Elicitation for decision problems under severe uncertainties

Nawapon Nakharutai[1] [0000-0002-3145-713X], Matthias C. M. Troffaes[3] [0000-0002-1294-600X], and Sébastien Destercke[2] [0000-0003-2026-468X]

[1] Department of Statistics, Faculty of Science, Chiang Mai University, Chiang Mai 50200, Thailand. nawapon.nakharutai@cmu.ac.th
[2] Department of Mathematical Sciences, Durham University, UK matthias.troffaes@durham.ac.uk
[3] UMR CNRS 7253 Heudiasyc, Sorbonne Université, Université de Technologie de Compiègne, France, sebastien.destercke@hds.utc.fr

Abstract. In this paper, we investigate the problem of eliciting information from an expert, where the assumed uncertainty model is a coherent upper prevision (or equivalently a closed convex set of probabilities). The goal is to solve a decision problem under the maximality decision rule, with as few queries to the expert as possible. To address this, we study the range of coherent upper bounds an expert may give on a given query. In doing so, we provide new results and characterisations for this range. We then use these results to provide an algorithm of elicitation. We illustrate the algorithm on an example.

Keywords: decision · elicitation · coherence

1 Introduction

In many applications, to support effective decision making under uncertainty, expert elicitation is critical. Elicitation faces several important challenges [16]. These include how to ask queries to experts to avoid psychological biases [9], how to elicit beliefs in addition to payments to eliminate hedging opportunities [4], how to translate communicated information into mathematical models, and how to elicit both preferences and beliefs about decision problems [11, 15].

While the existing literature on the topic is very rich, most authors focus on modelling expert knowledge in terms of precise probabilities (see for example Renooij [18] and Gartwaithe *et al.* [8]). This includes situations where the expert only provides partial information, as for instance in [14, Sec. 6.]. However, in case of severe uncertainty, it has been argued by many authors [27, 2, 22] that allowing the experts to provide bounded probability assessments rather than precise probabilities is more reasonable. Indeed, by allowing experts to provide bounded probability assessments, they can more honestly represent their uncertainty, especially in the face of severe uncertainty [27]. In this paper too, we will consider eliciting bounded probabilities.

Additionally, most existing elicitation procedures do not directly consider the decision problem for which the elicitation is eventually needed. This means that the elicitation procedure may be sub-optimal, i.e., we may ask the expert for information that is not needed to solve the decision problem at hand. While often elicitation of uncertainty is of interest by itself regardless of any decision problem, there are plenty of situations where uncertainties are elicited to solve a specific decision problem known at the time of elicitation. One aim of this paper is to explicitly consider the decision problem to more effectively target the elicitation, and thereby reduce the cognitive burden on the expert.

The problem of eliciting bounded probabilities has been studied extensively [21, 26], including elicitation in related uncertainty theories such as evidence theory [29] and possibility theory [20]. Applications include environmental modelling [19] and medical diagnosis [6]. However, here too, most of these approaches focus on eliciting the models without considering a possible underlying decision problem. Two exceptions are the work of Jansen *et al.* [10] and T'Joens [23]. However, these works differ from ours in several aspects. Jansen *et al.* [10] aim to find a preference system as efficiently as possible, whilst we aim at finding an optimal solution to the decision problem as efficiently as possible. T'Joens [23] frames the problem within the framework of desirability whilst we use upper previsions, and additionally makes the strong assumption that the unknown uncertainty model elicited from the expert leads to a unique optimal recommendation which we do not assume.

In this paper, we study the problem of elicitation in the context of a given (but arbitrary) decision problem: what are the best questions we can ask an expert, to identify the best possible actions for a specific decision problem? To tackle this, we treat the elicitation problem as an optimization problem, similarly to what is sometimes done when eliciting multi-criteria preferences [5, 3, 1]. Specifically, we solve the following problem:

> Given current expert assessments, and a set of potential queries that we can ask the expert, which query should we present next to gain the most useful information for solving the decision problem?

The questions we ask, called queries, are in the form of uncertain rewards, also called gambles, and the expert answers us by giving us their upper bound on the expectation of such a gamble; such upper bound is sometimes also called an upper prevision [28, 27, 25].

To determine which query is most useful, we use an objective that directly relates to the decision problem: which query has the most potential for reducing the set of so-called maximal [24] actions? To determine this, we need to know in advance the range of upper bounds on the expectation that the expert might provide. For this purpose, we develop the theory concerning the range of coherent extensions, expanding on some results from Walley [27, §3.1.6, p. 126]. We note that verification of coherence of conditional probability bounds was also studied in [6].

