



Handling inconsistency in (numerical) preferences using possibility theory

Loïc Adam, Sébastien Destercke

► To cite this version:

Loïc Adam, Sébastien Destercke. Handling inconsistency in (numerical) preferences using possibility theory. Information Fusion, 2024, 103, pp.102089. 10.1016/j.inffus.2023.102089 . hal-04335693

HAL Id: hal-04335693

<https://hal.science/hal-04335693>

Submitted on 11 Dec 2023

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.

Handling inconsistency in (numerical) preferences using possibility theory

Loïc Adam^a, Sebastien Destercke^{a,*}

^aUMR CNRS 7253 Heudiasyc, Sorbonne Université, Université de Technologie de Compiègne, 57 avenue de Landshut, Compiègne, CS 60319 - 60203, France

Abstract

In this paper, we address the issues of gathering the preferences of a user when they may be uncertain, and of handling the possible ensuing inconsistency. We suggest using possibility theory as a means of modelling this uncertainty and making inferences despite the possible presence of inconsistencies due to user errors. While some parts of our approach are general, we specifically apply it to the case of numerical models and show through synthetic experiments the potential benefits of our approach.

Keywords: Possibilities, Preference, Inconsistency

1. Introduction

Modelling the preferences of individuals or populations has been a long-standing topic in fields such as economics, statistics, artificial intelligence, and operational research. However, the challenges of dealing with uncertainty, inconsistency, and incomplete knowledge in preference modelling are still actively discussed, particularly in artificial intelligence [1, 2].

In this paper, we are interested in the problem of handling uncertainty and inconsistency when eliciting the preferences of a single user. This type of scenario is often encountered in multi-criteria decision aid or artificial intelligence, while other fields such as statistics, machine learning, or economics tend to focus on learning from populations of individuals rather than a single one. There are two main approaches for dealing with the uncertainty and inconsistency (with respect to the chosen preference model) present in the preferential information of the user:

- The first approach involves using robust, set-based methods, in which we consider sets of potential models and make inferences from these sets. Each new piece of preferential information narrows down the set of possible models, with the updated set representing our uncertainty about the optimal model. These approaches are similar to the concept of version space in machine learning, as noted by some authors [3].

*Corresponding author

Email addresses: `loic.adam@hds.utc.fr` (Loïc Adam),
`sebastien.destercke@hds.utc.fr` (Sebastien Destercke)

When the available information results in a set of possible models, different inference strategies can be used. One is to only consider those that are true for all possible models, resulting in so-called robust methods [4] that may sometimes deliver partial results in the form of multiple recommended alternatives. Another one is to still make precise inferences in the presence of uncertainty, usually by adopting a cautious attitude in the form of a Maximin or a Minimax regret method [5, 6, 7]. These approaches consider worst-case scenarios to gather new preferential information and provide inference tools (*e.g.*, to choose a recommended alternative) with strong guarantees. We will focus on this latter approach in our experiments, as comparing different approaches is easier when making precise inferences and recommendations.

In this context, inconsistency occurs when none of the models in the selected family can match the observed preferences of a user. It is an all-or-nothing binary signal, and common strategies for addressing inconsistency include removing or modifying some preferential information, either by minimizing an objective function such as the number of violated preferences [8, 9], or by interacting with the decision maker [10].

It is worth noting that elicitation approaches using the Minimax regret avoid inconsistencies by collecting preferential information that is necessarily consistent with the previous answers and the chosen model class (see for example [6]). However, this approach relies on very strong assumptions, and notably that the decision maker and analyst are assumed, respectively, to make no mistake in their answers or model selection.

- The second approach for handling inconsistencies involves finding a model that minimizes some form of average error. This can be achieved through optimization techniques such as least squares minimization [11] or margin maximization [12], or by using Bayesian probabilities to maintain a probability distribution over the possible models and update it as new information becomes available, and making inferences using the classical expectation operator [13, 14, 15].

While such approaches are effective in handling inconsistencies, they typically cannot provide the same level of strong guarantees as set-based approaches. They only offer statistical or expected guarantees, with the *proviso* that the uncertainty model used is correct. This is particularly constraining for probability models, as they often require a large amount of data or a strong inductive bias for specification and statistical tools for validation, which may be difficult to reconcile with learning individual preferences. For instance, a common assumption is that the probability of error is proportional to the distance between two alternatives [15], but this may not always be evident. For instance, it may be easier to compare similar alternatives (*e.g.*, those that differ only on two criteria [16]) than very different ones. Additionally, probabilistic approaches do not clearly distinguish between situations with and without inconsistencies, as probability masses must be shared among multiple competing and exclusive

hypotheses¹.

In this paper, we study a third approach, consisting in using possibility distributions to model uncertain preferential information, and possibility theory [17] for reasoning with the provided information. The interest of such an approach is that it can remain consistent with set-based approaches, as sets are a special instance of possibility distributions, while producing a non-binary quantification of the inconsistency. It also gives at our disposal various tools for dealing with such an inconsistency, these tools extending set-theoretic and logical operations such as conjunction and disjunction rather than expectation-based operators. Such possibilistic approaches therefore constitute an interesting extension of the set-based approaches, while offering a view on uncertainty handling that differs from probabilities (as none of the two theories that are possibility theory and probability theory subsume the other). In particular, we will consider the resolution of inconsistencies from an information fusion perspective [18], as inconsistency handling is a typical problem of such approaches. Although the idea of applying possibilistic approaches to the modelling of preferences is not new [19], our contribution enriches this approach by considering methods issued from information fusion to deal with inconsistency, and by verifying through synthetic experiments that these methodological proposals have some empirical interest.

In Section 2, we describe our general possibilistic setting: how are modelled preferences, how can inconsistency be measured, how can it be dealt with. Section 3 then deals with possible inconsistencies in the observed preferential information, and how to handle them. Finally, Section 4 provides some experiments demonstrating the potential interest of our approach. Note that although we focus on numerical models associating a numerical value to each alternative, most parts of this paper can be readily applied to non-numerical models, especially Sections 2 and 3.

All along the paper, we will illustrate our approach through a simple running example using weighted averages as models.

2. Possibilistic modelling of preference models

In this section, we describe our general setting and illustrate it on a simple example.

2.1. Preferences and preference models

We consider that we want to describe the preferences of a user between various multi-criteria alternatives. The space of alternatives is a Cartesian product $\mathbb{X} = \prod_{i=1}^M \mathcal{X}_i$ where \mathcal{X}_i is the domain of values that the i th criterion can take. Such a domain can be discrete or continuous, and describes an aspect of the alternative. In practical multi-criteria decision problems, only a finite subset $\mathbb{A} \subset \mathbb{X}$ of all possible alternatives is considered.

We also assume that the preferences of the user can be described by some model $\omega \in \Omega$, where the set of possible models Ω is chosen by the analyst. Each

¹We illustrate this point at the end of Section 2

model Ω then induces a partial pre-order² over the set of alternatives. There are many such possible models, and we refer to the work of Pigozzi *et al.* [20] for a nice review.

In this work, we focus on numerical models, where $\omega : x \rightarrow \mathbb{R}$ is a real-valued function that maps any alternative $x \in \mathbb{X}$ to a corresponding value $\omega(x)$. For easiness, we also denote by $x \succeq_{\omega} y$ the fact that $\omega(x) \geq \omega(y)$. However, many of the ideas presented in this paper also apply to the case where ω is not a numerical model.

Example 1. *We suppose that a customer wants to buy a piece of cheese and wants to be sure she is making the best choice. For simplicity, we consider that she evaluates the score of a cheese through two criteria: the richness of the flavour and the price. The available cheeses are presented in Table 1.*

It is important to note that no alternatives are objectively worse than any other (i.e. all are Pareto undominated): American cheese may have no flavour, but it is the least expensive. On the other hand, truffle Brie is overpriced, but has the richest flavour. The other cheeses are tradeoffs between the two criteria.

Table 1: Set of alternatives \mathbb{X} and their scores, with $\omega = (0.6, 0.4)$

Name	Flavour	1/Price	Score
American cheddar	0	10	4
Emmental	4	6	4.8
Edam	5	5	5
Mozzarella	7	3	5.4
Truffle Brie	8	1	5.2

If we consider that the preferences of the user (the customer) are described by a model ω , a weighted sum with parameters $(0.6, 0.4)$, we obtain the scores presented in Table 1. Mozzarella is the preferred alternative for this user, as $\forall x \in \mathbb{A}, \omega(MO) \geq \omega(x)$.

2.2. Possibility theory reminder

A possibility distribution π over a space Ω is simply a mapping $\pi : \Omega \rightarrow [0, 1]$ where $\pi(\omega)$ measures how much the element³ ω is plausible. A distribution π is said *consistent* if $\max_{\omega \in \Omega} \pi(\omega) = 1$, that is, if at least one element is fully plausible. From a distribution π , one can then define two measures for any event or subset $A \subseteq \Omega$, called possibility and necessity measures, defined as:

$$\Pi(A) = \sup_{x \in A} \pi(x), \quad (1)$$

$$N(A) = 1 - \Pi(A^c) = 1 - \sup_{x \notin A} \pi(x). \quad (2)$$

Since Π and N are dual, working with one of them for every event A is sufficient. Possibility theory formally extends sets, in the sense that the information given by specifying a subset E is modelled by the distribution $\pi(x) = 1$ if $x \in E$,

²A transitive, antisymmetric relation on $\mathbb{X} \times \mathbb{X}$.

³In our case, a preference model.

zero otherwise, in which case $N(A) = 1$ for any A such that $E \subseteq A$, and zero otherwise. When π is consistent, the two bounds $[N(A), \Pi(A)]$ can be interpreted as probabilistic bounds, inducing in this case the probabilistic set

$$\mathcal{P} = \{P | N(A) \leq P(A) \leq \Pi(A), \forall A \subseteq \Omega\}.$$

135 Another important notion in possibility theory is the alpha-cut. Given a possi-
136 bility distribution π , its alpha-cut π_α is the subset

$$\pi_\alpha = \{\omega \in \Omega : \pi(\omega) \geq \alpha\} \quad (3)$$

137 which includes all the elements of Ω having a possibility degree higher than α .

138 Given a real-valued function $f : \Omega \rightarrow \mathbb{R}$, other notions used in this paper
139 are the one of lower and upper expectations $\underline{\mathbb{E}}_\pi(f), \bar{\mathbb{E}}_\pi(f)$ of f induced by a
140 distribution π . It corresponds to the lower and upper expected values of this
141 function over the set \mathcal{P} , and can be computed as

$$\underline{\mathbb{E}}_\pi(f) = \int_0^1 \min_{\omega \in \pi_\alpha} f(\omega) \, d\alpha, \quad (4)$$

142

$$\bar{\mathbb{E}}_\pi(f) = \int_0^1 \max_{\omega \in \pi_\alpha} f(\omega) \, d\alpha. \quad (5)$$

143 When the distribution π takes a finite number of distinct values $1 = \alpha_1 > \dots >$
144 $\alpha_n > \alpha_{n+1} = 0$ (as will be our case here), Equations (4)-(5) become:

$$\underline{\mathbb{E}}_\pi(f) = \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \min_{\omega \in \pi_{\alpha_i}} f(\omega), \quad (6)$$

145

$$\bar{\mathbb{E}}_\pi(f) = \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \max_{\omega \in \pi_{\alpha_i}} f(\omega). \quad (7)$$

146 In the rest of the paper, we will also consider unnormalized possibility distribu-
147 tions π such that $\max_{\omega \in \Omega} \pi(\omega) < 1$, in which case the value

$$\text{Inc}(\pi) = 1 - \max_{\omega \in \Omega} \pi(\omega) \quad (8)$$

148 will quantify the inconsistency of the available information.

