainty modelling. As the next example shows, failure in those hypotheses can lead to unwarranted situations.

Example 5 (A single error to ruin everything). Assume the user preferences are modelled by a WS with $\omega = (0.9, 0.1)$. If the agent must choose between x_D and x_F , we have already shown in example 2 that the duck should have been taken ($f_\omega(x_D) = 9$, $f_\omega(x_F) = 7.6$). Let us imagine the user is unfocused, unsure, or that the WS is not the good model, and assess $x_F \succeq x_D$. Then we have Ω' the set of models consistent with her known preferences, such that $\Omega' = \Omega_{x_F \succeq x_D} = \{\omega \in \Omega : \sum_{i=1}^2 \omega^i \cdot (x_F^i - x_D^i) \geq 0\} = \{\omega \in \Omega : 2\omega^2 \geq \omega^1\}$.

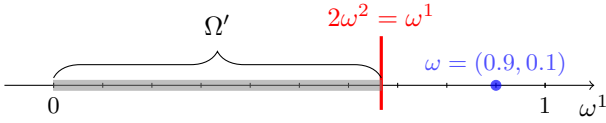


Figure 1: Wrong answer leading to a wrong model

As shown on Figure 1, where we only give the ω^1 value (since $\omega^1 + \omega^2 = 1$), the true model is definitely left out of Ω' . Whatever the next answers are, we cannot get to ω .

Probabilistic approaches by Bourdache et al. [2019] are one solution to this issue, yet they do not provide the same guarantees as the ones of Proposition 1. Also, they typically do not question the model family, as all probabilities are renormalised after updating: renormalisation prevents from identifying the presence of inconsistency, whatever its source is. One of our main motivation in this work is to let go strong hypotheses of the robust approach, while preserving its strong guarantees. For this, we will adopt a possibilistic approach.

Possibility theory basics Possibility theory [Denœux et al., 2020] is an uncertainty modelling theory that adds degree to classical, binary set-valued information. As such, it formally extends set-theory, a feature absent from probability theory that we will use at our advantage. It is also arguably the simplest uncertainty theory to do so, therefore being computationally easier to handle than, e.g., evidence theory, that has also been proposed as a way to extend robust incremental elicitation [Guillot and Destercke, 2019].

Possibility distributions $\pi : \Omega \rightarrow [0, 1]$ are the basic tools of possibility theory, which are said to be *consistent* if $\exists \omega \in \Omega$

such that $\pi(\omega) = 1$. Such a distribution induces two dual measures, the possibility and the necessity, over events $U \subseteq \Omega$ defined as:

$$\Pi(U) = \sup_{\omega \in U} \pi(\omega), \quad (13)$$

$$N(U) = 1 - \Pi(U^c) = \inf_{\omega \in U^c} (1 - \pi(\omega)). \quad (14)$$

There are two important notions issued from the theory that we will use. The first one is the α -cut of π that is defined as:

$$E_\pi^\alpha = \{\omega \in \Omega : \pi(\omega) \geq \alpha\}, \quad (15)$$

which is just the set of values having a possibility higher than α . The second one is the degree of inconsistency of π , defined as:

$$K = 1 - \max_{\omega \in \Omega} \pi(\omega), \quad (16)$$

which measures how far π is from being consistent.

Possibility theory extends set theory, in the sense that a standard set Ω' is modelled by the distribution $\pi_{\Omega'}(\omega) = \mathbb{I}_{\Omega'}(\omega)$ where \mathbb{I}_A is the indicator function of A . The notion of inclusion also has a straightforward extension, as two distributions π, π' will be denoted $\pi \subseteq \pi'$ iff $\pi \leq \pi'$. This is equivalent to have all α -cuts included in each others, as:

$$\pi \subseteq \pi' \iff E_\pi^\alpha \subseteq E_{\pi'}^\alpha, \quad \forall \alpha \in [0, 1]. \quad (17)$$