The rest of the paper is organised as follows. Section 2 reviews the theory of upper previsions and the concept of maximality in decision making. Section 3, proposes a new algorithm for elicitation, along with a fully worked example. Section 4 concludes the paper.

2 Preliminaries

2.1 Upper previsions

We denote by Ω a set of possible states. For simplicity, throughout this paper, we assume Ω is finite. We consider a decision problem where a decision maker can choose a finite number of acts from a set of acts $\mathcal{A}$. For each $a \in \mathcal{A}$, if $\omega \in \Omega$ happens to be the true state of nature, then $a(\omega)$ represents the reward to the decision maker, expressed as a real value on a utility scale. In this way, each act a is modelled as a real-valued function on a set of possible states. Such a function, representing an uncertain reward, is called a *gamble* [27]. Following [27], we will assume gambles are bounded, and $\mathcal{L}$ will denote the set of all gambles.

If a probability mass function p on Ω is available, the decision maker can simply pick an act that maximizes expected utility, i.e. $a^* = \arg\max_{a\in\mathcal{A}} \mathbb{E}_p(a)$, where $\mathbb{E}_p(a) \coloneqq \sum_{\omega\in\Omega} p(\omega)a(\omega)$. However, a complete probability mass function may not be available. In that case, the decision maker may ask an expert to state their knowledge about the unknown true value of $\omega \in \Omega$. We assume here that we can elicit some information in terms of an upper prevision $\overline{P}$ which is a real-valued function defined on some subset $\operatorname{dom}\overline{P}$ of $\mathcal{L}$. Specifically, for each gamble $f \in \operatorname{dom}\overline{P}$, the value $\overline{P}(f)$ represents the expert's infimum selling price for f. So, by specifying $\overline{P}$, the expert declares that $\overline{P}(f) - f + \epsilon$ is acceptable uncertain reward for all $\epsilon > 0$ and $f \in \operatorname{dom}\overline{P}$. As we shall see, $\overline{P}(f)$ can also be interpreted as an upper bound on the expectation of f. Note that upper previsions generalize a very wide range of uncertainty models used in practical problems, including belief functions and possibility distributions [25].

In this study, we are interested in iteratively making queries to the expert in order to learn their uncertainty in the form of upper previsions, so that the decision maker can solve the underlying decision problem. To do so, we must first state some properties and consistency conditions for upper previsions.

We say that $\overline{P}'$ *dominates* $\overline{P}$, and write $\overline{P}' \leq \overline{P}$, when $\operatorname{dom}\overline{P}' \supseteq \operatorname{dom}\overline{P}$ and $\overline{P}'(f) \leq \overline{P}(f)$ for all $f \in \operatorname{dom}\overline{P}$. When $\overline{P}'$ dominates $\overline{P}$, all selling prices implied by $\overline{P}$ are also implied by $\overline{P}'$, and so $\overline{P}'$ is more informative compared to $\overline{P}$.

Definition 1 ([27, §3.3.1, p. 133]). *The* credal set *of $\overline{P}$ is the set of probability mass functions on Ω whose expectations dominate $\overline{P}$, that is*

$$\mathcal{M}(\overline{P}) \coloneqq \left\{p \colon \forall f \in \mathcal{L},\, \mathbb{E}_p(f) \leq \overline{P}(f)\right\} \tag{1}$$

Note that $\mathcal{M}(\overline{P})$ is convex and closed. By $\mathcal{K} \Subset \operatorname{dom}\overline{P}$ we mean that $\mathcal{K}$ is a finite subset of $\operatorname{dom}\overline{P}$. We also define $\mathbb{R}_+ \coloneqq \{x \in \mathbb{R} \colon x \geq 0\}$.

Definition 2 ([27, §3.1.3(d)], [25, Cor. 8.38(ii)]). *The* natural extension *of* $\overline{P}$ *is the upper prevision* $\overline{\mathbb{E}}$ *defined on every* $f' \in \mathcal{L}$ *by*

$$\overline{\mathbb{E}}(f') := \inf_{\substack{\mathcal{K} \Subset \operatorname{dom} \overline{P} \\ \lambda\colon \mathcal{K} \to \mathbb{R}_+}} \sup \left[f' + \sum_{f \in \mathcal{K}} \lambda_f (\overline{P}(f) - f) \right] \tag{2}$$

$$= \max_{p \in \mathcal{M}(\overline{P})} \mathbb{E}_p(f') \tag{3}$$

In eq. (3), if $\mathcal{M}(\overline{P}) = \emptyset$, the maximum is taken to be $-\infty$. Equation (3) follows from eq. (2) by duality [27, §3.3.3, p. 134], but it also has an obvious direct interpretation as the upper expectation induced by the credal set of $\overline{P}$.

