



HAL
open science

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. 17th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU 2023), Sep 2023, Arras, France. pp.239-253, 10.1007/978-3-031-45608-4_19 . hal-04335669

HAL Id: hal-04335669

<https://hal.science/hal-04335669v1>

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 Adam¹[0000-0003-2349-4827] and Sébastien Destercke¹[0000-0003-2026-468X]

UMR CNRS 7253 Heudiasyc, Sorbonne Université, Université de Technologie de
Compiègne CS 60319 - 60203 Compiègne cedex, France
{loic.adam, sebastien.destercke}@hds.utc.fr

Abstract Gathering the preferences of a user in order to make correct recommendations becomes a difficult task in case of uncertain answers. Using possibility theory as a means of modelling and detecting this uncertainty, we propose methods based on information fusion to make inferences despite observed inconsistencies due to user errors. While the principles of our approach are general, we illustrate its potential benefits on synthetic experiments using weighted averages as preference models.

Keywords: Preferences · Inconsistency · Information Fusion · Possibility.

1 Introduction

This paper focuses on handling uncertainty and inconsistency in the observed preferences of a single user. While multi-criteria decision analysis often focuses on specific users, other fields such as statistics, machine learning, and economics tend to look at populations. Traditionally, uncertainty and inconsistency in single-user preferences are addressed through set-based approaches, relying on techniques like min-max regret bounds [2,3,6]; or through average error calculations [5,13,14,23,24]. However, set-based approaches rely on the strong assumptions that both the user and the model choice are always correct, while probabilistic and averaging methods lack strong guarantees, justifying new approaches.

This paper explores a third approach, using possibility theory [10] to process uncertain preferential information. This approach remains consistent with a set-based approach while providing a non-binary quantification of inconsistency. It also provides various tools for dealing with inconsistency, extending set and logic operations such as conjunction and disjunction [9], unlike expectation-based operators from probabilities. While using possibilities for preferences is not new [21], our contribution enriches such proposals by incorporating information fusion methods to address inconsistency, and by validating the proposed methodologies through synthetic experiments. Sec. 2 describes our general possibilistic setting. Sec. 3 provides strategies to deal with user inconsistency. Sec. 4 provides some experiments demonstrating the potential interest of our approach.

All along the paper, we will illustrate our approach through weighted averages, as they are widely used and simple to understand. However, other numerical models could be applied, such as the ordered weighted averages (OWA) [25].

2 Possibilistic Modelling of Preference Models

2.1 Preferences and Preference Models

In this paper, we consider 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. We also assume that user preferences can be described by some model $\omega \in \Omega$, the set of models Ω being chosen by the analyst. Each model ω then induces a partial pre-order¹ over the set of alternatives. We refer to [20] for a list of possible models.

In this work, we focus on numerical models, where $\omega : \mathbb{X} \rightarrow \mathbb{R}$ is a real-valued function² that maps any alternative $x \in \mathbb{X}$ to a corresponding value $\omega(x)$, denoted as the score of the alternative x given the preference model ω . We note by ω^i the i th parameter of the said function. For easiness, we also denote by $x \succeq_{\omega} y$ the relation $\omega(x) \geq \omega(y)$. However, many of the ideas in this paper also apply to the case where ω is not a numerical model.

Example 1. A user wants to buy cheese, and we suppose that she evaluates a cheese through two criteria: flavour and price. If her preferences are described by a weighted sum with parameters summing to one (0.6, 0.4), we obtain the scores presented in Table 1 for a set of cheese. Mozzarella is her preferred alternative.

Table 1. Set of alternatives \mathbb{X} and their scores, with ω having parameters (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

2.2 Possibility Theory Reminder

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

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

Π and N are dual, as $N(A) = 1 - \Pi(A^c)$. Therefore, working with one of them for every event A is sufficient. Possibility theory formally extends sets, as the information given by a subset E is modelled by the distribution $\pi(x) = 1$ if

¹ A transitive, antisymmetric relation on $\mathbb{X} \times \mathbb{X}$.

² For simplicity, we will use the same notation for the function and its parameters ω .

$x \in E$, zero otherwise. When π is consistent, the bounds $[N(A), \Pi(A)]$ induce the probabilistic set $\mathcal{P} = \{P | N(A) \leq P(A) \leq \Pi(A), \forall A \subseteq \Omega\}$.