149 2.3. Possibilistic preferential information

150 In the following, we consider elementary pieces of information taking the form
151 of (E, α) , where $E \subseteq \Omega$ is a subset of possible models and α is understood as
152 the certainty degree that the assertion E is true. It is interpreted as $N(E) \geq \alpha$,
153 to which we can associate a corresponding possibility distribution $\pi_{(E, \alpha)}$ that is
154 the least informative satisfying $N(E) \geq \alpha$. This distribution is:

$$\pi_{(E, \alpha)}(\omega) = \begin{cases} 1 & \text{if } \omega \in E, \\ 1 - \alpha & \text{otherwise.} \end{cases} \quad (9)$$

155 In particular, $\alpha = 1$ corresponds to a set-valued information where we are
156 certain that $\omega \in E$, while $\alpha = 0$ amounts to a void statement corresponding to

157 ignorance. Equation (9) can be interpreted as an item of information within
 158 possibilistic logic [21], and most reasoning tools used in this paper could be
 159 interpreted through the lens of such a logic⁴. As illustrates the next example, E
 160 will typically correspond to a subset of possible models resulting from an answer
 161 provided by the user.

162 **Example 2** (Piece of information). *Given Example 1 and Table 1, assuming*
 163 *that the user declares $TB \succeq_\omega EM$ with a certainty degree $\alpha_1 = 0.8$, we obtain the*
 164 *following decision frontier:*

$$\omega(TB) \geq \omega(EM) \Rightarrow 8\omega^1 + \omega^2 \geq 4\omega^1 + 6\omega^2 \Rightarrow 4\omega^1 \geq 5\omega^2, \quad (10)$$

corresponding to the information pictured in Figure 1




Figure 1: Preferential information $\pi_{(E_1, \alpha_1)}(\omega)$ of Example 2

165

166 In this paper, we will consider that a set E_i is the result of some pairwise
 167 comparison between a pair of alternatives $(x, y) \in \mathbb{X}^2$, where the user can either
 168 state $x \succeq y$ or $y \succeq x$. We denote by $E_{x \succeq y}$ and $E_{y \succeq x}$ the subsets of Ω resulting
 169 from each possible answer.

170 In practice, we collect multiple pieces of information (E_i, α_i) , $i = 1, \dots, n$
 171 during the elicitation process, each of them corresponding to a distribution
 172 $\pi_{(E_i, \alpha_i)}$. Note that those possible answers E_i will define a finite partition
 173 $\{\Omega_1, \dots, \Omega_P\}$ of Ω where Ω_i is of the kind $\cap_{\phi_i \in \{E_i, E_i^c\}} \phi_i$, where each element of
 174 the partition correspond to non-empty intersections of elementary answers or
 175 their complements. The distributions $\pi_{(E_i, \alpha_i)}$ can then be combined or fused
 176 together in a single distribution by extending classical set operators such as
 177 conjunction (logical AND) and disjunction (logical OR). The use of such operators
 178 also allows for an easier interpretation of the performed operations [23, 24, 18].
 179 In particular, if we have no reasons to think that the pieces of information
 180 $\pi_{(E_i, \alpha_i)}$ are unreliable⁵, the most sensible way to combine them is through
 181 conjunction, which in possibility theory is typically done through the use of a
 182 T-norm operator [25]. As our goal here is not to discuss the pros and cons of
 183 the different T-norms, we will focus on the product T-norm, resulting in the

⁴This should not be confused with the idea of using possibilistic logic to represent preferences [22], in which degrees α_i represent intensities of preferences and not uncertainty quantification.

⁵We will deal with this situation in Section 3.

184 distribution π_{\cap} such that:

$$\pi_{\cap}(\omega) = \prod_{i=1}^n \pi_{(E_i, \alpha_i)}(\omega). \quad (11)$$

Example 3 (Fusion of information). *Now we consider two pieces of information: $\pi_{(E_1, \alpha_1)}$ determined from the answer of the user in Example 2, and a new one denoted $\pi_{(E_2, \alpha_2)}$. We suppose the user now declares that $MO \succeq_{\omega} TB$ with a certainty degree $\alpha_2 = 0.6$. We obtain a new decision frontier: $\omega^1 \leq 2\omega^2$. The new piece of information is shown on Figure 2. We then apply a T-norm between the two pieces of information, obtaining:*

$$\pi_{\cap}(\omega) = \begin{cases} T(1 - \alpha_1, 1) = 0.2 & \text{if } \omega^1 < 5/9, \\ T(1, 1) = 1 & \text{if } 5/9 \leq \omega^1 \leq 2/3, \\ T(1, 1 - \alpha_2) = 0.4 & \text{if } \omega^1 > 2/3, \end{cases}$$

which is shown on Figure 3.




Figure 2: Preferential information $\pi_{(E_2, \alpha_2)}(\omega)$ of Example 3

185




Figure 3: Fusion π_{\cap} of two pieces of preferential information of Example 3 and resulting partition of Ω

186 We can now briefly illustrate the notions of lower and upper expectations,
187 using our previous example:

188 **Example 4** (Expectation bounds). *Let us now consider the possibility distribu-*
189 *tion obtained in Example 3, as well as the function $f(\omega) = \omega(MO) = 7\omega^1 + 3\omega^2$.*




Figure 4: Overlap of the possibility distribution π_\cap from Example 3 with the score of the alternative MO

Both the possibility distribution and the function are superimposed on Figure 4. We have three distinct values for π_\cap : 1, 0.4 and 0.2. We thus have three alpha-cuts: $\pi_1 = [5/9, 2/3]$, $\pi_{0.4} = [5/9, 1]$ and $\pi_{0.2} = [0, 1]$.

We can now determine the lower expectation:

$$\begin{aligned}\mathbb{E}_{\pi_\cap}[\omega(MO)] &= (1 - 0.4) \min_{\omega \in [5/9, 2/3]} \omega(MO) + (0.4 - 0.2) \min_{\omega \in [5/9, 1]} \omega(MO) \\ &\quad + (0.2 - 0) \min_{\omega \in [0, 1]} \omega(MO) \\ &= 0.6(7 \times 5/9 + 3 \times 4/9) + 0.2(7 \times 5/9 + 3 \times 4/9) + 0.2(3 \times 1) \\ &\approx 4.77.\end{aligned}$$

And for the upper expectation:

$$\begin{aligned}\overline{\mathbb{E}}_{\pi_\cap}[\omega(MO)] &= (1 - 0.4) \max_{\omega \in [5/9, 2/3]} \omega(MO) + (0.4 - 0.2) \max_{\omega \in [5/9, 1]} \omega(MO) \\ &\quad + (0.2 - 0) \max_{\omega \in [0, 1]} \omega(MO) \\ &= 0.6(7 \times 2/3 + 3 \times 1/3) + 0.2(7 \times 1) + 0.2(7 \times 1) \\ &= 6.2.\end{aligned}$$

193

2.4. Errors in set-wise and possibilistic approaches

As said in the introduction, set-wise approaches are quite useful and can provide strong guarantees regarding inferences as long as the information provided by the user is correct.

However, such hypotheses are often unrealistic in many applications, and it may be desirable to account for possible mistakes (in the user responses or in the analyst choice) through uncertainty modelling. As the next example shows, failure in those hypotheses can lead to unwarranted situations.

Example 5 (A single error to ruin everything). We take Example 2 with two small but important modifications: we do not consider a possibilist information, and thus only E_1 is considered (equivalent to $\alpha_1 = 1$); and the user is either unfocused or unsure and makes the erroneous claim that $TB \preceq EM$. We determine that $E_1 = E_{TB \preceq EM}$ is now $\{\omega \in \Omega : 4\omega^1 \leq 5\omega^2\}$, corresponding to information shown on Figure 5. As we can see, the true model is definitely left out of E_1 . Whatever the next answers are, we cannot get to ω^* .




Figure 5: Wrong answer leading to a wrong model in Example 5

209 As we have already shown elsewhere [26], using possibility theory is an
 210 interesting way to solve this issue, as illustrates the next example.

211 **Example 6** (An error no longer ruins everything). *Let us continue with Example*
 212 *5, but this time the user is providing a certainty degree $\alpha_1 = 0.7$ with her wrong*
 213 *answer. We obtain a possibilist information shown on Figure 6 such that it*
 214 *is possible to find ω^* with further questions, supposing the user stops making*
 215 *erroneous claims.*




Figure 6: Possibilistic preferential information in Example 6

216 By adding uncertainty to the user opinions, we can weaken the assumption
 217 of a set-based approach and only consider that some models are less plausible
 218 than the others, while remaining consistent with a set-based approach, that is
 219 retrieved when the specified level of certainty α is 1 (in which case some models
 220 are considered impossible). This is in contrast with probabilistic approaches,
 221 where increasing the plausibility of some models necessarily means decreasing
 222 the plausibility of others. For instance, in a probabilistic setting it would be
 223 impossible to have Figure 1, or to increase the degree α_1 on E_1^c without decreasing
 224 the plausibility on E_1 . Such a modelling is clearly less preferable under our single
 225 user and version space assumptions, where we assume that there is a unique true
 226 model, that should be able to remain fully plausible in theory.

227 Moreover, as we pointed out in the introduction, using probabilistic ap-
 228 proaches is not desirable in our context: while they may give similar or better
 229 recommendations than possibilistic approaches in terms of pure numerical per-
 230 formances (this is investigated in Section 4.5), they cannot be used to reliably
 231 detect some form of inconsistency, which is a problem if we want to repair it,
 232 like in Subsection 3.2, or if we are interested in detecting and analysing this
 233 inconsistency. As Example 7 illustrates, a probabilistic approach will inevitably
 234 generate some conflict ($P(\emptyset) > 0$) even in the case of coherent answers. This

shows that probabilistic approaches, while potentially showing good numerical results, are not well tailored for detecting inconsistencies and for analysing them.