Possibilistic preferential information We will use specific distributions to model the uncertain preferential information provided by a user. More specifically, we will link to any assessment of the kind $x \succeq y$ a degree $\alpha \in [0, 1]$ of certainty or necessity (in the sense of Eq. 14) that the assessment holds. Such assessments will be denoted $x \succeq_\alpha y$, and we will link it with the largest (in the sense of inclusion) distribution such that $N(\Omega_{x \succeq y}) = \alpha$. This distribution, denoted as $\pi_{x \succeq_\alpha y}$, is such that:

$$\pi_{x \succeq_\alpha y}(\omega) = \begin{cases} 1 & \text{if } \omega \in \Omega_{x \succeq y}, \\ 1 - \alpha & \text{if } \omega \in \Omega_{x \prec y}. \end{cases} \quad (18)$$

The degree α can either be provided directly by the user (possibly using a linguistic finite scale to transform into numbers), or be directly specified by the analyst, allowing uncertainty in the process without requiring additional cognitive efforts from the user. In practice, such an alpha would just be used as an artefact, possibly with a fixed value and a simple T-norm like T_{\min} , allowing contradiction between questions, making the detection of inconsistencies possible.

Example 6 (An error no longer ruins everything). Let us continue example 5, but this time the user says she is not totally certain that she prefers fish to duck. We suppose this corresponds to a confidence level $\alpha = 0.7$. We then have $\Pi(\Omega') = 1$ and $\Pi(\Omega \setminus \Omega') = 0.3$: there is still a small possibility that the correct model belongs to $\Omega \setminus \Omega'$, making it possible to find ω . Figure 2 pictures distribution $\pi_{x_F \succeq_{0.7} x_D}$.

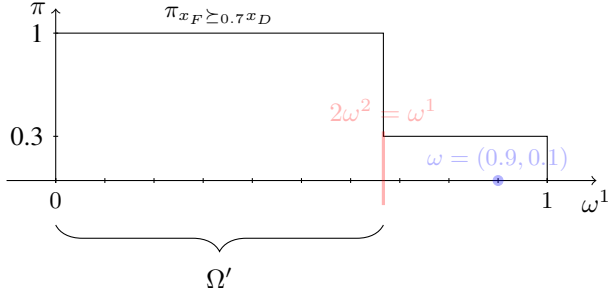


Figure 2: Possibilistic preferential information

Note that the two extreme values $\alpha = 1$ and $\alpha = 0$ for $x \succeq_\alpha y$ respectively correspond to specify a subset as in the robust approach (hence formally extending it), and to provide no information at all.

Combining possibilistic information An important aspect of the elicitation procedure is how to combine two preferential information. If we consider this as an information fusion problem, there are literally an infinity of ways to do that [Dubois et al., 2016]. In our case, since we want to extend the robust approach that use set intersection, we will adopt the possibilistic tools that extend such an operation: T-norms.

A T-norm is a function $T : [0, 1] \times [0, 1] \rightarrow [0, 1]$ that is commutative, associative, increasing, and has 1 as an identity element. It is upper bounded by the minimum operation. Three T-norms of particular interest are the following:

- Minimum T-norm: $T_{\min}(a, b) = \min(a, b)$,
- Product T-norm: $T_{\text{prod}}(a, b) = a \cdot b$,
- Łukasiewicz t-norm: $T_{\text{Luk}}(a, b) = \max(0, a + b - 1)$,

where $T_{\min} \geq T_{\text{prod}} \geq T_{\text{Luk}}$, with T_{\min} and T_{Luk} corresponding to the Fréchet-Hoeffding bounds. T_{\min} or T_{Luk} each can be linked to an extreme dependence assumption within probability theory, and working within them through a sensitivity analysis then amounts to making no assumption at all. T_{prod} , on its side can be linked to an independence assumption. Since they are associative and commutative, their extension to n dimensions is straightforward.