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

$$\pi_\alpha = \{\omega \in \Omega : \pi(\omega) \geq \alpha\}, \quad (2)$$

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

The last two significant notions, given a real-valued function $f : \Omega \rightarrow \mathbb{R}$, are the lower and the upper expectations $\underline{\mathbb{E}}_\pi(f), \overline{\mathbb{E}}_\pi(f)$ of f induced by a possibility distribution π . They are respectively the lower and the upper expected values of f over the set \mathcal{P} . When π takes a finite number of distinct values $1 = \alpha_1 > \dots > \alpha_n > \alpha_{n+1} = 0$ (being our case here), they can be computed as:

$$\underline{\mathbb{E}}_\pi(f) = \int_0^1 \min_{\omega \in \pi_\alpha} f(\omega) d\alpha = \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \min_{\omega \in \pi_{\alpha_i}} f(\omega), \quad (3)$$

$$\overline{\mathbb{E}}_\pi(f) = \int_0^1 \max_{\omega \in \pi_\alpha} f(\omega) d\alpha = \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \max_{\omega \in \pi_{\alpha_i}} f(\omega). \quad (4)$$

In this paper, we also consider unnormalized possibility distributions π such that $\max_{\omega \in \Omega} \pi(\omega) < 1$ to quantify the inconsistency:

$$\text{Inc}(\pi) = 1 - \max_{\omega \in \Omega} \pi(\omega). \quad (5)$$

2.3 Possibilistic Preferential Information

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

$$\pi_{(E, \alpha)}(\omega) = \begin{cases} 1 & \text{if } \omega \in E, \\ 1 - \alpha & \text{otherwise.} \end{cases} \quad (6)$$

In particular, $\alpha = 1$ corresponds to a set-valued information $\omega \in E$, while $\alpha = 0$ amounts to a void statement corresponding to ignorance. Equation (6) can be interpreted as an item of information within possibilistic logic [11], and most reasoning tools used in this paper could be interpreted through the lens of such a logic³. E is typically a subset of possible models resulting from a user answer.

Example 2 (Piece of information). Given Example 1 and Table 1, assuming that the user declares Truffle Brie \succeq_ω Emmental with a certainty degree $\alpha_1 = 0.8$, we obtain the following decision frontier:

$$\omega(\text{Truffle Brie}) \geq \omega(\text{Emmental}) \Rightarrow 8\omega^1 + \omega^2 \geq 4\omega^1 + 6\omega^2 \Rightarrow 4\omega^1 \geq 5\omega^2, \quad (7)$$

³ This should not be confused with possibilistic logic used to represent preferences [4], where α represents intensities of preferences.

corresponding to the possibilistic information $\pi_{(E_1, \alpha_1)}$ pictured in Figure 1, with $E_1 = \{\omega \mid \omega(\text{Truffle Brie}) \geq \omega(\text{Emmental})\}$ and $\alpha_1 = 0.8$.

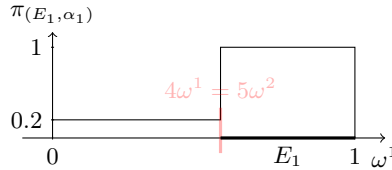


Figure 1. Preferential information $\pi_{(E_1, \alpha_1)}(\omega)$ of Example 2

In this paper, we will consider that a set E_i is the result of some pairwise comparison between a pair of alternatives $(x, y) \in \mathbb{X}^2$, where the user can either 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 from each possible answer.

In practice, we collect multiple pieces of information (E_i, α_i) , $i = 1, \dots, n$ during the elicitation process, each of them corresponding to a distribution $\pi_{(E_i, \alpha_i)}$. Note that those E_i will define a finite partition $\{\Omega_1, \dots, \Omega_P\}$ of Ω where Ω_i is of the kind $\cap_{\phi_i \in \{E_i, E_i^c\}} \phi_i$. The distributions $\pi_{(E_i, \alpha_i)}$ can then be combined or fused together into a single distribution by extending classical set operators such as conjunction (logical AND) and disjunction (logical OR). The use of such operators also allows for an easier interpretation of the performed operations [8,9,19]. In particular, if we have no reasons to think that the pieces of information $\pi_{(E_i, \alpha_i)}$ are unreliable⁴, the most sensible way to combine them is through conjunction, which in possibility theory is typically done through the use of a T-norm operator [17]. As our goal here is not to discuss the pros and cons of the different T-norms, we will focus on the product T-norm, resulting in the distribution π_\cap such that $\pi_\cap(\omega) = \prod_{i=1}^n \pi_{(E_i, \alpha_i)}(\omega)$.