Example 7 (Conflict with probabilistic information). *We take Example 3, where answers were consistent with the subset of models $5/9 \leq \omega^1 \leq 2/3$. We interpret the different pieces of preferential information as probabilities, with $P(E_i) = \alpha_i$ and $P(E_i^c) = 1 - \alpha_i$. This means that the first piece of information gives $P(\omega^1 < 5/9) = 0.2$ and $P(\omega^1 \geq 5/9) = 0.8$, while the second piece of information gives $P(\omega^1 \leq 2/3) = 0.6$ and $P(\omega^1 > 2/3) = 0.4$. As such, and supposing both pieces of information are independent, we obtain the following distribution:*

$$P(\omega^1) = \begin{cases} 0.12 & \text{for } 0 \leq \omega^1 < 5/9, \\ 0.48 & \text{for } 5/9 \leq \omega^1 \leq 2/3, \\ 0.32 & \text{for } 2/3 < \omega^1 \leq 1, \\ 0.08 & \text{for } \omega^1 < 5/9 \text{ and } 2/3 > \omega^1, \end{cases}$$

with the last condition leading to $P(\emptyset) = 0.08$, despite the fact that both answers were consistent together and with our model space. Therefore, while we could ignore this conflict and decide to normalise the distribution in order to have a probability distribution, probability theory cannot be used to reliably detect inconsistency.

3. Handling inconsistencies

As said in the previous section, one interest of the possibilistic approach is that it can model uncertainty in the user replies, and will avoid completely discarding the good model in case of error. In case where some answers are inconsistent between them, possibility distributions also quantify inconsistency gradually, rather than having an all-or-nothing information as set-based approaches do.

There are mainly two reasons for inconsistencies to be observed in the preferential information collection process and the inferences that ensue:

- **Model error** that arises when the choice of the space Ω is unable to account for given preferential information, even when those are all correct. When such errors happen, a possible strategy is to change or broaden the model space (*e.g.*, switching from weighted averages to k-monotone Choquet integrals [27]) so as to reduce inconsistencies;
- **User error** that comes from the user who committed some errors in the past when formulating her preferences, hence making the answers jointly incompatible with any possible model of Ω . Note that such errors can have different origins: the user was unsure of her answers, or maybe changed her viewpoint during the elicitation process. When this happens, a common strategy is to remove some answers of the user in some minimal way, making the remaining ones consistent [8].

Differentiating between the two types of errors without having interactions with the user or without additional information is in itself a challenging problem, which is not the main topic of this paper, even if information fusion tools can provide interesting answers to this question [28]. We therefore assume here that

errors originate from the user, and consider possible strategies to deal with such errors when considering possibility theory and associated information fusion tools. Concretely, we look at the case where π_\cap is subnormalized, *i.e.*, $\text{Inc}(\pi_\cap) > 0$.

We will detail two strategies: not questioning the conjunctive merging and adapting our inference tools based on Equations (4) and (5) ; or modifying the way we merge information, thus changing our assumptions about the sources of information. The first tool is an easy fix but does not provide much information about the source of inconsistencies, while the second is more involved but provides some analysis along with the fix. It should also be noted that both strategies are not incompatible, using for instance the first to have quick inference, and the second to make a final analysis, as studied in Section 4.5.

3.1. Inferring despite inconsistencies

Having a positive inconsistency $\text{Inc}(\pi_\cap) > 0$ implies that $E_{\pi_\cap}^1 = \emptyset$. This means that if one wants to make inferences over a given function $f(\omega)$ in Equations (6)-(7) without correcting inconsistencies, we need to define minima and maxima over the empty set. While it is possible to define virtually an infinity of strategies to account for that, we only present here two classical solutions, which differences are illustrated on Example 8:

- First way: consider that $\min_{\omega \in \emptyset} f(\omega) = \max_{\omega \in \emptyset} f(\omega) = 0$. This simply amounts to ignoring the inconsistent information. This is to some extent similar to inference procedures in possibilistic logic in presence of inconsistency [21]. One possible advantage of such an approach is that if $\pi' \subseteq \pi$, then $[\underline{\mathbb{E}}_{\pi'}(f(\omega)), \overline{\mathbb{E}}_{\pi'}(f(\omega))] \subseteq [\underline{\mathbb{E}}_\pi(f(\omega)), \overline{\mathbb{E}}_\pi(f(\omega))]$, keeping a certain monotonicity with respect to information gain, as more precise possibility distributions will lead to more precise inferences ;
- Second way: consider that $\min_{\omega \in \emptyset} f(\omega) = \min_{\omega \in \Omega} f(\omega)$ and $\max_{\omega \in \emptyset} f(\omega) = \max_{\omega \in \Omega} f(\omega)$. This amounts to transforming conflict into ignorance, and to have a very conservative view about it. It can also be viewed as normalizing the possibility distributions by taking $\pi' = \pi + \text{Inc}(\pi_\cap)$.

This way of resolving inconsistencies does not change our hypothesis with respect to the previously given answers and information (they are not modified), nor how we combine them (conjunctively). Note that this approach somehow avoids searching for the sources of inconsistency, and either ignores it or turn it into ignorance (a different concept than inconsistency). Therefore, such strategies appear legitimate only when inconsistency and its effects are likely to be limited, and when there is no need to analyse the details of the conflicting situation.

Example 8 (Inferring despite inconsistencies without modifying preferential information). *Let us take again Example 4, but this time the user is giving two incorrect preferential information: $EM \succeq_\omega TB$ with $\alpha_1 = 0.8$ and $TB \succeq_\omega MO$ with α_2 . The fusion of the two pieces of information (using a product T-norm) is:*

$$\pi_\cap(\omega) = \begin{cases} T(1, 1 - \alpha_2) = 0.4 & \text{if } \omega^1 < 5/9, \\ T(1 - \alpha_1, 1 - \alpha_2) = 0.08 & \text{if } 5/9 \leq \omega^1 \leq 2/3, \\ T(1 - \alpha_1, 1) = 0.2 & \text{if } \omega^1 > 2/3, \end{cases}$$




Figure 7: Overlap of the unnormalized possibility distribution π_Ω from Example 8 with the score of the alternative MO

which is represented on Figure 7 along with the score of MO. As we can see, the preferential information is incoherent, with $\text{Inc}(\pi_\Omega) = 0.6$. We thus need to define minima and maxima over the empty set to infer.

π_Ω has three distinct values: 0.4, 0.2 and 0.08. We thus have four alpha-cuts: $\pi_{1\cap} = \emptyset$, $\pi_{0.4\cap} = [0, 5/9]$, $\pi_{0.2\cap} = [0, 5/9] \cup [2/3, 1]$ and $\pi_{0.08\cap} = [0, 1]$. We define the minimum and the maximum on the empty set according to the two solutions presented earlier:

- $\min_{\omega \in \emptyset} f_\omega(MO) = \max_{\omega \in \emptyset} f_\omega(MO) = 0$.
- $\min_{\omega \in \emptyset} f_\omega(MO) = \min_{\omega \in \Omega} f_\omega(MO) = 3$ (and $\max_{\omega \in \emptyset} f_\omega(MO) = 7$).

We can now compute the lower expectation:

$$\begin{aligned} \mathbb{E}_{\pi_\Omega}[\omega(MO)] &= 0.6 \min_{\omega \in \emptyset} \omega(MO) + 0.2 \min_{\omega \in [0, 5/9]} \omega(MO) \\ &\quad + 0.12 \min_{\omega \in [0, 5/9] \cup [2/3, 1]} \omega(MO) + 0.08 \min_{\omega \in [0, 1]} \omega(MO) \\ &= 0.6 \min_{\omega \in \emptyset} \omega(MO) + 0.2 \times 3 + 0.12 \times 3 + 0.08 \times 3 \\ &= 0.6 \min_{\omega \in \emptyset} \omega(MO) + 1.2. \end{aligned}$$

The lower expectation is equal to 1.2 if we take $\min_{\omega \in \emptyset} = 0$ (we ignore the conflict), or 3 if we take $\min_{\omega \in \emptyset} \omega(MO) = \min_{\omega \in \Omega} \omega(MO) = 3$ (we transform the conflict into ignorance).

For the upper expectation:

$$\begin{aligned} \mathbb{E}_{\pi_\Omega}[\omega(MO)] &= 0.6 \max_{\omega \in \emptyset} \omega(MO) + 0.2 \max_{\omega \in [0, 5/9]} \omega(MO) \\ &\quad + 0.12 \max_{\omega \in [0, 5/9] \cup [2/3, 1]} \omega(MO) + 0.08 \max_{\omega \in [0, 1]} \omega(MO) \\ &= 0.6 \max_{\omega \in \emptyset} \omega(MO) + 0.2 \times 47/9 + 0.12 \times 7 + 0.08 \times 7 \\ &\approx 0.6 \max_{\omega \in \emptyset} \omega(MO) + 2.44. \end{aligned}$$

The upper expectation is approximatively 2.44 if we take $\max_{\omega \in \emptyset} = 0$, or approximatively 6.64 if we take $\max_{\omega \in \emptyset} \omega(MO) = \max_{\omega \in \Omega} \omega(MO) = 7$.

Another way to infer despite inconsistencies would consist in normalizing the distribution π_Ω , to come back to a consistent situation. There are a lot of ways to perform such a normalization [29], yet they may be harder to interpret than the two solutions we consider here. For this reason, we will not explore them here, although the second way of handling inconsistency can be seen as a specific normalization, as already mentioned.

3.2. Resolving inconsistencies through information fusion

A second strategy to resolve inconsistencies is to change the way we combine the different sources of information, so that the inconsistency disappears. Such an approach does not modify the preferential information we receive, but is a convenient tool to make or test different hypotheses about them. For instance, a conjunctive rule resulting in π_{\cap} makes the assumption that all sources are reliable and provide trustworthy information. Clearly, if $\text{Inc}(\pi_{\cap}) > 0$, this assumption cannot be true, and others may be investigated. Wanting the fusion to be consistent with every piece of initial information is debatable in preference modelling, where all information are issued from the same user, meaning that if the model is to be trusted, inconsistency necessarily results from some error in the user answers. However, it is also reasonable to assume that most of the user answers are correct. We explore below some fusion operators that are inline with such a setting.

ℓ -out-of- k . It is tempting to use a fusion operator that can resolve inconsistencies to some extent, but whose result tends to be consistent with most (but not necessarily all) of the initial information. Since it is also difficult to know which answer coming from the user is wrong, it is a natural thing to consider operators that treat sources anonymously (*i.e.*, whose result remains unchanged if the indices of the information pieces π_i are permuted). The assumption that ℓ sources among the k considered are correct corresponds to such an operator. If $\mathcal{S} = \{\pi_1, \dots, \pi_k\}$ is the set of the considered items of information, then the distribution resulting from an ℓ/k assumption is:

$$\pi_{\ell/k}(\omega) = \bigcup_{\mathcal{L} \subseteq \mathcal{S}, |\mathcal{L}|=\ell} \left(\bigcap_{\pi_i \in \mathcal{L}} \pi_i(\Omega) \right), \quad (12)$$

where \cap and \cup are replaced by a T-norm and its dual T-conorm (in our case, the product T-norm and the probabilistic sum T-conorm). This fusion operator is an example of a ℓ -quota operator [18], applied to possibility theory. Ideally, a minimal repair should consist in finding a value ℓ as close as possible to k . Two steps are then required to perform such an idea:

1. find the highest ℓ^* such that $\pi_{\ell^*/k}$ is normalized,
2. compute the resulting distribution $\pi_{\ell^*/k}$.