Remark 1. T_{\min} being idempotent and applicable to ordinal scales results in some advantages and drawbacks. Purely ordinal, linguistic information can be used, which is attractive. Moreover, the distribution $\pi_{\min} = T_{\min}(\pi, \pi')$ will take the same values as π, π' , therefore not creating new levels. Keeping the same levels has some computational advantage, but that may lead to a reduced expressiveness and to drowning effects [Benferhat et al., 1993].

T-norms have two important features in our framework:

- If all combined distributions π_1, \dots, π_n are such that $\exists \omega$ with $\pi_i(\omega) = 1$, then $T(\pi_1(\omega), \dots, \pi_n(\omega)) = 1$. In particular, if all pieces of preferential information are consistent with at least one model of Ω , then the resulting distribution is normalised. Again, this contrasts with a probabilistic treatment.
- An observed sub-normalised result, i.e. $\max_{\omega \in \Omega} T(\pi_1(\omega), \dots, \pi_n(\omega)) < 1$ is a reliable indication that there are some inconsistencies between the delivered preferential information and the model family. The inconsistencies can come either from a bad model choice or from mistakes in the user answers. The measure K given by Eq (16) can then be used to assess how significant this inconsistency is.

Example 7 (Inconsistency illustrated). Let us pursue Example 6, and consider that when comparing x_H to x_D , the user now answers with a strong conviction that $x_H \preceq_{0.9} x_D$, with $\Omega_{x_H \preceq_{0.9} x_D} = \{\omega : \omega^1 \geq 0.7\}$. Applying the minimum T-norm T_{\min} , we get the distribution pictured in Figure 3. We see that the inconsistency $K = 0.7$ is quite severe, but that the true model $\omega = (0.9, 0.1)$ is now among the most plausible.

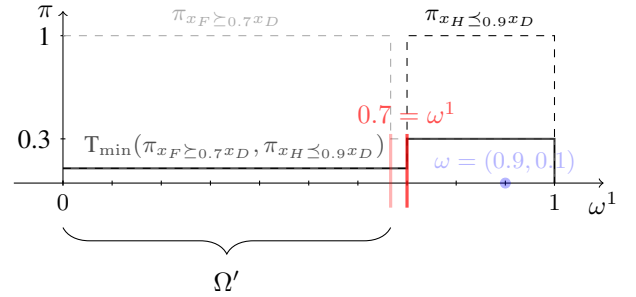


Figure 3: Inconsistency after merging information

3 EXTENSION

3.1 EXTENDING REGRET NOTIONS

We are now ready to extend the notions used in regret-based elicitation to our possibilistic framework. Since here Ω is a convex space, and preferential assessments of Eq. (1) induce linear constraints on Ω , those constraints will induce a partition P_1, \dots, P_L of Ω , over which elements π will be constant, i.e., $\pi(\omega) = \pi(\omega')$ for any $\omega, \omega' \in P_i$. In Figure 3, the partition would be $P_1 = [0, 2/3]$, $P_2 = [2/3, 0.7]$ and $P_3 = [0.7, 1]$ with π being constant on $P_1 \cup P_2$ due to the use of the minimum T-norm.

The elicitation framework remains the same: at a step k , we propose a pair (x, y) to the user who provides a possibilistic answer of the form $x \succeq_\alpha y$ resulting in a distribution $\pi_{x \succeq_\alpha y}$. We then combine the distribution (using a T-norm) with π^k ,

the possibility function over Ω obtained at the step k of the elicitation process with $\pi^0 = \mathbb{I}_\Omega$ being equivalent to the set Ω of all possible models. We denote P_1, \dots, P_{L_k} the corresponding partition.

Extending PMR Our extension of PMR, named EPMR, averages the PMR over the different α -cuts:

$$\text{EMR}(x, y, \pi^k) = \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \text{PMR}(x, y, E_{\pi^k}^{\alpha_i}), \quad (19)$$

where $1 = \alpha_1 > \dots > \alpha_n > \alpha_{n+1} = 0$ are the distinct values of π^k . If $\pi^k = \mathbb{I}_{\Omega^k}$, we retrieve the standard PMR of Eq. (6). Note that such an averaging is standard in possibilistic approaches (see, e.g., [Hüllermeier, 2014]).