Example 3 (Fusion of information and expectation bounds). Now we consider two pieces of information: $\pi_{(E_1, \alpha_1)}$ from Example 2, and a new one denoted $\pi_{(E_2, \alpha_2)}$ obtained from the affirmation of the user that Mozzarella \succeq_ω Truffle Brie with a certainty degree $\alpha_2 = 0.6$. We obtain a new decision frontier: $\omega^1 \leq 2\omega^2$. $\pi_{(E_2, \alpha_2)}$ is shown on Figure 2, and the resulting fused distribution is shown on Figure 3.

As an illustration of Equation (3), we can consider the function $f(\omega) = \omega(\text{Mozzarella}) = 7\omega^1 + 3\omega^2$. The lower expectation is:

$$\underline{\mathbb{E}}(f(\omega)) = (1-0.4) \times (5/9 \times 7 + 4/9 \times 3) + (0.4-0.2) \times (5/9 \times 7 + 4/9 \times 3) + 0.2 \times 3 \approx 4.77.$$

2.4 Errors in Set-wise and Possibilistic Approaches

As recalled in the introduction, set-wise approaches are especially useful when needing strong guarantees, as long as the information provided by the user is

⁴ We will deal with this situation in Section 3.

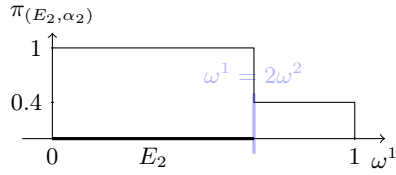


Figure 2. Preferential information $\pi_{(E_2, \alpha_2)}(\omega)$ of Example 3

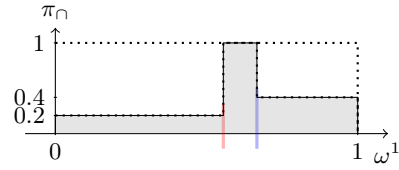


Figure 3. Fusion π_{\cap} of two preferential information of Example 3

correct. Yet, such hypotheses are often unrealistic, in which case using sets can lead to unwarranted situations, hence the need to account for possible mistakes through refined uncertainty modelling. As we have shown [1], using possibility theory is an interesting solution to this issue, as shows the next example.

Example 4 (A single error to ruin everything). We take Example 2 with two small but important modifications: we do not consider a possibilistic information, and thus only E_1 is considered (equivalent to $\alpha_1 = 1$); and the user is unfocused or unsure and makes the erroneous claim that Truffle Brie \preceq Emmental. We determine that $E_1 = E_{\text{Truffle Brie} \preceq \text{Emmental}}$ is now $\{\omega \in \Omega : 4\omega^1 \leq 5\omega^2\}$.

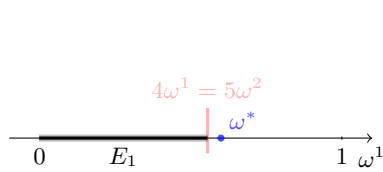


Figure 4. Wrong answer leading to a wrong model

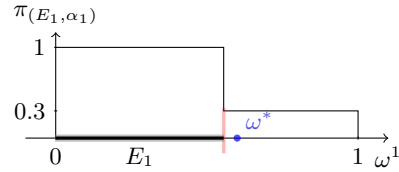


Figure 5. Possibilistic preferential information

As shown on Figure 4, the true preference model of the user, denoted by ω^* , is definitely left out of E_1 . Whatever the next answers are, we cannot get to ω^* . Now, if the user provides a certainty degree $\alpha_1 = 0.7$, we obtain the distribution shown on Figure 5, and ω^* is still reachable with further questions. With additional correct answers, the possibility of E_1 will decrease to a point that E_1 has a lower possibility than E_1^c , suggesting that ω^* is more likely to belong to E_1^c . In such a case, we detect some inconsistency, as $\max_{\omega \in \Omega} \pi(\omega) < 1$.

Using a possibilistic approach allows us to enrich set-based approaches while remaining consistent with it, as it is retrieved when giving $\alpha = 1$ as certainty degrees. This is in contrast with probabilistic approaches, where increasing the plausibility of some models necessarily means decreasing the plausibility of others.