Definition 3 ([27, §2.4.4, p. 69] and [25, Cor. 8.38(i)]). *We say that* $\overline{P}$ avoids sure loss *when* $\overline{\mathbb{E}}(0) = 0$*, i.e. when*

$$\inf_{\substack{\mathcal{K} \Subset \operatorname{dom} \overline{P} \\ \lambda\colon \mathcal{K} \to \mathbb{R}_+}} \sup \left[\sum_{f \in \mathcal{K}} \lambda_f (\overline{P}(f) - f) \right] \geq 0. \tag{4}$$

or equivalently, when $\mathcal{M}(\overline{P}) \neq \emptyset$*.*

Definition 4 ([27, §2.5.4, p. 75], [25, Cor. 8.38(iii)]). *We say that* $\overline{P}$ *is* coherent *when* $\overline{\mathbb{E}}(f') = \overline{P}(f')$ *for all* $f' \in \operatorname{dom} \overline{P}$*, i.e.*

$$\inf_{\substack{f' \in \operatorname{dom} \overline{P},\, \mathcal{K} \Subset \operatorname{dom} \overline{P} \\ \lambda' \in \mathbb{R}_+,\, \lambda\colon \mathcal{K} \to \mathbb{R}_+}} \sup \left[\sum_{f \in \mathcal{K}} \lambda_f (\overline{P}(f) - f) - \lambda' (\overline{P}(f') - f') \right] \geq 0. \tag{5}$$

or equivalently, when for every $f' \in \operatorname{dom} \overline{P}$ *there is a* $p \in \mathcal{M}(\overline{P})$ *such that* $\overline{P}(f') = \mathbb{E}_p(f')$*.*

We assume that the decision maker remains coherent throughout the process of elicitation, i.e. that their $\overline{P}$ is a coherent upper prevision. Otherwise, $\overline{P}$ must be fixed through natural extension.

As we shall see next, the natural extension is also an important tool for solving decision problems.

2.2 Decision making

To select the best possible acts in $\mathcal{A}$, there are many decision criteria that can be used with upper previsions [24]. Here, we will consider *maximality*:

$$\operatorname{opt}(\overline{P}, \mathcal{A}) := \{a \in \mathcal{A} \colon \forall a' \in \mathcal{A}, \overline{\mathbb{E}}(a - a') \geq 0\} \tag{6}$$

The next property of maximality will be very useful later when we do elicitation.

Theorem 1 ([24]). *If* $\overline{P}'$ *dominates* $\overline{P}$ *then* $\operatorname{opt}(\overline{P}', \mathcal{A}) = \operatorname{opt}(\overline{P}', \operatorname{opt}(\overline{P}, \mathcal{A}))$.

Example 1. Consider a situation where we want to invest money in a stock fund. The return depends on the inflation rate across the next year. Suppose that next year's inflation rate could be either low (L), moderate (M) or high (H), that is, $\Omega = \{L, M, H\}$. There are six stock funds available, so $\mathcal{A} = \{1, \dots, 6\}$. The net profit under each fund and inflation rate is presented in table 1. Suppose that, prior to investment, we elicit an expert's information about next year's inflation rate. The expert states:

- L is at least as probable as M.
- L is at least as probable as H.

This information can be phrased in terms of an upper prevision:

$$\overline{P}(\mathbb{I}_M - \mathbb{I}_L) = 0 \text{ and } \overline{P}(\mathbb{I}_H - \mathbb{I}_L) = 0, \tag{7}$$

Here, we use the notation $\mathbb{I}_{\omega'}$ to denote the gamble formed by the indicator function of the event $\{\omega'\}$, i.e. $\mathbb{I}_{\omega'}(\omega) = 1$ if $\omega' = \omega$, else 0. For every pair of acts a and a' in $\mathcal{A}$, we can calculate $\overline{\mathbb{E}}(a - a')$. These values are listed in table 2. We can see that the set of maximal acts is $\operatorname{opt}(\overline{P}, \mathcal{A}) = \{3, 4, 5, 6\}$ since these acts have corresponding rows in table 2 with no negative entries.