If $E_{\pi^k}^{\alpha_1} = \emptyset$, we need to define $\text{PMR}(x, y, \emptyset)$. There are different options to do so [Guillot and Destercke, 2019]. Here we will consider $\text{PMR}(x, y, \emptyset) = 0$, to ensure convergence.

Extending MR We propose an extension of the MR, the EMR, averaging the MR on each focal set:

$$\text{EMR}(x, \pi^k) = \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \max_{y \in \mathbb{X}} \text{PMR}(x, y, E_{\pi^k}^{\alpha_i}), \quad (20)$$

corresponding to the average over cuts of the maximal pairwise regret. Again, if $\pi^k = \mathbb{I}_{\Omega^k}$, we retrieve the standard MR of Eq. (7).

Extending mMR We propose to extend the mMR with the mEMR:

$$\text{mEMR}(\pi^k) = \min_{x \in \mathbb{X}} \text{EMR}(x, \pi^k), \quad (21)$$

which, since $\text{EMR}(x, \pi^k)$ reduces to Eq. (7) when $\pi^k = \mathbb{I}_{\Omega^k}$, also reduces to (8) in the same case.

When π is normalised and in the absence of inconsistency, the necessity (N) and possibility (II) measures can be viewed as lower and upper bounds of unknown probabilities, respectively. This means that EPMR, EMR and mEMR can be interpreted as upper expectations over a set of probabilities. In absence of inconsistencies, mEMR remains an upper bound of the real regret. This connection is lost once errors happen.

3.2 EXTENDING CSS

PCSS We now propose the Possibility Current Solution Strategy (PCSS) in order to select questions, which extends the CSS one. The strategy is summarised in Algorithm 1. We assume that the user provides a unique choice when being presented with a pair, that is translated as a non-strict preference. Moreover, with polytopes and linear programs, whether including borders or not will have no consequences on the practical computations and numerical results.

Algorithm 1: PCSS algorithm

Data: Max Inconsistency Max_I , Threshold τ , Max number of queries Max_q , T-norm T

Result: $x^* = \arg \text{mEMR}(\pi^k)$

$k = 0, \pi^0 = \mathbb{I}_\Omega$;

while $\max \pi^k \geq 1 - Max_I, k \leq Max_q$ **and**

$mEMR \geq \tau$ **do**

 Compute $x^* = \arg \text{mEMR}(\pi^k)$;

 Compute $y^* = \arg \max_{y \in \mathbb{X}} \text{EMR}(x^*, y, \pi^k)$;

 User provide answer $x^* \succeq_\alpha y^*$ or $x^* \preceq_\alpha y^*$;

if User answer is $x^* \succeq_\alpha y^*$ **then**

$\pi^{k+1} = T(\pi^k, \pi_{x^* \succeq_\alpha y^*})$

else $\pi^{k+1} = T(\pi^k, \pi_{x^* \preceq_\alpha y^*})$;

$k = k + 1$;

end

Given that EPMR, EMR and mEMR all reduce to their robust counter-parts when $\alpha = 1$ for all answers, we do have that PCSS extend standard CSS.

As we said in the introduction, our goal is to retain the nice properties of the robust and CSS approaches when their assumptions hold (right model choice and correct answers), that is converging towards the right recommendation with high guarantees, while allowing for mistakes. We will now show that the first requirement is satisfied, while further experiments in Section 4 will demonstrate that we do detect inconsistencies and mistakes when those exist.

Firstly, we demonstrate that we retrieve convergence when no inconsistencies in the answers are observed.

Proposition 2 (Monotonicity). *Given π^k and π^{k-1} two successive possibility distributions issued from PCSS. Assuming both are normalised, then:*

1. $\text{EPMR}(x, y, \pi^k) \leq \text{EPMR}(x, y, \pi^{k-1})$,
2. $\text{EMR}(x, \pi^k) \leq \text{EMR}(x, \pi^{k-1})$,
3. $\text{mEMR}(\pi^k) \leq \text{mEMR}(\pi^{k-1})$.