3 Handling Inconsistencies

Another interest of the possibilistic approach is that when some answers are inconsistent between them, possibility distributions quantify inconsistency gradu-

ally, rather than having an all-or-nothing information as set-based approaches do. We assume in this section (and in the rest of the paper) that errors originate only 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.*, $\text{Inc}(\pi_{\cap}) > 0$.

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 $f(\omega)$ in Equations (3)-(4) without correcting inconsistencies, we need to define minima and maxima over the empty set. While an infinity of strategies could be considered, the two following ones are classical solutions:

- 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 somewhat similar to inference procedures in possibilistic logic in presence of inconsistency [11].
- 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 either a very conservative or optimistic 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 previously given answers and information (they are not modified), nor the way we combine them (conjunctively). This approach somehow avoids searching for the sources of inconsistency, and either ignores it or turn it into ignorance (a different concept). Therefore, such strategies appear legitimate only when inconsistency is likely to be limited, and when there is no need to analyse the details of the conflicting situation.

Another way to infer despite inconsistencies would consist in normalizing the distribution π_{\cap} , to come back to a consistent situation. There are a lot of ways to perform such a normalization [18], 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 combine differently the pieces of information, so that the inconsistency disappears. Such an approach does not modify the preferential information we receive, but is a convenient tool to test different hypotheses about them. For instance, a conjunctive rule resulting in π_{\cap} makes the assumption that all pieces of information are reliable. Clearly, if $\text{Inc}(\pi_{\cap}) > 0$, this assumption is not true, and others should be considered.

Compared to the previous approaches of Subsection 3.1, modifying the way we combine information pieces is usually computationally more intensive, but has the advantage of potentially providing interesting insights to the user or the analyst, for instance by giving us a lower bound of the number of errors committed, or giving a subset of answers to examine with the user.

ℓ -out-of- n In preference modelling, it is reasonable to assume that most of the user answers are correct, but not all (otherwise they would be consistent). Naturally, we want a fusion operator whose result can resolve inconsistencies while remaining 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 natural to consider operators that treat sources anonymously (*i.e.*, whose result is invariant under indices permutation). If $\mathcal{S} = \{\pi_1, \dots, \pi_k\}$ are the considered items of information, then the distribution resulting from a ℓ/k assumption is:

$$\pi_{\ell/k}(\omega) = \bigcup_{\mathcal{L} \subseteq \mathcal{S}, |\mathcal{L}|=\ell} \left(\bigcap_{\pi_i \in \mathcal{L}} \pi_i(\Omega) \right), \quad (8)$$

where \cup and \cap are replaced by a corresponding T-norm and dual T-conorm. This fusion operator is an example of a k -quota operator [9], applied to possibility theory. Ideally, a minimal repair should consist in finding a value ℓ as close as possible to k . We propose here an efficient method to determine such a ℓ , 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 $\pi_i = (\pi_1(\Omega_i), \dots, \pi_k(\Omega_i))$.

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 an element Ω_i for which at least ℓ possibility degrees have 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.

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

Example 5 (ℓ -out-of- k).

We now suppose that the user gives 4 answers along with the certainty degrees $\alpha = \{0.9, 0.5, 0.7, 0.3\}$, as shown on Figure 6. Moreover, answer 4 is wrong because the user was uncertain or unfocused, leading to some inconsistency being detected, as shown on Figure 7, with $\text{Inc}(\pi_{\cap}) = 0.3$. Our objective is to handle inconsistency, and more specifically to resolve current inconsistency through information fusion.

Here we will use ℓ -out-of- k 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 (8). Given $\mathcal{S} = \{\pi_1, \dots, \pi_4\}$, the first step is to determine all the subsets