While such operations may seem difficult to achieve in practice, we propose here an efficient method to achieve the first step, assuming that for each element Ω_i of the partition $\Omega_1, \dots, \Omega_P$ mentioned in Section 2.3 we do have an associated vector $\vec{\pi}_i = (\pi_1(\Omega_i), \dots, \pi_k(\Omega_i))$.

Algorithm 1 provides an easy way to find ℓ^* , and is based on the simple idea that $\pi_{\ell/k}$ will be normalized if there is at least an element Ω_i such that at least ℓ possibility degrees have a value one on this element (otherwise, we cannot select ℓ values such that applying a T-norm on these values will result in value one). Algorithm 1 then consists in finding the highest value satisfying this constraint. It is of linear complexity in the number P of elements, hence is quite fast once P is fixed.

The strategy consisting in performing an ℓ -out-of- k repair then consists of simply computing $\pi_{\ell^*/k}$ and making inference using this distribution. Since

Algorithm 1: Algorithm to find ℓ^*

Data: Sources $\mathcal{S} = \{\pi_1, \dots, \pi_k\}$
Result: Maximal ℓ^* to reach consistency
 $\ell^* = k$;
for $j \in \{1, \dots, P\}$ **do**
 if $|\{\pi_i(\Omega_j) : \pi_i(\Omega_j) = 1, i = 1, \dots, k\}| < \ell^*$ **then**
 $\ell^* = |\{\pi_i(\Omega_j) : \pi_i(\Omega_j) = 1, i = 1, \dots, k\}|$
 end
end

365 $\pi_{\ell^*/k}$ can be evaluated element-wise, computing it remains affordable as long as
 366 k is not too high, which is typically the case when querying information from
 367 the user.

368 **Example 9** (ℓ -out-of- k repair). We now suppose that the user gives 4 answers
 369 along with the certainty degrees $\alpha = \{0.9, 0.5, 0.7, 0.3\}$, as shown on Figure 8.
 370 Moreover, answer 4 is wrong because the user was either uncertain or unfocused,
 371 leading to some inconsistency being detected, as shown on Figure 9, with
 372 $\text{Inc}(\pi_\cap) = 0.3$. Our objective is to handle inconsistency, and more specifically to
 373 resolve current inconsistency through information fusion.




Figure 8: Answers given by the user in Example 9, answer 4 being wrong




Figure 9: Preferential information with inconsistency in Example 9

Here we will use ℓ -out-of- k repair algorithm. In this case, we can easily determine the maximal ℓ to reach consistency, which is $\ell = 3$: consistency is reached by removing a single answer, either answer 2 or 4. We then compute $\pi_{3/4}$ according to Equation (12). Given $\mathcal{S} = \{\pi_1, \dots, \pi_4\}$, the first step is to determine all the subsets $\mathcal{L} \subseteq \mathcal{S}$ such that $|\mathcal{L}| = 3$, obtaining $\mathcal{L}_1 = \{\pi_1, \pi_2, \pi_3\}$, $\mathcal{L}_2 = \{\pi_1, \pi_2, \pi_4\}$, $\mathcal{L}_3 = \{\pi_1, \pi_3, \pi_4\}$ and $\mathcal{L}_4 = \{\pi_2, \pi_3, \pi_4\}$. We then compute the 4 associated possibility distributions $\pi_{\mathcal{L}_i}$ through a product T-norm. For example, $\pi_{\mathcal{L}_1}(\omega) = \prod_{i=1}^3 \pi_i(\omega)$. After that, we compute $\pi_{3/4}$ through a probabilistic sum

T-conorm. Owing to its commutativity and its associativity, computing the *T-conorm* can be done iteratively through pairs of distributions. We recall that combining two distributions π_1 and π_2 through a *T-conorm* results in:

$$\pi_{\cup}(\omega) = \pi_1(\omega) + \pi_2(\omega) - \pi_1(\omega) \cdot \pi_2(\omega).$$

374 The final result is shown on Figure 10. As expected, consistency is restored and
 375 the resulting possibility distribution reaches 1 on two distinct subsets of Ω , in
 376 which at least 3 answers are consistent, which is the case for answers 1, 2 and 3
 377 when $\omega^1 \in [0.5, 0.65]$, and answers 1, 3 and 4 when $\omega^1 \in [0.7, 0.75]$. As indicates
 378 this remark, this approach does not guarantee that the set of most plausible models
 379 will be convex, even when each individual answer points out to a convex set of
 380 most plausible models. However, non-convex sets of most plausible answers will
 381 only happen in case of disagreement, and could be shown to the user for further
 382 investigations.




Figure 10: Preferential information corrected with ℓ -out-of- k repair algorithm in Example 9

Maximal Coherent Subsets. Rather than using an intersection operator that amounts to a logical AND, one can use other logical operators that will try to deal with conflicting and inconsistent situations, mostly by finding compromises between conjunctive and disjunctive behaviours. A common approach is for example to use the notion of maximal coherent subsets (MCS) [23]. In our context, and given a set $\mathcal{S} = \{\pi_1, \dots, \pi_k\}$ of considered items of information, we define a subset $\mathcal{L} \subseteq \mathcal{S}$ as a MCS if the result

$$\pi_{\mathcal{L}} = \bigcap_{\pi_i \in \mathcal{L}} \pi_i$$

383 of their combination is such that⁶ $\text{Inc}(\pi_{\mathcal{L}}) = 0$ and $\text{Inc}(\pi_{\mathcal{K}}) > 0$ for any $\mathcal{K} \supset \mathcal{L}$.
 384 A classical way to restore consistency through information fusion, inherited from
 385 ideas in logic [30], is simply to consider all MCS and take the disjunctions of
 386 all the MCS's conjunctions. Yet such an approach will typically deliver quite
 387 imprecise results in the presence of outliers or errors, mainly because it results in
 388 a combined distribution whose intersection with any of the initial (preferential)
 389 information is non-empty. We thus want to consider only one MCS, containing
 390 at least all the correct answers from the user.

391 Listing all the MCSs of a set of information is very costly: unlike Algorithm 1,
 392 we have to consider all possible subsets of information, thus at worst 2^K subsets.

⁶A weaker notion would be to require $\text{Inc}(\pi_{\mathcal{K}}) < 1$ and $\text{Inc}(\pi_{\mathcal{L}}) = 1$, but this would not restore full consistency, and would not be useful here, as distributions have Ω for support.

393 Supposing the number of information stays reasonable, listing all the MCSs
 394 is doable. A strategy would be to consider only the MCSs of size ℓ given by
 395 Algorithm 1, supposing most pieces of information are correct. However, as
 396 we will see in Section 4, while this heuristic can be interesting when paired
 397 with the associated average confidence degree, it usually does not give the most
 398 interesting MCS.

399 **Example 10** (MCS repair). *Keeping the same setting as Example 9, this time*
 400 *we want to resolve current inconsistency through a MCS, specifically a MCS of*
 401 *maximal size.*




Figure 11: Preferential information corrected through a MCS in Example 10

402 *As on the previous example, we have 4 answers, one of them being incorrect,*
 403 *and we know that $\ell = 3$. Therefore, we first need to determine all the MCSs*
 404 *\mathcal{L} such that $|\mathcal{L}| = 3$. Since it is not possible to have a MCS \mathcal{L} with $|\mathcal{L}| > 3$*
 405 *(otherwise ℓ would not be 3), it is sufficient to check only for coherent subsets,*
 406 *i.e. subsets \mathcal{L} such that $\text{Inc}(\pi_{\mathcal{L}}) = 0$. $\mathcal{L}_1 = \{\pi_1, \pi_2, \pi_3\}$ and $\mathcal{L}_3 = \{\pi_1, \pi_3, \pi_4\}$*
 407 *are the only coherent subsets of the specified size. We then need to pick the*
 408 *MCS that maximizes the average of the associated certainty degrees. We have*
 409 *$\alpha_{\mathcal{L}_1} = \frac{\alpha_1 + \alpha_2 + \alpha_3}{3} = 0.7$, and $\alpha_{\mathcal{L}_3} = \frac{\alpha_1 + \alpha_3 + \alpha_4}{3} \approx 0.63$, indicating $\hat{\mathcal{L}} = \mathcal{L}_1$.*

410 $\pi_{\hat{\mathcal{L}}}$ is shown on Figure 11. As we can see, consistency is restored and unlike
 411 ℓ -out-of- k repair algorithm, we are guaranteed that the set of most plausible
 412 models form a convex set if it is the case for each individual answers, thanks to
 413 the sole use of conjunctive operators.

414 Compared to the previous approaches of Section 3.1, modifying the way we
 415 combine information pieces is usually computationally more intensive, but has
 416 the advantage of potentially providing interesting insights to the user or the
 417 analyst. For instance, the number ℓ resulting from Algorithm 1 gives us a lower
 418 bound of the number of errors committed, while the set of answers constituting
 419 a MCS provides an interesting subset of answers that one could submit to the
 420 scrutiny of the user.

421 4. Experiments

422 In this section, we perform some synthetic experiments⁷ to see how our
 423 various approaches perform when inconsistency appears. As such, they provide
 424 proofs of concept that the approaches proposed in this paper have some interest

⁷https://github.com/LoicAdam/Possibilist_Elicitation_Fusion and https://github.com/LoicAdam/Possibilist_Elicitation_Fusion_Random

when treating uncertain preferential information. Confirming this interest in applied situations would require a real-world experiment, something that goes beyond the scope of this paper. Note that our experiments focus on assessing our inconsistency handling methods, as experiments in some of our previous work [26] already showed that the possibilistic approach can outperform the set-based approach in the presence of errors. In order to evaluate our approaches, we will consider a recommendation problem, where one item should be recommended based on the available information. In order to do this, we will first introduce some decision rules used to make such a recommendation (Section 4.1), as well as the elicitation approaches we will consider in the experiments (Section 4.2).

4.1. Decision rules

There are many decision rules when considering uncertain information in the form of possibility distributions, and we will only recall the ones we use here (the interested reader in other rules can check [31]).

Given a subset $\mathbb{A} \subseteq \mathbb{X}$ of available alternatives, the goal of the decision rules considered here is to make a recommendation $x^* \in \mathbb{A}$.

Maximin. Given an alternative x and a model ω , the function $\omega(x)$ provides an evaluation of the quality of x . When $\omega : \mathbb{X} \rightarrow \mathbb{R}$ is real-valued and our knowledge about it is encoded through a possibility distribution π , we can use Equation (6) to obtain $\mathbb{E}_\pi(\omega(x))$, and the corresponding Maximin recommendation:

$$x_{Mm}^* = \arg \max_{x \in \mathbb{A}} \mathbb{E}_\pi(\omega(x)), \quad (13)$$

which provides the strongest guarantees about its performance, as we adopt a pessimistic view.

Example 11. Given the available alternatives of Example 1, we want to determine the best alternative given the Maximin decision rule. On Figure 12 is shown the score of each alternative given $\omega \in \Omega$, supposing the score is determined by a weighted sum of unknown weights and that we have no information on the set of possible models. For example, MO has a minimal score of 3, reached for $\omega = (0, 1)$.