Proof. Denote by $\mathcal{A}^k = \{\alpha_1, \dots, \alpha_m\}$ the ordered distinct values of π^k , and $\mathcal{A}^{k-1} = \{\alpha_1, \dots, \alpha_n\}$ those of π^{k-1} . Consider the ordered set $\mathcal{B} = \mathcal{A}^k \cup \mathcal{A}^{k-1} = \{\beta_1, \dots, \beta_\ell\}$:

$$\text{EPMR}(x, y, \pi^k) = \sum_{i=1}^{\ell} (\beta_i - \beta_{i+1}) \text{PMR}(x, y, E_{\pi^k}^{\beta_i}),$$

$$\text{EPMR}(x, y, \pi^{k-1}) = \sum_{i=1}^{\ell} (\beta_i - \beta_{i+1}) \text{PMR}(x, y, E_{\pi^{k-1}}^{\beta_i}),$$

$\pi^k \subseteq \pi^{k-1}$ with Eq. (17) implies $E_{\pi^k}^{\beta_i} \subseteq E_{\pi^{k-1}}^{\beta_i}$, thus $\text{PMR}(x, y, E_{\pi^k}^{\beta_i}) \leq \text{PMR}(x, y, E_{\pi^{k-1}}^{\beta_i})$, ending the proof.

Similar arguments can be used to prove the inequalities for EMR and mEMR.

□

We show that our proposal preserves the strong guarantees of the robust CSS, and is even more cautious in case of uncertain answers, as the corresponding regrets will be higher.

Proposition 3 (Strong guarantees). *Let π^k and Ω^k be the distribution and subset obtained through the same sequence of alternative pairs and answers, respectively with $\alpha < 1$ for all answers and with $\alpha = 1$ for all answers. Then:*

1. $PMR(x, y, \Omega^k) \leq EPMR(x, y, \pi^k)$,
2. $MR(x, \Omega^k) \leq EMR(x, \pi^k)$,
3. $mMR(\Omega^k) \leq mEMR(\pi^k)$.

Proof. We observe that $\pi_{x \succeq_{\alpha} y} \subseteq \pi_{x \succeq_{\beta} y}$ whenever $\alpha \geq \beta$. This means that all answers with $\alpha = 1$ are included in the answers with $\alpha < 1$.

We therefore have that $\pi_{\Omega^k} = \mathbb{I}_{\Omega^k} \subseteq \pi^k$. From this, we can simply reproduce the proof of Proposition 2.

□

On computational complexity While our approach is computationally more costly than the robust one, as it has to do with partition P_1, \dots, P_L of polytopes, the increase of computational cost is limited through several factors:

- Computing the PMR for a given α -cut or subset Ω can be done by using $PMR(x, y, \Omega) = \max_{P_i \in \Omega} PMR(x, y, P_i)$, possibly using computations done for P_i at previous steps.
- For any query, the constraint (1) will only cut in half a limited number of elements P_i . In fact, it can be shown that the number of new elements will be bounded above by a polynomial number [Schläfli and Wild, 2013, P 39]. We can readily see that in Figure 3, where a new constraint can only cut one of the interval.
- The fact that α -cuts are nested (i.e., $E_{\pi}^{\alpha} \subseteq E_{\pi}^{\beta}$ if $\alpha > \beta$) means that computation done for one cut can be partially re-used for other cuts.

This means that we must perform the same computations than for the robust approach on a partition of polytopes, but whose size increases polynomially after each question. Exact computations are possible as long as q and the number of queries remain small. In other cases, using approximated computations is doable (e.g., ignoring α -cuts having very small weights).