stock fund	L	M	H
1	2	5	7
2	3	7	3
3	5	2	7
4	7	4	2
5	5	6	5
6	3	8	4

Table 1. Total net profit of investing money under each stock fund and inflation rate for Example 1.

a \ a'	1	2	3	4	5	6
1	0.0	1.5	0.0	0.33	−0.5	1.0
2	1.5	0.0	1.5	0.0	−.5	0.0
3	3.0	3.0	0.0	1.5	1.0	2.5
4	5.0	4.0	2.0	0.0	2.0	4.0
5	3.0	2.0	2.0	1.0	0.0	2.0
6	2.0	0.67	2.0	0.67	0.0	0.0

Table 2. Values of $\overline{\mathbb{E}}(a - a')$ in Example 1.

As we can see in example 1 the information from the expert allow us to suppress two non-maximal acts and leave us with four remaining maximal acts. Suppose we want to draw additional information from the expert by asking them to state an upper prevision of additional gambles which we will call *queries*. This new information may reduce the values of $\overline{\mathbb{E}}(a - a')$ and therefore the number of maximal acts (see theorem 1). Specifically, we would like to optimally pick an additional query that maximize our chances to suppress as many additional acts as possible. To do so, we must first study the range of possible coherent extensions.

2.3 Range of coherent extensions

We will aim to extend the domain of an expert's coherent upper prevision by one gamble at a time. To do so, we introduce the following notation. Given any coherent upper prevision $\overline{P}$, $g \in \mathcal{L}$, and $\beta \in \mathbb{R}$ we define $\overline{P}_{g,\beta}$ on $\operatorname{dom}\overline{P} \cup \{g\}$ as follows:

$$\overline{P}_{g,\beta}(f) := \begin{cases} \overline{P}(f) & \text{if } f \in \operatorname{dom}\overline{P} \setminus \{g\}, \\ \beta & \text{if } f = g. \end{cases} \tag{8}$$

By $\overline{\mathbb{E}}_{g,\beta}$ we denote the natural extension of $\overline{P}_{g,\beta}$.

We are interested in knowing, for a given $g \notin \operatorname{dom}\overline{P}$, the range of β for which $\overline{P}_{g,\beta}$ is coherent. A formula for this was given in [27, §3.1.6, p. 126]. The next theorem provides new formulas for the lower bound of this range, which are more concise and easier to interpret compared to [27, §3.1.6, p. 126].

Theorem 2. *Let $g \notin \operatorname{dom}\overline{P}$ and $\beta \in \mathbb{R}$. Then $\overline{P}_{g,\beta}$ is coherent if and only if*

$$\widetilde{\mathbb{E}}(g) \leq \beta \leq \overline{\mathbb{E}}(g), \tag{9}$$

where

$$\widetilde{\mathbb{E}}(g) := \sup_{f' \in \operatorname{dom}\overline{P}} -\overline{\mathbb{E}}_{-f',-\overline{P}(f')}(-g) \tag{10}$$

$$= \sup_{f' \in \operatorname{dom}\overline{P}} \min_{\substack{\mu \in \mathcal{M}(\overline{P}) \\ \mathbb{E}_\mu(f') = \overline{P}(f')}} \mathbb{E}_\mu(g). \tag{11}$$

Note that $\overline{P}_{-f',-\overline{P}(f')}$ is not necessarily coherent, but it avoids sure loss due to the coherence of $\overline{P}$ [27, p. 123]. So, for any given $g \notin \operatorname{dom}\overline{P}$, any $\overline{P}_{g,\beta}$ is coherent only for $\beta \in [\widetilde{\mathbb{E}}(g), \overline{\mathbb{E}}(g)]$. Therefore, this is the range of coherent value that an expert can provide, assuming the expert is coherent (an assumption we make here). If the bounds coincide, then we already have full knowledge about the upper prevision of g, and querying it would be useless.

Note that the upper bound $\overline{\mathbb{E}}(g)$ does not add any information to what we already know since $\overline{\mathbb{E}} = \overline{\mathbb{E}}_{g,\overline{\mathbb{E}}(g)}$ [25, Cor. 4.32], whilst the lower bound $\widetilde{\mathbb{E}}(g)$ is the most informative value a subject could give if questioned about g, by lemma 1.