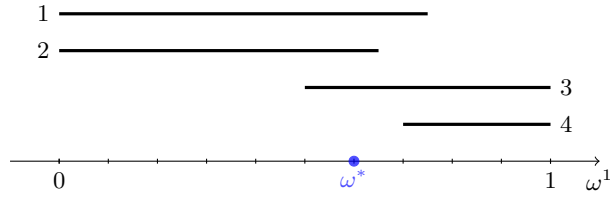


Figure 6. Answers given by the user in Example 5, answer 4 being wrong

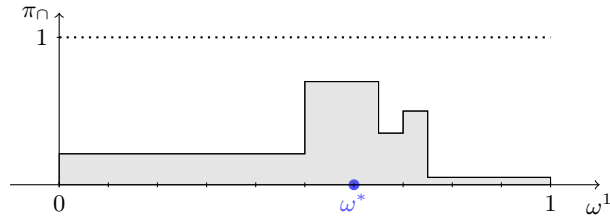


Figure 7. Preferential information with inconsistency in Example 5

$\mathcal{L} \subseteq \mathcal{S}$ 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. As such a T-conorm is commutative and associative, this can be done by iteratively applying it to pairs of distributions, recall that combining two distributions π_1 and π_2 through it results in:

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

The final result is shown on Figure 8. As expected, consistency is restored and the resulting possibility distribution reaches 1 on two distinct subsets of Ω , in which at least 3 answers are consistent, which is the case for answers 1, 2 and 3 when $\omega^1 \in [0.5, 0.65]$ and 1, 3 and 4 when $\omega^1 \in [0.7, 0.75]$. As indicates this remark, this approach does not guarantee that the set of most plausible models will be convex, even when each individual answer points out to a convex set of most plausible models. However, non-convex sets of most plausible answers will only happen in case of disagreement, and could be shown to the users for further investigations.

Maximal coherent subsets Another strategy for dealing with conflict is to use the notion of maximal coherent subsets (MCS) [8]. 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 $\pi_{\mathcal{L}} = \bigcap_{\pi_i \in \mathcal{L}} \pi_i$ of their combination is such that $\text{Inc}(\pi_{\mathcal{L}}) = 0$ and $\text{Inc}(\pi_{\mathcal{K}}) > 0$ for any $\mathcal{K} \supset \mathcal{L}$.

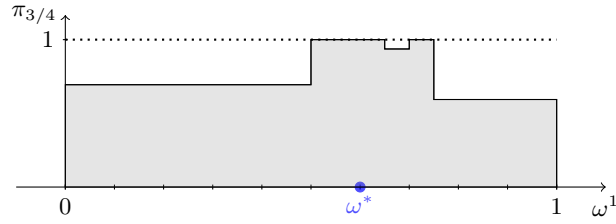


Figure 8. Preferential information corrected with ℓ -out-of- k repair algorithm in Example 5

A classical way to restore consistency through information fusion, inherited from ideas in logic [22], is simply to consider all MCSs and take the disjunctions of all the MCSs' conjunctions. An operator fusion based on MCSs will typically deliver quite imprecise results in the presence of outliers or errors, as the resulting distribution will have a non-empty intersection with any of the initial (preferential) information. In our setting, it seems more natural to consider only one MCS, hopefully containing all the correct answers from the user.

Listing all the MCSs can be very costly: unlike Algorithm 1, we have to consider all possible subsets of information: at worst 2^K subsets. Supposing the number of information stay reasonable, listing all the MCSs is doable. A strategy would be to consider only the MCSs of size ℓ^* given by Algorithm 1, supposing most pieces of information are correct. However, as we will see in Subsection 4.4, while this heuristic can be interesting when paired with the associated average confidence degree, it usually does not give the most interesting MCS.

Example 6 (MCS repair).

Keeping the same setting as Example 5, this time we want to resolve current inconsistency through a MCS, specifically a MCS of maximal size.

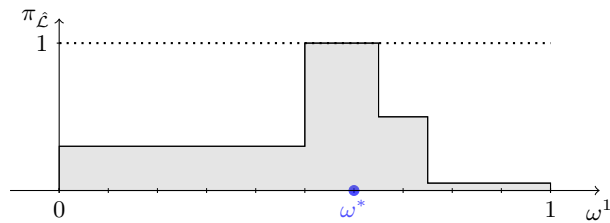


Figure 9. Preferential information corrected through a MCS in Example 6

As on the previous example, we have 4 answers, one of them being incorrect, and we know that $\ell = 3$. Therefore, we first need to determine all the MCSs \mathcal{L} such that $|\mathcal{L}| = 3$. Since it is not possible to have a MCS \mathcal{L} with $|\mathcal{L}| > 3$ (otherwise ℓ would not be 3), it is sufficient to check only for coherent subsets,

i.e. subsets \mathcal{L} such that $\text{Inc}(\pi_{\mathcal{L}}) = 0$. $\mathcal{L}_1 = \{\pi_1, \pi_2, \pi_3\}$ and $\mathcal{L}_3 = \{\pi_1, \pi_3, \pi_4\}$ are the only coherent subsets of the specified size. In this case, we can pick the MCS that maximizes the average of the associated certainty degrees. We have $\alpha_{\mathcal{L}_1} = \frac{\alpha_1 + \alpha_2 + \alpha_3}{3} = 0.7$, and $\alpha_{\mathcal{L}_3} = \frac{\alpha_1 + \alpha_3 + \alpha_4}{3} \approx 0.63$, indicating $\hat{\mathcal{L}} = \mathcal{L}_1$.