Compared to the previously proposed evidential approach Guillot and Destercke [2019], our approach based on possibility theory is more efficient thanks to two factors:

- As we already said, α -cuts and thus focal sets are nested, reducing the computational complexity. In evidential theory, focal sets are not necessary nested.
- The merging rules in our possibilist approach, based on T-norms, create less focal sets over which the PMR must be computed, in opposition to the evidential approach. This is especially true when using T_{\min} , where the final number of focal sets is equivalent to the number of different α used.

4 EXPERIMENTS

We perform several experiments on synthetic data to study the behaviour of our method and to demonstrate its advantages with respect to a standard robust approach. In the experiments presented in our paper, we conducted 200 simulations. For each simulation, we generated randomly 30 alternatives x_i with 4 criteria according to a uniform distribution such that $x_i \in [0, 1]^4$. Moreover, the simulated error rate is always the same, which may be different from the used error rate in the method, as in practice one does not have access to the "true" error rate.

Influence of T and of α The first experiment illustrates the influence of the T-norm and of the confidence level on the convergence of our method. During this experiment, we first used a specific T-norm, and then a fixed confidence level α . For each simulation, a WS model with 4 criteria was generated randomly according to a Dirichlet distribution $\alpha_{dirich} = (1/4, \dots, 1/4)$, and we computed the EMR at each step, and divided it by the initial mEMR. The ratios were then averaged for each degree and T-norm. In this experiment, the user always delivers the right answer. Remember that $\alpha = 1$ is equivalent to robust CSS, and $\alpha = 0$ means no information is provided at all.

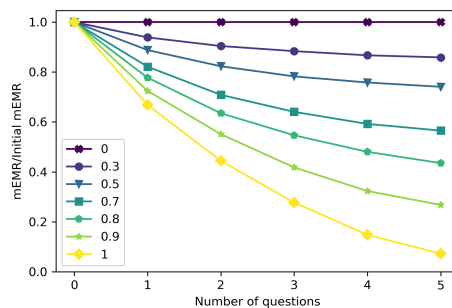


Figure 4: Decrease of computed mEMR with T_{\min}

On Figure 4 in which are given the results for T_{\min} , we see that high confidence leads to faster convergence, as expected. Results for other T-norms are similar. Figure 5 shows the results when the T-norm varies. As expected, the smaller the T-norm is, the faster π^k converges to low values for $\omega \notin$

$E_{\pi_k}^1$, and the faster the method converges. A bit surprising is the small gap between T_{prod} and T_{Luk} , suggesting that taking T_{prod} is sufficient, but that T_{min} could be overcautious. However, the quicker the convergence to small distribution values is, the lower the chances are to query models $\omega \notin E_{\pi_k}^1$ and to investigate potential inconsistencies.

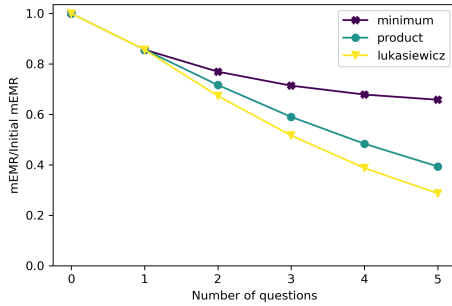


Figure 5: Decrease of computed mEMR with $\alpha = 0.6$

Handling of the inconsistency The second experiment highlights how our method is effective for managing inconsistency. In this experiment, when inconsistency is detected, the elicitation is stopped. The idea is to prove that our method is able to detect inconsistency and thus stops before it diverges too much from the optimal model. We did not try different ways to handle observed inconsistency, but it is topic of further research.

For each simulation, we generated a WS model, again with 4 criteria, according to a Dirichlet distribution with $\alpha_{\text{dirich}} = (1/4, \dots, 1/4)$, we picked the product T-norm and a fixed confidence level. We model the inconsistency of the user by supposing she answers correctly 70% of the time, and gives a random answer 30% of the time, meaning a true $\alpha = 0.7$. Each simulation ends when an inconsistency is detected, when 15 questions were answered, or when mEMR = 0. We determine $x^* = \arg \text{mEMR}$ and compute the regret of choosing it over the best alternative y_ω given the true model ω . The real regret and the inconsistency are then averaged for each degree, and we compute a confidence interval of 95% for the real regret.