Lemma 1. *$\overline{\mathbb{E}}_{g,\beta} \geq \overline{\mathbb{E}}_{g,\alpha}$ for all $\beta \geq \alpha$.*

Proof. $\overline{P}_{g,\alpha}$ dominates $\overline{P}_{g,\beta}$ whenever $\beta \geq \alpha$. So, by [25, Prop. 4.27], $\overline{\mathbb{E}}_{g,\alpha}$ must also dominate $\overline{\mathbb{E}}_{g,\beta}$.

3 Elicitation for decision problems

This section introduces and analyzes a new procedure to perform elicitation for decision problems under severe uncertainty. Unlike most elicitation procedures, this procedure tries to gain as much information as quickly as possible that is relevant to the decision problem at hand.

Algorithm 1 Elicitation

Require: $\mathcal{A}$, $\overline{P}$, $\mathcal{Q}$
Ensure: $\text{opt}(\overline{P}, \mathcal{A})$
1: **while** $\mathcal{Q} \neq \emptyset$ **do**
2: $\quad \mathcal{A} \leftarrow \text{opt}(\overline{P}, \mathcal{A})$
3: $\quad$ **for all** $q \in \mathcal{Q}$ **do**
4: $\quad\quad$ Compute $\widetilde{\mathbb{E}}(q)$ and $\overline{\mathbb{E}}(q)$
5: $\quad\quad$ Remove q from $\mathcal{Q}$ if $\widetilde{\mathbb{E}}(q) = \overline{\mathbb{E}}(q)$
6: $\quad\quad$ Otherwise, calculate $\text{opt}(\overline{P}_{q,\widetilde{\mathbb{E}}(q)}, \mathcal{A})$
7: $\quad$ **end for**
8: $\quad q^* \leftarrow \arg\min |\text{opt}(\overline{P}_{q,\widetilde{\mathbb{E}}(q)}, \mathcal{A})|$
9: $\quad$ Present q^* to the expert, who returns its upper prevision value, β
10: $\quad \overline{P} \leftarrow \overline{P}_{q^*,\beta}$
11: $\quad$ Remove q^* from $\mathcal{Q}$
12: **end while**
13: **return** $\text{opt}(\overline{P}, \mathcal{A})$

3.1 Elicitation procedure

The proposed procedure is given in algorithm 1.

We used theorem 1 in line 2 of the algorithm. Specifically, since all upper previsions considered later in the process will always dominate $\overline{P}$, all non-maximal acts from $\mathcal{A}$ with respect to $\overline{P}$ can be removed because these can never become maximal at any later stage of the elicitation process. This can significantly reduce the number of pairs a and a' for which we have to calculate $\overline{\mathbb{E}}(a-a')$ (each $\overline{\mathbb{E}}(a-a')$ requires solving a linear program).

By coherence, the expert will never give a value for β that is less than $\widetilde{\mathbb{E}}(q)$. Consequently, by using $\overline{P}_{q,\widetilde{\mathbb{E}}(q)}$ to pick q^*, we adopt the most optimistic attitude towards the expert's possible answers, leading to the largest potential reduction in the number of maximal acts. Interestingly, such an optimistic view is at work in other selection settings such as racing algorithms [12].

To reduce cognitive burden on the expert, the condition $\mathcal{Q} \neq \emptyset$ could be further expanded, for instance, by simply limiting the total number of queries to be asked, or by considering that having almost indifferent options in $\mathcal{A} \leftarrow \text{opt}(\overline{P}, \mathcal{A})$ is sufficient to stop the procedure. As there are many sensible ways to do this, we have omitted this from the algorithm description.

Another aspect that is important is the set of queries $\mathcal{Q}$. One could for instance query the expert on the upper and lower probabilities of events $A \subseteq \Omega$, which correspond to indicator functions $\mathbb{I}_A$ in terms of gambles. For instance, asking about the lower probability of A is equivalent to asking about the upper prevision of $-\mathbb{I}_A$, and of $\mathbb{I}_A$ if one is interested in the upper probability. Other common queries include comparative ones between events A and B, e.g. $\mathbb{I}_A - \mathbb{I}_B$ [13, 7], or the direct comparisons of two alternatives within $\mathcal{A}$. There is clearly a balance to find between the expressivity of $\mathcal{Q}$, and the cognitive accessibility of the queries for the decision maker. However, exploring such a question is out of the scope of the current initial study.