$\pi_{\hat{\mathcal{L}}}$ is shown on Figure 9. As we can see, consistency is restored and unlike ℓ -out-of- k repair algorithm, we are guaranteed that the set of most plausible models form a convex set if it is the case for each individual answers, thanks to the sole use of conjunctive operators.

4 Experiments

In this section, we perform some synthetic experiments⁵ to see how our various approaches perform when inconsistency appears.

4.1 Decision rules

Many decision rules exist when using possibility theory [15], and we will only recall the ones we use here. Given a subset $\mathbb{A} \subseteq \mathbb{X}$ of available alternatives, the goal of the decision rules we consider is to make a recommendation $x^* \in \mathbb{A}$. Given an alternative x and a model ω , the function $\omega(x)$ provides an evaluation of the quality of x . We consider three decision rules:

- Maximin, adopting a pessimistic view and providing safe recommendations: $x_{Mm}^* = \arg \max_{x \in \mathbb{A}} \underline{\mathbb{E}}_{\pi}(\omega(x))$,
- Maximax, adopting an optimistic view: $x_{MM}^* = \arg \max_{x \in \mathbb{A}} \overline{\mathbb{E}}_{\pi}(\omega(x))$,
- Minimax regret, less conservative than Maximin, still providing rather safe recommendations and widely used in recommendation problems (*e.g.* [3]): $x_{mMR}^* = \arg \min_{x \in \mathbb{A}} \sum_{i=1}^n (\alpha_i - \alpha_{i+1}) \max_{y \in \mathbb{A}} [\max_{\omega \in \pi_{\alpha_i}} (\omega(y) - \omega(x))]$.

4.2 Experimental Protocol

We simulate 50 multi-criteria alternatives. Each alternative has 4 criteria with $\mathcal{X}_i = [0, 1]$. For each alternative x , $x_i \sim U(0, 1)$ and $\sum_{i=1}^4 x_i \approx 2$ (so they are not Pareto-dominated). The true preference model ω^* of the user (a set of weights summing up to one) is randomly generated according to a Dirichlet distribution with hyperparameter $(1, 1, 1, 1)$, ensuring a uniform sampling of models.

To simulate a user elicitation process, we pick 15 pairs of alternatives. In the experiments below, the certainty degrees α_i provided with each answer are generated randomly according to a beta distribution $\mathcal{B}(7, 2)$, corresponding to an optimist scenario where the user is confident of her choices. To model the uncertainty of the user, given a certainty degree α_i , we consider that the user answers correctly with a probability α_i , and randomly (so sometimes correctly) with probability $1 - \alpha_i$. Overall, the user has a final probability $\alpha_i + \frac{(1 - \alpha_i)}{2}$ to answer correctly. Simulations are repeated 300 times, to have a reasonable sample size, and we consider only experiments with errors.

⁵ https://github.com/LoicAdam/Possibilist_Elicitation_Fusion_Random

4.3 Number of Errors Detected

We first want to see if Algorithm 1 is able to detect the number of incorrect answers. It should be reminded that we have no information on whether an answer given by the user is wrong or not.

Table 2. Number of errors detected given the real number of errors

Nb detected	Number errors					
	1	2	3	4	5	6
4					2	
3			8	8	<u>3</u>	
2		23	<u>24</u>	<u>9</u>	1	1
1	<u>84</u>	<u>38</u>	<u>24</u>	5	1	
0	54	12	2	1		

Table 2 shows the difference between the real number of incorrect answers and the number ℓ returned by Algorithm 1. The difference is explained easily: a wrong answer does not necessarily contradict all the correct answers, and does not automatically create inconsistencies (think for example of the case where the first answer is wrong). We can see that the higher the number of errors is, the more difficult it gets to assess correctly the number of errors. This is natural, as more errors are likely to be consistent between themselves.