Figure 6 shows that inconsistency in answers is quickly detected (sometimes as early as the third question), confirming that our framework is practically relevant to investigate potential user mistakes, without the need to change the elicitation strategy to actively look for mistakes. Without surprises, higher values of α tend to create higher inconsistencies when they appear.

We observe on Figure 7 and on table 4 that, with respect to real regret, the classic robust elicitation ($\alpha = 1$) performed the worst. This behaviour was expected, as there are no ways for detecting inconsistency, as shown on Figure 6. At each step the computed mEMR decreases, but the real regret

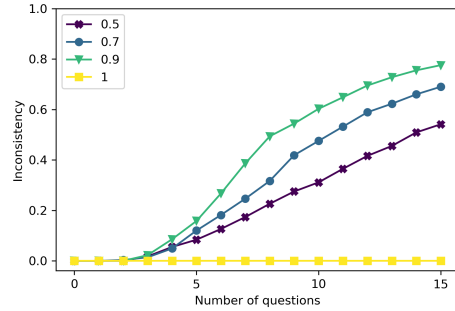


Figure 6: Detection of the inconsistency with different α

can increase, as we are getting further away from the true model. We also observe the lower the level of confidence is, the lower the real regret is, as the more cautious the PCSS is. In any case, Figure 7 shows that when we have uncertain answers, our approach will outperform the robust one.

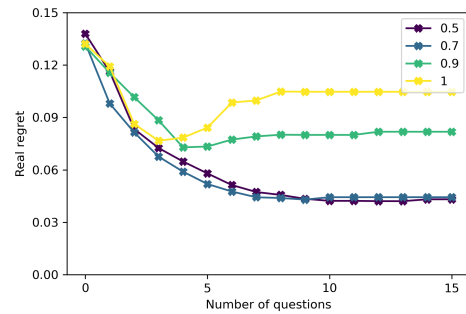


Figure 7: Real regret for different α

	Lower	Mean	Upper
$\alpha = 0.5$.032	.043	.055
$\alpha = 0.7$.030	.044	.058
$\alpha = 0.9$.063	.081	.100
$\alpha = 1$.082	.105	.127

Table 4: Real regret given a confidence level α

Detection of wrong models The third experiment stresses that we can detect a wrong preference model, especially if it is very distant from the true model. For each simulation, the wrong preference model is assumed to be a WS model with 4 criteria. We also used the product T-norm and a fixed confidence level $\alpha = 0.7$. We model the inconsistency of the agent by supposing she answers correctly 70% of the time, and gives a random answer 30% of the time. Each simulation ends when 15 questions were answered, or when mEMR = 0. The inconsistency is averaged over the different simulations. We tested three types of model, and incorrectly searched the true model in WS family:

- Random models: the user always answers randomly.
- Fairness-orientated OWA models: weights were generated randomly by a Dirichlet distribution with hyperparameter $\alpha_{dirich} = (0.85, 0.05, 0.05, 0.05)$.
- Random OWA models: weights were generated randomly by a Dirichlet distribution with hyperparameter $\alpha_{dirich} = (0.25, 0.25, 0.25, 0.25)$. Such models are close to WS models, as they will most of the time be close to arithmetic average.

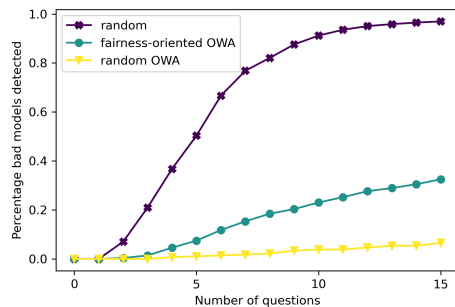


Figure 8: Detection of wrong models with $\alpha = 0.7$

Figure 8 shows that when the true model is very far from the postulated one (random and Fairness-orientated OWA), we can effectively detect it. When the models behave similarly (random OWA), such a detection can take longer.