Figure 12: Illustration of Maximin approach

We have $ED \succ EM \succ MO \succ TB \succ AC$, and alternative ED is suggested to the user according to the Maximin decision rule. Let us note that this alternative minimizes the loss in the worst-case scenario, but for any model ω , there is at least one alternative that does better than ED .

457 **Maximax.** The Maximax rule to recommend an alternative adopts an optimistic
 458 attitude, opposite to the Maximin one. Applying this rule comes down to
 459 consider Equation (7) to obtain $\bar{\mathbb{E}}_\pi(\omega(x))$, and the corresponding Maximax
 460 recommendation is:

$$x_{MM}^* = \arg \max_{x \in \mathbb{A}} \bar{\mathbb{E}}_\pi(\omega(x)). \quad (14)$$

461 Such an attitude provides fewer guarantees than the Maximin approach, but is
 462 more likely to recommend the best alternative.

463 **Example 12.** Given the available alternatives of Example 1, we want to deter-
 464 mine the best alternative given the Maximax decision rule. On Figure 13 is shown
 465 the score of each alternative given $\omega \in \Omega$, supposing the score is determined by
 466 a weighted sum of unknown weights and that we have no information on the
 467 set of possible models. For example, TB has a maximal score of 8, reached for
 468 $\omega = (1, 0)$.




Figure 13: Illustration of Maximax approach

469 We have $AC \succ TB \succ MO \succ EM \succ ED$, and alternative AC is suggested to
 470 the user according to the Maximax decision rule. Let us note that this alternative
 471 maximizes the gain in the best-case scenario, but can be bad or even the worst
 472 for other scenarios (here AC being the worst alternative whenever the price is
 473 not important).

474 **Minimax regret.** While Maximin approaches are known to provide safe rec-
 475 ommendations, in the sense that they try to maximize the gain in worst-case
 476 scenarios, they are often criticized for their too strong conservatism. Regret-based
 477 rules preserve the idea of making safe recommendation with strong guarantees,
 478 while limiting the potential conservatism of Maximin. As recalled in the in-
 479 troduction, they are often used in incremental elicitation procedures, and in
 480 recommendation problems [6]. Let us introduce the main ideas behind Minimax
 481 regret recommendations.

482 The regret of choosing an alternative x over an alternative y for a specific
 483 model ω is defined by:

$$R_\omega(x, y) = \omega(y) - \omega(x). \quad (15)$$

484 It expresses the difference between what we would have obtained by picking y (the
 485 value $\omega(y)$), and what we actually obtain by picking x (the value $\omega(x)$). Given a
 486 set E of possible models, the value of Equation (15) is no longer well-defined,
 487 and the pairwise maximal regret over E is then defined as:

$$\text{PMR}(x, y, E) = \max_{\omega \in E} R_\omega(x, y), \quad (16)$$

488 corresponding to the maximum regret of choosing x over y for any model $\omega \in E$.
 489 The maximal regret of choosing x is then defined as:

$$\text{MR}(x, E) = \max_{y \in \mathbb{A}} \text{PMR}(x, y, E), \quad (17)$$

490 corresponding to the regret of choosing x in the worst case scenario, *i.e.*, against
 491 its worst opponent. Lastly, the minimal maximal regret of a set \mathbb{A} of alternatives
 492 given a set E of possible models is defined as:

$$\text{mMR}(E) = \min_{x \in \mathbb{A}} \text{MR}(x, E), \quad (18)$$

493 and $x^* = \arg \text{mMR}(E)$ is an alternative that gives the minimal regret in the
 494 worst-case scenario, corresponding to the optimal recommendation (in terms of
 495 regret) if no further information can be collected.

496 **Example 13** (Initial choice). *From Example 1, we first compute $\max_{\omega \in \Omega}$*
 497 *$R_\omega(x_i, x_j) \forall i, j \in \mathbb{A}^2$. In absence of information, the initial set of possible models*
 498 *is Ω , which is only constrained by $0 \leq \omega^i \leq 1$ and $\sum_i \omega^i = 1$, $i \in \{1, 2\}$. We com-*
 499 *pute the pairwise maximal regrets by optimizing $\max_{\omega \in \Omega} (\omega(x_i) - \omega(x_j)) \forall i, j \in$*
 500 *\mathbb{A}^2 . Since we optimize a linear function over a convex polytope Ω , the optimiza-*
 501 *tion problem is solved easily and exactly using linear programming (LP). We*
 502 *thus obtain the PMR, as shown in Table 2. For example, if the user chooses the*
 503 *truffle Brie over the Emmental, her maximal regret is $\text{PMR}(TB, EM) = 5$: for*
 504 *the model $\omega = (0, 1)$, we have $\omega(TB) = 1$ and $\omega(EM) = 6$, giving a maximal*
 505 *regret of 5.*

Table 2: Initial mMR determination of Example 13

x/y	AC	EM	ED	MO	TB		MR
AC	0	4	5	7	8		8
EM	3	0	1	3	4		4
ED	4	1	0	2	3		4
MO	6	3	2	0	1		6
TB	8	5	4	2	0		8
mMR							4

506 *The corresponding MR is given in Table 2. We obtain $\text{mMR} = 4$, the best*
 507 *initial choice being $x^* = EM$ or ED , which are the least regretted in the worst*
 508 *case scenario when having no information on the preferences of the user.*

509 In a previous paper [26], we extended minimax-regret notions to the pos-
 510 sibilistic setting, in the following way: our extension of PMR, named EPMR,
 511 averages the PMR over the different α -cuts:

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

512 where $1 = \alpha_1 > \dots > \alpha_n > \alpha_{n+1} = 0$ are the distinct values of π^k . If $\pi^k = \mathbb{I}_{E_k}$,
 513 we obtain the standard PMR of Equation (16). Note that such an averaging is
 514 standard in possibilistic approaches (see for example [32]).

515 If $E_{\pi^k}^{\alpha_1} = \emptyset$, we need to define $\text{PMR}(x, y, \emptyset)$. There are different options to
 516 do so [33], and we have discussed the main ones in Section 3.
 517 Similarly, our extension of the MR, the EMR, averages the MR over α -cuts:

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

518 corresponding to the average over cuts of the maximal pairwise regret. Again, if
 519 $\pi^k = \mathbb{I}_{E_k}$, we obtain the standard MR of Equation (17).
 520 Finally, we propose to extend the mMR with the mEMR:

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

521 which, since $\text{EMR}(x, \pi^k)$ reduces to Equation (17) when $\pi^k = \mathbb{I}_{E_k}$, also reduces
 522 to Equation (18) in the same case. When looking at Equation (7), it is clear
 523 that EPMR and EMR can be interpreted as upper expected values of regret,
 524 given our possibilistic knowledge and as long as this latter is normalized.

Example 14. From Example 6, we have 3 different values for the possibility distribution: 1, 0.3 and 0 (when $\omega \notin \Omega$). We thus have two different alpha cuts, and we deduce the EPMR for any pair (x, y) as follows:

$$\begin{aligned} \text{EPMR}(x, y, \pi) &= (1 - 0.3) \text{PMR}(x, y, E_{\pi}^1) + (0.3 - 0) \text{PMR}(x, y, E_{\pi}^{0.3}) \\ &= 0.7 \text{PMR}(x, y, E_{EM \succeq TB}) + 0.3 \text{PMR}(x, y, \Omega), \end{aligned}$$

and the corresponding EMR:

$$\text{EMR}(x, \pi) = 0.7 \text{MR}(x, \Omega_{EM \succeq TB}) + 0.3 \text{MR}(x, \Omega).$$

525 4.2. Retained elicitation procedures

526 In the experiments described in the sequel, we decided to use two elicitation
 527 strategies: one where we consider batch, non-incremental elicitation, that are
 528 typical of those cases where we collect information to estimate a possible model;
 529 and another where we consider incremental, optimized elicitation procedures
 530 whose goal is not especially to estimate a precise model, but rather to converge
 531 as quickly as possible to a good solution and recommendation. We now describe
 532 those two procedures.

533 **Batch elicitation.** Batch or non-incremental elicitation means that all prefer-
 534 ential information is given at once, in contrast to incremental elicitation where
 535 the requested preferential information in a given step depends on the questions
 536 and answered received in the previous step.

537 Considering the subset $\mathbb{A} \subseteq \mathbb{X}$ of available alternatives, batch elicitation in
 538 this context amounts to selecting pairs (x, y) such that the user provides an
 539 answer $x \succeq_{\alpha} y$ or $y \succeq_{\alpha} x$, that is either x is preferred to y with a certainty
 540 degree α , or the reverse. We will explain in Section 4.3 how those answers are
 541 simulated.

542 **Incremental regret-based strategy.** We now recall the Possibility Current
 543 Solution Strategy (PCSS) in order to select questions, extending the Current
 544 Solution Strategy (CSS) one. The CSS strategy was initially proposed as an
 545 efficient way to efficiently converge to a recommendation minimizing regret, and
 546 we refer to [34] for details. Here, we will only recall its possibilistic adaptation,
 547 originally introduced in [26]. The strategy is summarized in Algorithm 2. We
 548 assume that the user provides a unique choice when being presented with a pair,
 549 that is translated as a non-strict preference.

Algorithm 2: PCSS algorithm

Data: Max number of queries Max_q , set Ω of models, set \mathbb{A} of alternatives

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

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

while $k \leq \text{Max}_q$ **do**

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

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

 User provides answer $x^* \succeq_{\alpha_k} y^*$ or $x^* \preceq_{\alpha_k} y^*$;

if *User answer is* $x^* \succeq_{\alpha_k} y^*$ **then**

$\pi^{k+1} = \pi^k \cdot \pi_{x^* \succeq_{\alpha_k} y^*}$

else $\pi^{k+1} = \pi^k \cdot \pi_{x^* \preceq_{\alpha_k} y^*}$;

$k = k + 1$;

end

550 The interest of such an approach is that it retains the nice properties of
 551 the robust and CSS approaches when their assumptions hold (right model
 552 choice and correct answers), in particular the convergence towards the right
 553 recommendation with high guarantees as the regret will decrease after each
 554 iteration (PCSS regret bounds being more conservative than the ones of CSS). In
 555 the case where their assumptions do not hold, PCSS may identify it through the
 556 obtention of unnormalised distributions. We refer to our previous study [26] and
 557 its experiments that shows that the PCSS approach can be efficient to detect
 558 mistakes in the user answers or in the model assumptions.

559 4.3. Experimental protocol

560 Our goal here is to show that including uncertainty modelling as well as refined
 561 strategies to handle inconsistency can be helpful when treating preferences, and
 562 more particularly when recommending alternatives from real-valued aggregation
 563 models. To do so, we will compare different situations using the decision rules
 564 and elicitation strategies described in the previous sections.