4.4 Uncertainty Management Methods and Decision Rules

Using a numerical model allows us to easily measure the performance of the different approaches. To do so, we compute over the repeated experiments⁶ the average of the real regret $R_{\omega_i}(x^*, x_i^{\text{opt}}) = \omega(x_i^{\text{opt}}) - \omega(x^*)$ 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 $\bar{x} = \frac{1}{n} \sum_{i=1}^n R_{\omega_i}(x^*, x_i^{\text{opt}})$. We also determined a confidence interval $\text{IC} = [\bar{x} \pm t_{n-1, 1-\frac{\alpha}{2}} \frac{S^*}{\sqrt{n}}]$ on \bar{x} , S^* being the corrected standard deviation of the real regrets, with $\alpha = 0.05$ and $n = 231$.

Our question is to know whether there is a difference, in terms of recommendation quality, between merely handling inconsistency through the redefinition of $\min_{\omega \in \emptyset} f(\omega)$ and $\max_{\omega \in \emptyset} f(\omega)$, and using more elaborated fusion strategies. Figure 10 summarizes the results when comparing the redefinitions of Section 3.1 to the fusion approaches of Section 3.2. We also added a naive consistency restoration method, consisting in relaxing the linear constraints obtained from each answer until consistency is reached. Let us note we used different heuristics to pick a MCS (one without the confidence degrees, two based on them) and we also selected the MCS that truly minimizes the real regret.

We can see on this graph that not accounting for uncertainty degrees (the first rows) tend to provide worse results, as well as transforming inconsistency into

⁶ We only kept the repetitions for which inconsistency was detected, here 231.

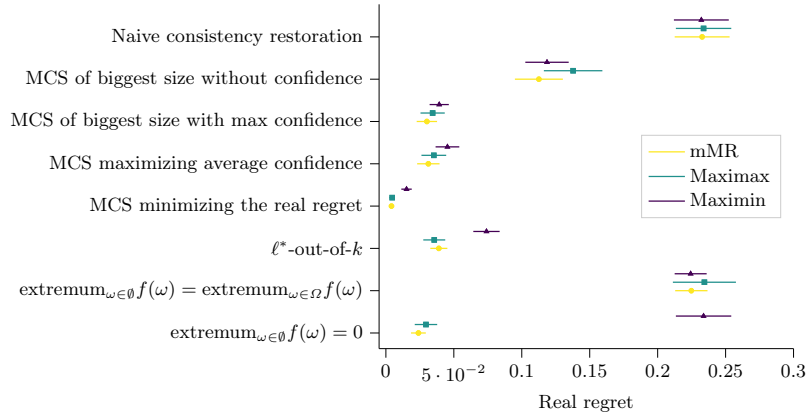


Figure 10. Confidence intervals on real regret

ignorance (before the last row), despite the fact that this is a common strategy. Simply ignoring the inconsistency ($\min_{\omega \in \emptyset} f(\omega) = \max_{\omega \in \emptyset} f(\omega) = 0$) gives much better results, but that are not robust to all decision rules. In contrast, using fusion rules to handle inconsistency provides good and stable results across all decision rules. In addition to this stability, such fusion rules also provide some additional insights to the user, such as an estimation of the number of errors, or some interesting sets of answers (*i.e.*, the MCS) to examine in more details.

These encouraging results should be further validated through additional synthetic or real-world experiments. The current results are based on random pairs of alternatives presented to the user.

5 Conclusion

In this paper, we have discussed integrating uncertainties in preferential information through possibility theory. Our experiments confirm the potential benefits of our approaches and some of their limitations.

Many aspects of the proposed framework can be easily extended to other situations beyond the scope of this paper. This includes multi-objective combinatorial problems [2] or more complex numerical models such as Choquet integrals [3,16]. The framework can also be applied to tasks like ranking alternatives or sorting them into ordered categories (see for instance [7, Ch. 7]).

Finally, our framework shares similarities with possibilistic logic. It would therefore be quite interesting to see how the handling of inconsistency in such logics [12] can help in our current framework.