5 CONCLUSION

We have introduced a possibilistic extension of robust incremental elicitation of preferences, demonstrating that such a tool can effectively solve existing issues of the robust approach, while remaining consistent with it. Experiments on simulated data sets show that we have better performances when user assessments can be wrong.

Our future works will focus on two issues:

- Providing means to repair the observed inconsistency, either by removing provided information or by selecting different model spaces.
- Extending the current approach to other models that have very different structures, such as lexicographic ones [Booth et al., 2010] that are not even based on numerical evaluation (thus requiring to adapt the notion of regret, or to use a different one to pick the query and recommendation).

Acknowledgements

This work benefited from the financial support of the project PreServe ANR-18-CE23-0008 of the French National Research Agency (ANR).

We want to thank the anonymous reviewers, whose comments have helped us to significantly improve our manuscript quality.

References

- Nawal Benabbou, Patrice Perny, and Paolo Viappiani. Incremental elicitation of choquet capacities for multicriteria choice, ranking and sorting problems. *Artificial Intelligence*, 246:152–180, 2017.
- Salem Benferhat, Claudette Cayrol, Didier Dubois, Jérôme Lang, and Henri Prade. Inconsistency management and prioritized syntax-based entailment. In *IJCAI*, volume 93, pages 640–645, 1993.
- Richard Booth, Yann Chevaleyre, Jérôme Lang, Jérôme Mengin, and Chattrakul Sombatheera. Learning conditionally lexicographic preference relations. In *ECAI*, volume 10, pages 269–274, 2010.
- Nadjet Bourdache and Patrice Perny. Anytime algorithms for adaptive robust optimization with owa and wowa. In *International Conference on Algorithmic Decision Theory*, pages 93–107. Springer, 2017.
- Nadjet Bourdache, Patrice Perny, and Olivier Spanjaard. Incremental elicitation of rank-dependent aggregation functions based on bayesian linear regression. In *IJCAI-19-Twenty-Eighth International Joint Conference on Artificial Intelligence*, pages 2023–2029. International Joint Conferences on Artificial Intelligence Organization, 2019.
- Craig Boutilier. Computational decision support: Regret-based models for optimization and preference elicitation. *Comparative Decision Making: Analysis and Support Across Disciplines and Applications*, pages 423–453, 2013.
- Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8-9):686–713, 2006.
- Thierry Denœux, Didier Dubois, and Henri Prade. Representations of uncertainty in ai: Beyond probability and possibility. In *A Guided Tour of Artificial Intelligence Research*, pages 119–150. Springer, 2020.
- Didier Dubois and Henri Prade. Possibility theory and data fusion in poorly informed environments. *Control Engineering Practice*, 2(5):811–823, 1994.
- Didier Dubois, Weiru Liu, Jianbing Ma, and Henri Prade. The basic principles of uncertain information fusion. an organised review of merging rules in different representation frameworks. *Information Fusion*, 32:12–39, 2016.

- Pierre-Louis Guillot and Sébastien Destercke. Preference elicitation with uncertainty: Extending regret based methods with belief functions. In *International Conference on Scalable Uncertainty Management*, pages 289–309. Springer, 2019.
- Eyke Hüllermeier. Learning from imprecise and fuzzy observations: Data disambiguation through generalized loss minimization. *International Journal of Approximate Reasoning*, 55(7):1519–1534, 2014.
- Leonard J Savage. The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67, 1951.
- Ludwig Schläfli and H Wild. *Theorie der vielfachen Kontinuität*, volume 38. Springer-Verlag, 2013.
- Ivan Vendrov, Tyler Lu, Qingqing Huang, and Craig Boutilier. Gradient-based optimization for bayesian preference elicitation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10292–10301, 2020.
- Ronald R Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on systems, Man, and Cybernetics*, 18(1): 183–190, 1988.