565 We want to recommend to a user the alternative she should prefer among
 566 50 multi-criteria alternatives that are Pareto undominated (no alternative is
 567 worse than the others, regardless of the preferences of a user). Each alternative
 568 has 4 criteria with $\mathcal{X}_i = [0, 1]$. Each alternative x is generated randomly with
 569 respect to a uniform distribution, *i.e.* $x_i \sim \mathcal{U}(0, 1) \forall i \in \{1, \dots, 4\}$. We also add
 570 another constraint on the criteria to ensure that the alternatives are not Pareto
 571 Dominated: for each alternative x , the sum of the criteria is approximatively 2,
 572 *i.e.* $\sum_{i=1}^4 x_i \approx 2$. We also performed experiments on less and more criteria (*i.e.*,

3 and 5), but the conclusions reached during those experiments did not differ from the ones presented here.

The aggregating function modelling user preferences is a weighed sum parametrized by a set of weights such that $0 \leq \omega^i \leq 1$ and $\sum_{i=1}^4 \omega^i = 1$. A weighed sum is one of the most simple and used aggregating function of preferences, and we refer to the book of Grabisch *et al.* [35] for a detailed presentation of the different aggregating functions. With a weighted sum, each answer of the user, given as a comparison between two alternatives, can be represented by a linear constraint refining the space of possible models Ω . This is also true for any model that becomes linear in its parameters, such as OWA models [36] or Choquet integrals [6]. It should also be noted that, in this case, the size of the partition Ω induced by pairwise choices only increases polynomially with each question [37, P. 39], rather than the worst case exponential increase. We also refer to [33, Sec. 3.4.] for a discussion about computational issues when using models such as possibility distributions and belief functions.

The weights of each weighted sum are randomly generated according to a Dirichlet distribution with hyperparameter $(1, 1, 1, 1)$. The Dirichlet distribution gives us sets of weights that are summing up to one, and this specific hyperparameter choice guarantees us that the sets of weights are generated uniformly⁸ on Ω .

To find the optimal recommendation, we apply the PCSS algorithm presented in Algorithm 2. The user has to answer 15 questions. The certainty degrees α_i provided with each answer are generated randomly depending on the scenario, using either a beta distribution $\mathcal{B}(a, b)$ or a uniform one $\mathcal{U}(0, 1)$:

- $\alpha_i \sim \mathcal{B}(7, 2)$ in an optimist scenario, where the user is confident of her choices;
- $\alpha_i \sim \mathcal{U}(0, 1)$ in an intermediate scenario.

Let us note that a pessimistic scenario, in which the user is very unconfident of her choices and makes a lot of errors, is unrealistic, as the user is supposed to do her best to help us and not be very adversarial.

We model the uncertainty of the user this way: given a certainty degree α_i , the likeliness that the user answers necessarily correctly is α_i , and randomly $1 - \alpha_i$. When the user answers randomly, we consider that the probability of an incorrect answer is 50%. Overall, the user has a probability $\alpha_i + \frac{(1-\alpha_i)}{2}$ to answer correctly, and a probability $\frac{(1-\alpha_i)}{2}$ to answer incorrectly. We also make sure that at least one answer is incorrect to retain the experiment (otherwise, no inconsistency is observed). Such a sampling is quite common in possibility and belief function theory, where this probability corresponds to the so-called Pignistic probability [38], and is equivalent to the Shapley value in game theory. Of course, to achieve simulations, any choice of probabilities within the set \mathcal{P} induced by the possibility distribution would also be valid, such as first sampling a probability β_i of being correct within $[\alpha_i, 1]$ for each question, and then having

⁸When a Dirichlet distribution has $(1, \dots, 1)$ as its hyperparameters, it is equivalent to a uniform distribution over the open standard $(K-1)$ -simplex. Using a uniform distribution and dividing the generated weights by their sum is not equivalent, as this latter generation process will produce models clustered around the gravity centre of the simplex.

615 β_i chance to answer correctly. In our experiment, modifying such a choice of the
 616 sampling probability within \mathcal{P} , the probability set induced by π , only marginally
 617 modified the results and did not change the main conclusions. We therefore only
 618 report results for the classical choice of the Pignistic probability.

619 In the possibilist elicitation without correction, we first need to infer despite
 620 inconsistencies, like presented in Subsection 3.1, using one of the two strate-
 621 gies: either ignoring conflict with $\min_{\omega \in \emptyset} f(\omega) = \max_{\omega \in \emptyset} f(\omega) = 0$, or consider
 622 conflict and transform it into ignorance with $\min_{\omega \in \emptyset} f(\omega) = \min_{\omega \in \Omega} f(\omega)$ and
 623 $\max_{\omega \in \emptyset} f(\omega) = \max_{\omega \in \Omega} f(\omega)$. We will make experiment for the two elicitation
 624 regime we mentioned before: the incremental PCSS strategy that aims at opti-
 625 mising the recommendation, and the batch setting. Under these two frameworks,
 626 we then compare multiple elicitation algorithms and fusion methods:

- 627 • In the experiments using PCSS, we also consider the classic CSS elicitation
 628 strategy, a robust approach in which the user is supposed to make no
 629 errors. This method is expected to under perform when the user gives
 630 wrong answers, as it did in previous experiments [26]. We also provide
 631 a strategy similar to the PCSS strategy, but using probabilities as a way
 632 to handle uncertainty, meaning that the answer (E_i, α_i) is translated into
 633 $P(E_i) = \alpha_i$ and $P(E_i^c) = 1 - \alpha_i$. Minimax regret and other elements
 634 of the CSS strategy then become simple expectations over the obtained
 635 partition, which is the same as for PCSS (meaning that the possibilistic
 636 and probabilistic approaches have the same complexity). This allows to
 637 compare the PCSS strategy with other common base lines;
- 638 • The PCSS strategy, with one of the strategy to infer despite inconsistencies
 639 (ignore conflict or transform conflict into ignorance), without any additional
 640 correction, as presented in Algorithm 2;
- 641 • Our ℓ -out-of- k repair algorithm (both in the PCSS and batch settings),
 642 presented in Subsection 3.2, applied to the elicitation result, as presented
 643 in Algorithm 1;
- 644 • Different MCS strategies (both in the PCSS and batch settings), presented
 645 in Subsection 3.2, again applied the elicitation result. For three MCS
 646 strategies, a MCS is selected by a heuristic: either 1) a random MCS of
 647 maximal size, not using the provided confidence degrees and that allows
 648 us to measure the usefulness of these confidence degrees; 2) the MCS
 649 whose answers maximize the average confidence $\bar{\alpha} = \frac{1}{|\text{MCS}|} \sum_{i \in \text{MCS}} \alpha_i$; or
 650 3) the MCS among the MCSs of the biggest size, whose answers maximize
 651 the average confidence. We also consider *) the MCS that corrects the
 652 inconsistency the most in order to have an idea of how well we could do
 653 by picking the best MCS (that is in theory unknown to the analyst);
- 654 • A simple and naive algorithm to restore consistency by relaxing all the
 655 linear constraints $Ax < b$ associated to each answer by the user. In order
 656 to do so, for each answer i we add a variable δ_i such that $A_i x - \delta_i < b_i$
 657 and we try to find the values of δ_i such that consistency is restored and
 658 $\sum_{i=1}^n \delta_i$ is minimal.

659 Each scenario is then repeated 300 times, to have a reasonable sample size.

4.4. Number of incorrect answers detected

We first want to see if our ℓ -out-of- k fusion method is able to detect the number of incorrect answers given by the user, as it is useful information to know, both for the analyst and the user. We only report results for the PCSS method and the strategy ignoring inconsistency, as the other settings provide similar results. It should be reminded that we have no information on whether an answer given by the user is wrong or not (and thus no *a priori* information about the number of wrong answers). Without further information, an analyst can only rely on the provided answers and their associated certainty degrees: a high certainty degree means that we are confident that the answer is correct, while a low certainty simply means that we have no idea whether the information is correct or not. It should be stressed that a low confidence degree is not interpreted here as a sign that the given information is likely to be false, as they are necessity degrees, in which case a low degree indicates that we do not know whether the given information is true or false. So a number $\alpha = 0$ here means that we simply have no idea whether the answer is correct or not. It contrasts with a probabilistic interpretation of the number $P(E_i) = 0$, in which case one is sure that E_i^c is true. The closest probabilistic statement to our $\alpha = 0$ would be $P(E_i) = 1/2$, but with a different semantic.




Figure 14: Number of errors detected given the real number of errors (optimist scenario)

On Figure 14 is shown the number of errors detected ℓ returned by Algorithm 1, given the real number of errors that and supposing the user is very confident of her choices ($\alpha_i \sim \mathcal{B}(7, 2)$). Given this setting, a user usually makes between 1 and 4 errors, rarely more, out of 15 answers. As we can see, there is a positive correlation between ℓ^* and the real number of errors, hence ℓ^* can be used as a reasonable proxy. The difference between the number returned by Algorithm 1 and the real number is explained easily: a wrong answer does not necessarily contradict all the correct answers, meaning that a wrong answer does not automatically create inconsistencies (think for example of the case where

the first answer is wrong).




Figure 15: Number of errors detected given the real number of errors (uniform scenario)

On Figure 15 is shown the number of errors detected like on Figure 14, but given the user is not necessary confident of her choices ($\alpha_i \sim \mathcal{U}(0, 1)$). This time, the number of incorrect answers is way higher, with 1 to 7 incorrect choices among 15 answers most of the time. Our fusion method is less effective for detecting the real number of incorrect answers, but this is in accordance with the previous results: given that the number of incorrect choices is this time way higher, our method is less able to detect all incorrect answers. As the number of mistakes becomes higher, the chance that multiple mistakes are consistent between them increases, making their detection more difficult.

4.5. Performance of the different methods with PCSS algorithm

In this paper, we focused on the problem of recommending an item using a numerical model. Using a numerical model is advantageous, as we have a direct numerical measure of the performance of the different approaches. To obtain this, we compute over the repeated experiments the average of the real regret $R_{\omega_i}(x^*, x_i^{\text{opt}})$ between the alternative recommended by a method x^* and x_i^{opt} the best alternative given the true model ω_i of an experiment. We denote this average by:

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n R_{\omega_i}(x^*, x_i^{\text{opt}}). \quad (22)$$

We only kept the repetitions for which inconsistency was detected (otherwise all inconsistency handling methods coincide), so 210 repetitions for the optimistic scenario and 266 for the intermediate scenario. We also determined a confidence interval over the average \bar{x} : $\text{IC} = [\bar{x} \pm t_{n-1, 1-\frac{\alpha}{2}} \frac{S^*}{\sqrt{n}}]$ on \bar{x} , S^* being the corrected standard deviation of the real regrets, n being the number of repetitions kept, and $\alpha = 0.05$.

In order to compare the different methods and determine whether the differences are significative, we did some statistical paired difference tests with a significance level of 5%. As we do not want to assume that the differences are normally distributed (which is confirmed on most examples by Shapiro-Wilk's tests), we use non-parametric paired Wilcoxon signed-rank tests to determine whether the real regret differences between the two methods are negligible. We provide the p -value p associated to the tests, and if $p < 0.05$, we can assume that the differences are statically significant, especially given our sample size.