Regarding the complexity of algorithm 1, computing eq. (11) in line 4 requires solving one linear program per item in dom $\overline{P}$, and this for each remaining item in $\mathcal{Q}$, meaning that the number of linear programs to solve roughly increases quadratically as we collect queries. We also assume that we can compute opt$(\overline{P}, \mathcal{A})$ and opt$(\overline{P}_{q,\widetilde{\mathbb{E}}(q)}, \mathcal{A})$, meaning that we are assuming that we can enumerate acts explicitly, and that those remain in reasonable number (say, less than 100). Of course, a lot of those computations can be parallelised to ensure scalability, yet considering combinatorial optimisation problems would still require some adaptations.

3.2 Fully worked example

We now demonstrate algorithm 1 through examples 2 to 4, where we iteratively make queries to the expert.

Example 2. Recall the situation from example 1 where we elicited an expert's knowledge about next year's inflection rate through an upper prevision $\overline{P}$. Based on this information, the maximal acts were 3, 4, 5 and 6. Consider the set of gambles $\mathcal{Q} = \{-\mathbb{I}_L, -\mathbb{I}_M, -\mathbb{I}_H\}$. We want to elicit the expert's upper prevision of one of these queries; this is equivalent to asking about lower probability of each atom. For each query, we can compute the coherent range using eqs. (3) and (11):

$$\widetilde{\mathbb{E}}(-\mathbb{I}_L) = -0.5 \qquad \overline{\mathbb{E}}(-\mathbb{I}_L) = -0.33 \tag{12}$$

$$\widetilde{\mathbb{E}}(-\mathbb{I}_M) = -0.33 \qquad \overline{\mathbb{E}}(-\mathbb{I}_M) = 0.0 \tag{13}$$

$$\widetilde{\mathbb{E}}(-\mathbb{I}_H) = -0.33 \qquad \overline{\mathbb{E}}(-\mathbb{I}_H) = 0.0 \tag{14}$$

So, for each question $q \in \mathcal{Q} = \{-\mathbb{I}_L, -\mathbb{I}_M, -\mathbb{I}_H\}$, we assume that the expert will state a value $\beta \in [\widetilde{\mathbb{E}}(q), \overline{\mathbb{E}}(q)]$ as their upper prevision for q. For example, for the atom L, any coherent lower probability must be between 0.33 and 0.5. If the expert's answer is not in this range, then the expert is incoherent. If they avoid sure loss, we may adjust the expert's answer through natural extension to make them coherent. If they do not avoid sure loss, we must ask the expert to revise their answer until they do avoid sure loss.

Here are two extreme cases. If the expert gives $\beta = \overline{\mathbb{E}}(q)$, then $\overline{\mathbb{E}} = \overline{\mathbb{E}}_{q,\beta}$, then no information is gained to assist the decision maker. For the other extreme case, if the expert gives $\beta = \widetilde{\mathbb{E}}(q)$, then this is the most informative case that reduces as much as possible the credal set while still being coherent.

We can see from table 3 that for $-\mathbb{I}_L$, all rows remain positive. Consequently, regardless of the expert answer, as long as the expert remains coherent, asking for their upper prevision of $-\mathbb{I}_L$ will not change the number of maximal elements. On the other hand, $-\mathbb{I}_M$ has more potential: table 4 shows that stock fund 3 may no longer be maximal since there is a negative value in the corresponding row. Finally, $-\mathbb{I}_H$ has the most potential: table 4 shows that stock funds 4 and 6 may no longer be maximal.

$a \backslash a'$	3	4	5	6
3	0.0	1.5	1.0	2.5
4	2.0	0.0	2.0	4.0
5	2.0	0.5	0.0	2.0
6	2.0	0.0	0.0	0.0

Table 3. Values of $\overline{\mathbb{E}}_{-\mathbb{I}_L,-0.5}(a-a')$ in Example 2.

$a \backslash a'$	3	4	5	6
3	0.0	0.33	−0.67	−0.33
4	2.0	0.0	0.67	1.33
5	2.0	1.0	0.0	0.67
6	2.0	0.67	0.0	0.0

Table 4. Values of $\overline{\mathbb{E}}_{-\mathbb{I}_M,-0.33}(a - a')$ in Example 2.

$a \backslash a'$	3	4	5	6
3	0.0	1.5	1.0	2.5
4	−0.33	0.0	0.33	2.0
5	0.67	1.0	0.0	1.67
6	0.33	0.67	−0.33	0.0

Table 5. Values of $\overline{\mathbb{E}}_{-\mathbb{I}_H,-0.33}(a - a')$ in Example 2.