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

References

1. Adam, L., Destercke, S.: Possibilistic preference elicitation by minimax regret. In: *Uncertainty in Artificial Intelligence*. pp. 718–727. PMLR (2021)
2. Benabbou, N., Perny, P.: Interactive resolution of multiobjective combinatorial optimization problems by incremental elicitation of criteria weights. *EURO journal on decision processes* **6**(3), 283–319 (2018)
3. Benabbou, N., Perny, P., Viappiani, P.: Incremental elicitation of Choquet capacities for multicriteria choice, ranking and sorting problems. *Artificial Intelligence* **246**, 152–180 (2017)
4. Benferhat, S., Dubois, D., Prade, H.: Towards a possibilistic logic handling of preferences. In: *Aiding Decisions with Multiple Criteria*, pp. 315–337. Springer (2002)
5. Bourdache, N., Perny, P., Spanjaard, O.: Incremental elicitation of rank-dependent aggregation functions based on Bayesian linear regression. In: *IJCAI-19-Twenty-Eighth International Joint Conference on Artificial Intelligence*. pp. 2023–2029. International Joint Conferences on Artificial Intelligence Organization (2019)
6. Boutilier, C.: Computational decision support: Regret-based models for optimization and preference elicitation. *Comparative Decision Making: Analysis and Support Across Disciplines and Applications* pp. 423–453 (2013)
7. Bouyssou, D., Marchant, T., Pirlot, M., Tsoukias, A., Vincke, P.: *Evaluation and decision models with multiple criteria: Stepping stones for the analyst*, vol. 86. Springer Science & Business Media (2006)
8. Destercke, S., Dubois, D., Chojnacki, E.: Possibilistic information fusion using maximal coherent subsets. *IEEE Transactions on Fuzzy Systems* **17**(1), 79–92 (2008)
9. Dubois, D., Liu, W., Ma, J., Prade, H.: The basic principles of uncertain information fusion. an organised review of merging rules in different representation frameworks. *Information Fusion* **32**, 12–39 (2016)
10. Dubois, D., Prade, H.: Possibility theory: qualitative and quantitative aspects. In: *Quantified representation of uncertainty and imprecision*, pp. 169–226. Springer (1998)
11. Dubois, D., Prade, H.: Possibilistic logic: a retrospective and prospective view. *Fuzzy sets and Systems* **144**(1), 3–23 (2004)
12. Dubois, D., Prade, H.: Inconsistency management from the standpoint of possibilistic logic. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* **23**(Suppl. 1), 15–30 (2015)
13. Eric, B., Freitas, N., Ghosh, A.: Active preference learning with discrete choice data. *Advances in neural information processing systems* **20** (2007)
14. Grabisch, M., Kojadinovic, I., Meyer, P.: 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), 766–785 (2008)
15. Huntley, N., Hable, R., Troffaes, M.C.M.: *Decision making*, chap. 8, pp. 190–206. John Wiley & Sons, Ltd (2014)
16. Martin, H., Perny, P.: Incremental preference elicitation with bipolar Choquet integrals. In: *International Conference on Algorithmic Decision Theory*. pp. 101–116. Springer (2021)
17. Mesiar, R.: Triangular norms—An overview. *Computational intelligence in theory and practice* pp. 35–54 (2001)

18. Oussalah, M.: On the normalization of subnormal possibility distributions: New investigations. *International Journal of General Systems* **31**(3), 277–301 (2002)
19. Pichon, F., Destercke, S., Burger, T.: A consistency-specificity trade-off to select source behavior in information fusion. *IEEE transactions on cybernetics* **45**(4), 598–609 (2014)
20. Pigozzi, G., Tsoukias, A., Viappiani, P.: Preferences in artificial intelligence. *Annals of Mathematics and Artificial Intelligence* **77**(3), 361–401 (2016)
21. Prade, H., Rico, A., Serrurier, M., Raufaste, E.: Eliciting Sugeno integrals: Methodology and a case study. In: *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*. pp. 712–723. Springer (2009)
22. Rescher, N., Manor, R.: On inference from inconsistent premisses. *Theory and decision* **1**(2), 179–217 (1970)
23. Teso, S., Passerini, A., Viappiani, P.: Constructive preference elicitation for multiple users with setwise max-margin. In: *International Conference on Algorithmic Decision Theory*. pp. 3–17. Springer (2017)
24. Viappiani, P., Boutilier, C.: Optimal bayesian recommendation sets and myopically optimal choice query sets. *Advances in neural information processing systems* **23** (2010)
25. Yager, R.R.: On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Transactions on systems, Man, and Cybernetics* **18**(1), 183–190 (1988)