On the Figures 16 and 17 are shown confidence intervals on the average real regret, depending on the method and the confidence of the user. Each figure is divided into three parts:

- The top part contains inconsistency correction strategies that make no use of the possibilistic information: the naive algorithm that relax constraints, a random MCS having maximal size, and the theoretical MCS that minimizes the real regret. The first two give us set-based baselines that a method integrating confidence degrees should outperform, while the last one helps us to determine how good our heuristics are for finding MCSs;
- The middle part contains our different fusion repair algorithms (ℓ -out-of- k and the MCSs using confidence degrees obtained through heuristics), and should be compared with PCSS algorithm to determine whether the fusion strategies give better recommendations or not than simply specifying the value of $\min_{\omega \in \emptyset} f(\omega)$.
- The bottom part contains the CSS, the probabilist and the PCSS algorithms, to determine whether simply handling inconsistency through PCSS is helpful or not, and to compare it with standard approaches.




Figure 16: Confidence intervals on real regret for each method in the optimist scenario, the lower the real regret is, the better the final recommendation is.

A first question is whether the way we handle inconsistency in the incremental choice of questions (*i.e.*, ignoring conflict or consider conflict as ignorance during the possibilist elicitation) in Subsection 3.1 has an impact on the final result and recommendation. Given the average real regret with the classic elicitation

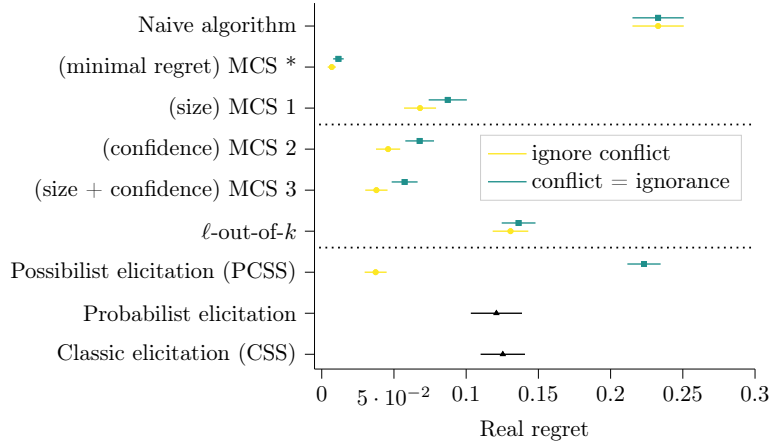


Figure 17: Confidence intervals on real regret for each method in the intermediate scenario, the lower the real regret is, the better the final recommendation is

741 ($\bar{x} = 0.0504$ in the strong scenario, 0.125 in the uniform scenario), we can see that
 742 ignoring inconsistency is a rather good strategy during the elicitation process
 743 (reduction of real regret in the strong scenario: 0.0423, $p = 1.59e - 16$. Uniform
 744 scenario: 0.0880, $p = 5.79e - 24$), as it presents a systematic and statistically
 745 significant gain when compared to the set-based approach, as well as with the
 746 probabilistic approach in the uniform scenario, as in this case probabilities are
 747 too uniformly distributed amongst true and false information). It is on par with
 748 the probabilistic approach in the case of confident users.

749 On the converse, we can see that transforming the inconsistency into ignor-
 750 ance during the elicitation process generally degrades the results in terms
 751 of recommendation (reduction of real regret in the strong scenario: -0.146 ,
 752 $p = 4.647e - 32$. Uniform scenario: -0.0978 , $p = 4.449e - 23$). So, while
 753 this option is quite common in the literature, our results clearly indicate that
 754 confusing inconsistency and ignorance is here a very risky and detrimental choice,
 755 at least when picking the questions to be presented to the user. This is due
 756 to the strong adopted strong bias that can vary between different alternatives
 757 (as each can have different maximal and minimal values). In contrast, ignoring
 758 conflict means that all alternatives are treated in the same way.

759 A second question is to know whether there is a difference, in terms of
 760 recommendation quality, between merely handling inconsistency through the
 761 redefinition of $\text{PMR}(x, y, \emptyset)$, and using more elaborated fusion strategies. Before
 762 going into details, we can see some first elements:

- 763 • Restoring consistency naively always gives the worst recommendations,
 764 and all methods are better regardless of the redefinition strategy;
- 765 • All fusion methods give better recommendations than the classic elicitation,
 766 except when picking a random MCS of the biggest size in the optimist
 767 scenario, and ℓ -out-of- k fusion method in the intermediate scenario. The
 768 fact that the results are much more robust to the conflict handling strategy
 769 for fusion rules also suggests that changing the conflict strategy has a major
 770 effect at the final inference time and not so much during the elicitation
 771 process;

Table 3: Difference in real regret between ignoring conflict and information fusion methods, a positive value meaning the fusion methods have a lower real regret in average

	Optimist	Uniform
ℓ -out-of- k	-0.00270	-0.0934
p -value	0.0430	4.21e-31
MCS from heuristics	-0.00499	-0.008312
p -value	0.122	0.00483
MCS minimizing \bar{x}	0.00666	0.0304
p -value	1.78e-08	4.76e-23

- All fusion methods give better recommendations than the ones when only considering conflict as ignorance. Given this result, we will focus on simply ignoring conflict in the rest of this subsection.

In addition to the previous figures, Table 3 summarizes the results when considering as reference method $\text{PMR}(x, y, \emptyset) = 0$ (ignoring conflict), and as compared methods the different fusion strategies. For simplicity, when it comes to MCSs we only give the results for the MCS of maximal size whose answers minimizes the average confidence, and the one minimizing the real regret. Given a fusion strategy and a scenario, we provide two numbers. The first one is the difference in average in real regret between the fusion strategy and just inferring with $\text{PMR}(x, y, \emptyset) = 0$, computed as $\frac{1}{n} \sum_{i=1}^n R_{\omega_i}(x_i^{\text{fusion}}, x_i^{\text{infer}})$. A negative number means that in average the fusion strategy gives worse recommendations, while a positive number means that in average the fusion strategy gives better recommendations. The second number is the p -value, to determine whether the difference in recommendation quality is statistically significative or not.

We can see that the results are quite mixed: ℓ -out- k fusion method and the MCS determined by a heuristic in average do not reduce the real regret, and even worse in average they slightly increase the real regret. However, this behaviour is in average, and for some instances they do reduce the real regret. Let us note though that the ℓ -out-of- k fusion strategy perform poorly on the uniform scenario. Table 3 also shows that the MCS minimizing the real regret reduces the real regret in average, regardless of the scenario, which is interesting.

4.6. Influence of the decision rule with a batch elicitation

In order to support the results found with the PCSS algorithm, we ran a simple test in which the questions were not chosen according to the PCSS algorithm, but completely at random, corresponding to a batch elicitation (preferential information is given at once, and the questions do not depend on the previous steps). Moreover, we want to know the influence of the decision rule on the results. Figures 18 and 19 show the performances obtained by the various approaches for the optimist and the uniform scenarios, respectively. They are divided in three parts like Figures 16 and 17, with the only difference being in the bottom part: it only contains the PCSS algorithm with the two strategies to handle inconsistency (ignore conflict, or transform conflict into ignorance).

We can observe the following results:

- If we focus on the Minimax regret decision rule (mMR), we can see that even in a batch setting, the results observed in Subsection 4.5 hold;
- Maximax decision rule (MM) provides in general recommendations very similar to the Minimax regret, albeit with a slightly higher real regret in average;
- Maximin decision rule (Mm) provides the worst recommendations regardless of the method (except for the random MCS of biggest size). However, when corrected by a fusion method, the recommendations become similar to the other decision rules;
- Strategies not using the confidence degrees (random biggest MCS or Naive restoration) perform badly, suggesting that the use of possibilistic information presents an advantage;
- Again, using fusion rules, and in particular MCS approaches using confidence degrees, gives much more stable as well as performant results across experiments. The performances of the other approaches can indeed vary a lot with respect to the chosen decision rule. This is a strong point of such approaches, as picking a particular decision rule is not always obvious.




Figure 18: Confidence intervals on real regret for each method in the optimist scenario in a batch setting with different decision rules

4.7. Overall result discussion

With the previous results, we can try to draw some first general conclusions:

- It seems that fusion methods can improve recommendations over inconsistency tolerant inference strategies when those latter provide rather bad recommendations to start with. In the other cases, at best the use of information fusion strategies provides no further improvement (with one exception), which seems normal if we start from an already good situation;




Figure 19: Confidence intervals on real regret for each method in the intermediate scenario in a batch setting with different decision rules

- The raw numerical benefits of using information fusion strategies appear to be limited for most of the fusion methods, as the increased requirement in computational power does not seem to be always paid-off by a significant increase in recommendation quality. It should however be reminded that those fusion rules also have an analytical interest, in the sense that they can give us useful information about the potential number of errors, as shows Figure 14. However, the true interest of such analytical properties is hard to assess in purely synthetic experiments, and testing them in practical scenario is something we would like to pursue in further research;
- The MCS that minimizes the real regret cannot be always be found by some simple heuristics. However, by sorting the MCSs based on their size and then their average confidence, a MCS minimizing the real regret is usually found among the first (mean rank in the optimist scenario: 1.5, the specific MCS ranked between 1st and 2nd for 90 % of repetitions ; mean rank in the intermediate scenario: 3.02, the specific MCS ranked between 1st and 4th for 80 % of repetitions). Therefore, it is worthwhile to present such sets of answers to the user, so that she can examine those more closely, and possibly determine the good answers;
- Using the Maximax decision rule instead of the Minimax regret only slightly impact the quality of recommendations. Maximin decision rule does impact negatively the quality of the recommendations, especially when no correction is applied. However, our fusion methods can restore the quality of the recommendations from Maximin decision rule to acceptable levels

It should be noted that the generalization of the above remarks, if they appear intuitively reasonable, should be checked by further synthetic or real-world experiments. Indeed, we considered a specific incremental elicitation method that is known to provide good performances in general, and we considered very simple questions which answers provide very limited information about the model: this means on the one hand that corrective actions will have a limited

860 impact (as the starting point is likely to already show good performances), and
861 on the other hand that many pieces of information may be consistent with each
862 others, even in the case of wrong answers.

863 5. Conclusions

864 In this paper, we have discussed general ways to integrate uncertainties in
865 preferential information through possibility theory. We see three main advantages
866 of using this theory:

- 867 • As a representation of uncertainty, it formally extends sets, and is therefore
868 coherent with robust approaches and their associated properties. It is also
869 quite relevant when looking at non-statistical problems, which is the case
870 when looking at a single user;
- 871 • A gradual assessment of the inconsistency present in preferential infor-
872 mation is doable, and can benefit from a very rich literature on how
873 to deal with such situations. In this paper, we have mainly considered
874 inconsistency-tolerant inference rules, as well as information fusion ap-
875 proaches allowing restoration of consistency before making inference. It
876 is also possible to obtain some interesting analysis of the situation. In
877 particular, the fact that most fusion rules can be directly associated to
878 logical operators or clear assumptions about the sources, provide further
879 readability that should be appreciated by a user;
- 880 • Compared to other uncertainty theories extending set-based approaches
881 such as belief functions or imprecise probabilities [39], the additional com-
882 putational cost of using possibility theory is limited, as the computational
883 complexity of the uncertainty representations grows linearly with the
884 number of collected information items.