In conclusion, it is not worth asking about $-\mathbb{I}_L$, but asking about $-\mathbb{I}_M$ or $-\mathbb{I}_H$ can potentially reduce the number of maximal acts. Since adding $-\mathbb{I}_H$ to the domain has the potentially largest reduction in the number of maximal acts, we will ask the expert for their upper prevision of $-\mathbb{I}_H$.

Suppose the expert specifies a value $\beta = -0.2$ for the upper prevision of $-\mathbb{I}_H$ (so a lower probability of 0.2 for H). So, we now have $\operatorname{dom} \overline{P} = \{\mathbb{I}_M - \mathbb{I}_L, \mathbb{I}_H - \mathbb{I}_L, -\mathbb{I}_H\}$ with

$$\overline{P}(\mathbb{I}_M - \mathbb{I}_L) = 0, \qquad \overline{P}(\mathbb{I}_H - \mathbb{I}_L) = 0, \qquad \overline{P}(-\mathbb{I}_H) = -0.2 \tag{15}$$

Next, we can calculate $\overline{\mathbb{E}}(a - a')$ based on this updated $\overline{P}$; see table 6. As we can see from table 6, the maximal acts are stock funds 3, 4, and 5.

$a \backslash a'$	3	4	5	6
3	0.0	1.5	1.0	2.5
4	0.6	0.0	1.0	2.8
5	1.2	1.0	0.0	1.8
6	1.0	0.67	−0.2	0.0

Table 6. Values of $\overline{\mathbb{E}}(a - a')$ after adding $-\mathbb{I}_H$ to $\operatorname{dom} \overline{P}$ in example 2.

Example 3. We can now continue the process by considering the remaining queries, that is, $\mathcal{Q} = \{-\mathbb{I}_L, -\mathbb{I}_M\}$. As before, we first compute the range of coherent extensions for these queries:

$$\widetilde{\mathbb{E}}(-\mathbb{I}_L) = -0.4 \qquad \overline{\mathbb{E}}(-\mathbb{I}_L) = -0.33 \tag{16}$$

$$\widetilde{\mathbb{E}}(-\mathbb{I}_M) = -0.33 \qquad \overline{\mathbb{E}}(-\mathbb{I}_M) = 0.0 \tag{17}$$

Again, for each $q \in \mathcal{Q}$, suppose that the expert will provide a value $\beta \in [\widetilde{\mathbb{E}}(q), \overline{\mathbb{E}}(q)]$. As before, we use the most optimistic case where $\beta = \widetilde{\mathbb{E}}(q)$ to choose a query. We compute the values for $\overline{\mathbb{E}}_{q,\widetilde{\mathbb{E}}(q)}(a - a')$, where $q = -\mathbb{I}_L$ and $q = -\mathbb{I}_M$ and present results in tables 7 and 8. Since all rows in table 7 are non-negative, it is not worth asking the expert about $-\mathbb{I}_L$. However, as there is a strictly negative number in table 8, asking about $-\mathbb{I}_M$ might make stock fund 4 non-maximal.

Suppose that the expert says $\overline{P}(-\mathbb{I}_M) = -0.2$. So, we now have $\operatorname{dom}\overline{P} = \{\mathbb{I}_M - \mathbb{I}_L, \mathbb{I}_H - \mathbb{I}_L, -\mathbb{I}_H, -\mathbb{I}_M\}$ with

$$\overline{P}(\mathbb{I}_M - \mathbb{I}_L) = 0, \quad \overline{P}(\mathbb{I}_H - \mathbb{I}_L) = 0, \quad \overline{P}(-\mathbb{I}_H) = -0.2, \quad \overline{P}(-\mathbb{I}_M) = -0.2. \tag{18}$$

We then can calculate $\overline{\mathbb{E}}(a - a')$ for the remaining acts and present them in table 9. Unfortunately, all entries in table 9 are non-negative, so the maximal acts remain the same.

$a \backslash a'$	3	4	5
3	0.0	1.5	1.0
4	0.6	0.0	1.0
5	1.2	0.8	0.0

Table 7. Values of $\overline{\mathbb{E}}_{-L,-0.4}(a - a')$ in example 3.

$a \backslash a'$	3	4	5
3	0.0	0.33	−0.67
4	0.6	0.0	−0.33
5	1.2	1.0	0.0

Table 8. Values of $\overline{\mathbb{E}}_{-M,-0.33}(a - a')$ in example 3.