885 To illustrate those aspects, we have concentrated on a weighted average aggrega-
886 tion function and some selected fusion rules in our empirical experiments. Those
887 experiments confirm the potential interests of our approaches and allowed us to
888 identify some of their limitations, *i.e.*, by identifying some situations where they
889 showed poor performances.

890 As already mentioned, many aspects of the current proposal, from Section 2
891 to Section 3, can easily be extended to situations other than the specific ones we
892 have considered in our empirical study. This includes, for example:

- 893 • Applying inconsistency handling techniques to other models or situations.
894 For instance, the regret-based approach has also been applied to multi-
895 objective combinatorial problems [7], more complex numerical models such
896 as Choquet integrals [6, 40], or even more qualitative models such as Sugeno
897 Integrals [41];
- 898 • Extending the proposed framework to consider other kinds of tasks or
899 inferences other than recommending a single alternative. Classical tasks
900 include ranking all alternatives, possibly partially, and sorting them into
901 ordered categories (see for instance [42, Ch. 7]). One simple starting
902 point would be considering minimax regret elicitation strategies for such
903 tasks [6, 43], and simply attach our possibilistic extension into those. It

would however require specifying how the corresponding inference (either ranking or assigning a category) can be performed given the final possibility distribution over models.

Finally, we already mentioned that our framework is very close in spirit to possibilistic logic, and could in fact be read as an instantiation of it if we restrict ourselves to Section 2. It would therefore be quite interesting to see how the handling of inconsistency in such logics [44] can help in our current framework.

Funding

This work was done within the PreServe Project, funded by grant ANR-18-CE23-0008 of the national research agency (ANR).

References

- [1] J. Fürnkranz, E. Hüllermeier, Preference learning, Springer, 2011.
- [2] P. Viappiani, C. Boutilier, On the equivalence of optimal recommendation sets and myopically optimal query sets, *Artificial Intelligence* 286 (2020) 103328.
- [3] H. Prade, A. Rico, M. Serrurier, Elicitation of sugeno integrals: A version space learning perspective, in: *International Symposium on Methodologies for Intelligent Systems*, Springer, 2009, pp. 392–401.
- [4] S. Corrente, S. Greco, M. Kadziński, R. Słowiński, Robust ordinal regression in preference learning and ranking, *Machine Learning* 93 (2) (2013) 381–422.
- [5] C. Boutilier, Computational decision support: Regret-based models for optimization and preference elicitation, *Comparative Decision Making: Analysis and Support Across Disciplines and Applications* (2013) 423–453.
- [6] N. Benabbou, P. Perny, P. Viappiani, Incremental elicitation of choquet capacities for multicriteria choice, ranking and sorting problems, *Artificial Intelligence* 246 (2017) 152–180.
- [7] N. Benabbou, P. Perny, Interactive resolution of multiobjective combinatorial optimization problems by incremental elicitation of criteria weights, *EURO journal on decision processes* 6 (3) (2018) 283–319.
- [8] V. Mousseau, J. Figueira, L. Dias, C. G. da Silva, J. Climaco, Resolving inconsistencies among constraints on the parameters of an mcda model, *European Journal of Operational Research* 147 (1) (2003) 72–93.
- [9] M. Schmitt, L. Martignon, D. Ron, On the complexity of learning lexicographic strategies., *Journal of Machine Learning Research* 7 (1) (2006).
- [10] B. Mayag, M. Grabisch, C. Labreuche, Dealing with inconsistencies in the representation of ordinal information by a 2-additive choquet integral, in: *Proceedings of the 7th conference of the European Society for Fuzzy Logic and Technology*, Atlantis Press, 2011, pp. 119–126.

- [11] M. Grabisch, I. Kojadinovic, P. Meyer, A review of methods for capacity identification in choquet integral based multi-attribute utility theory: Applications of the kappalab r package, *European journal of operational research* 186 (2) (2008) 766–785.
- [12] S. Teso, A. Passerini, P. Viappiani, Constructive preference elicitation for multiple users with setwise max-margin, in: *International Conference on Algorithmic Decision Theory*, Springer, 2017, pp. 3–17.
- [13] B. Eric, N. Freitas, A. Ghosh, Active preference learning with discrete choice data, *Advances in neural information processing systems* 20 (2007).
- [14] P. Viappiani, C. Boutilier, Optimal bayesian recommendation sets and myopically optimal choice query sets, *Advances in neural information processing systems* 23 (2010).
- [15] N. Bourdache, P. Perny, O. Spanjaard, Incremental elicitation of rank-dependent aggregation functions based on bayesian linear regression, in: *IJCAI-19-Twenty-Eighth International Joint Conference on Artificial Intelligence*, International Joint Conferences on Artificial Intelligence Organization, 2019, pp. 2023–2029.
- [16] J. S. Hammond, R. L. Keeney, H. Raiffa, Even swaps: A rational method for making trade-offs, *Harvard business review* 76 (1998) 137–150.
- [17] D. Dubois, H. Prade, Possibility theory: qualitative and quantitative aspects, in: *Quantified representation of uncertainty and imprecision*, Springer, 1998, pp. 169–226.
- [18] D. Dubois, W. Liu, J. Ma, H. Prade, The basic principles of uncertain information fusion. an organised review of merging rules in different representation frameworks, *Information Fusion* 32 (2016) 12–39.
- [19] H. Prade, A. Rico, M. Serrurier, E. Raufaste, Eliciting sugeno integrals: Methodology and a case study, in: *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, Springer, 2009, pp. 712–723.
- [20] G. Pigozzi, A. Tsoukias, P. Viappiani, Preferences in artificial intelligence, *Annals of Mathematics and Artificial Intelligence* 77 (3) (2016) 361–401.
- [21] D. Dubois, H. Prade, Possibilistic logic: a retrospective and prospective view, *Fuzzy sets and Systems* 144 (1) (2004) 3–23.
- [22] S. Benferhat, D. Dubois, H. Prade, Towards a possibilistic logic handling of preferences, in: *Aiding Decisions with Multiple Criteria*, Springer, 2002, pp. 315–337.
- [23] S. Destercke, D. Dubois, E. Chojnacki, Possibilistic information fusion using maximal coherent subsets, *IEEE Transactions on Fuzzy Systems* 17 (1) (2008) 79–92.
- [24] F. Pichon, S. Destercke, T. Burger, A consistency-specificity trade-off to select source behavior in information fusion, *IEEE transactions on cybernetics* 45 (4) (2014) 598–609.

- 984 [25] R. Mesiar, Triangular norms—an overview, *Computational intelligence in*
985 *theory and practice* (2001) 35–54.
- 986 [26] L. Adam, S. Destercke, Possibilistic preference elicitation by minimax regret,
987 in: *Uncertainty in Artificial Intelligence*, PMLR, 2021, pp. 718–727.
- 988 [27] M. Herin, P. Perny, N. Sokolovska, Learning sparse representations of
989 preferences within choquet expected utility theory, in: *Uncertainty in*
990 *Artificial Intelligence*, PMLR, 2022, pp. 800–810.
- 991 [28] L. Adam, S. Destercke, Multi-dimensional maximal coherent subsets made
992 easy: Illustration on an estimation problem, in: *International Conference*
993 *on Soft Methods in Probability and Statistics*, Springer, 2023, pp. 1–8.
- 994 [29] M. Oussalah, On the normalization of subnormal possibility distributions:
995 New investigations, *International Journal of General Systems* 31 (3) (2002)
996 277–301.
- 997 [30] N. Rescher, R. Manor, On inference from inconsistent premisses, *Theory*
998 *and decision* 1 (2) (1970) 179–217.
- 999 [31] N. Huntley, R. Hable, M. C. Troffaes, Decision making, *Introduction to*
1000 *imprecise probabilities* (2014) 190–206.
- 1001 [32] E. Hüllermeier, Learning from imprecise and fuzzy observations: Data
1002 disambiguation through generalized loss minimization, *International Journal*
1003 *of Approximate Reasoning* 55 (7) (2014) 1519–1534.
- 1004 [33] P.-L. Guillot, S. Destercke, Preference elicitation with uncertainty: Extend-
1005 ing regret based methods with belief functions, in: *International Conference*
1006 *on Scalable Uncertainty Management*, Springer, 2019, pp. 289–309.
- 1007 [34] C. Boutilier, R. Patrascu, P. Poupart, D. Schuurmans, Constraint-based
1008 optimization and utility elicitation using the minimax decision criterion,
1009 *Artificial Intelligence* 170 (8-9) (2006) 686–713.
- 1010 [35] M. Grabisch, J.-L. Marichal, R. Mesiar, E. Pap, *Aggregation functions*, Vol.
1011 127, Cambridge University Press, 2009.
- 1012 [36] R. R. Yager, On ordered weighted averaging aggregation operators in
1013 multicriteria decisionmaking, *IEEE Transactions on systems, Man, and*
1014 *Cybernetics* 18 (1) (1988) 183–190.
- 1015 [37] L. Schläfli, H. Wild, *Theorie der vielfachen kontinuierität*, vol. 38 (2013).
- 1016 [38] P. Smets, et al., Constructing the pignistic probability function in a context
1017 of uncertainty., in: *UAI*, Vol. 89, 1989, pp. 29–40.
- 1018 [39] S. Destercke, A generic framework to include belief functions in preference
1019 handling and multi-criteria decision, *International Journal of Approximate*
1020 *Reasoning* 98 (2018) 62–77.
- 1021 [40] H. Martin, P. Perny, Incremental preference elicitation with bipolar choquet
1022 integrals, in: *International Conference on Algorithmic Decision Theory*,
1023 Springer, 2021, pp. 101–116.

- 1024 [41] A. Rico, P. Viappiani, Incremental elicitation of capacities for the sugeno
1025 integral with a maximin approach, in: International Conference on Scalable
1026 Uncertainty Management, Springer, 2020, pp. 156–171.
- 1027 [42] D. Bouyssou, T. Marchant, M. Pirlot, A. Tsoukias, P. Vincke, Evaluation
1028 and decision models with multiple criteria: Stepping stones for the analyst,
1029 Vol. 86, Springer Science & Business Media, 2006.
- 1030 [43] Ö. Özpeynirci, S. Özpeynirci, V. Mousseau, A decomposition based minimax
1031 regret approach for inverse multiple criteria sorting problem, 4OR (2022)
1032 1–25.
- 1033 [44] D. Dubois, H. Prade, Inconsistency management from the standpoint
1034 of possibilistic logic, International Journal of Uncertainty, Fuzziness and
1035 Knowledge-Based Systems 23 (Suppl. 1) (2015) 15–30.