$a \backslash a'$	3	4	5
3	0.0	0.8	0.0
4	0.6	0.0	0.2
5	1.2	1.0	0.0

Table 9. Values of $\overline{\mathbb{E}}(a - a')$ after adding $-M$ to $\operatorname{dom}\overline{P}$ in example 3.

$a \backslash a'$	3	4	5
3	0.0	0.8	0.0
4	0.6	0.0	0.2
5	1.2	0.8	0.0

Table 10. Values of $\overline{\mathbb{E}}_{-L,-0.4}(a - a')$ in example 4.

Example 4. Finally, we consider the last query $-\mathbb{I}_L$ and compute its range of coherent extensions:

$$\widetilde{\mathbb{E}}(-\mathbb{I}_L) = -0.4 \quad \overline{\mathbb{E}}(-\mathbb{I}_L) = -0.33. \tag{19}$$

A coherent expert will give a value $\beta \in [\widetilde{\mathbb{E}}(-\mathbb{I}_L), \overline{\mathbb{E}}(-\mathbb{I}_L)]$. Again, we consider the most optimistic case where $\beta = \widetilde{\mathbb{E}}(-\mathbb{I}_L)$ and compute the values for $\overline{\mathbb{E}}_{-\mathbb{I}_L,-0.4}(a - a')$; see table 10. Since all entries in table 10 are non-negative, no answer from the expert about $-\mathbb{I}_L$ will change the number of maximal elements. Therefore, we can stop the elicitation process here.

4 Conclusion

In this paper, we studied the problem of elicitation in the context of decision making. We aimed to gain the most useful information to efficiently identify queries for a specific decision problem. To do so, we treated the expert's information as an upper prevision and proposed a new algorithm to most effectively pick a query to present to an expert. Assuming the expert remains coherent at all times, we used the lower bound of the coherent range to select the query that has most potential to reduce the number of maximal acts. A fully worked example demonstrated the algorithm.

In future work, we might explore additional stopping criteria to terminate elicitation. For example, we may consider to terminate the elicitation process when we can no longer remove any non-maximal acts regardless of the next query. This could save cognitive burden on the expert since we can present fewer queries. Moreover, instead of using the lower bound of the coherent range, we can look at other values, such for instance a mixture between the upper and lower bounds. Similarly, it would be desirable to test the impact of choosing different sets of possible queries, both in terms of convergence to the desired numerical outcome, and in terms of easiness for the decision maker. Such comparisons would necessitate extensive experiments.

Finally, it would be interesting to test the procedure on an actual decision problem with an actual expert, for instance to see if our assumption of coherence of the expert is reasonable in practice. In order to do this, we could take inspiration from previous experimental campaigns aiming at eliciting actual imprecise probabilities [17].

Acknowledgments. This project is funded by National Research Council of Thailand (NRCT): Contract number N42A660964. This research was supported by Chiang Mai University.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Bibliography

[1] Adam, L., Destercke, S.: Possibilistic preference elicitation by minimax regret. In: Uncertainty in artificial intelligence. pp. 718–727. PMLR (2021)

[2] Augustin, T., Coolen, F.P.A., de Cooman, G., Troffaes, M.C.M.: Introduction to imprecise probabilities, vol. 591. John Wiley & Sons (2014)

[3] Benabbou, N., Gonzales, C., Perny, P., Viappiani, P.: Minimax regret approaches for preference elicitation with rank-dependent aggregators. EURO journal on Decision processes **3**, 29–64 (2015)

[4] Blanco, M., Engelmann, D., Koch, A.K., Normann, H.T.: elief elicitation in experiments: is there a hedging problem? Experimental Economics **13**, 412–438 (2010). https://doi.org/10.1007/s10683-010-9249-1

[5] Boutilier, C.: Computational decision support: Regret-based models for optimization and preference elicitation (2013)

[6] Coletti, G., Scozzafava, R.: The role of coherence in eliciting and handling imprecise probabilities and its application to medical diagnosis. Information Sciences **130**(1), 41–65 (2000). https://doi.org/10.1016/S0020-0255(00)00085-2

[7] Erreygers, A., Miranda, E.: A graphical study of comparative probabilities. Journal of Mathematical Psychology **104**, 102582 (2021)

[8] Garthwaite, P.H., Kadane, J.B., O'Hagan, A.: Statistical methods for eliciting probability distributions. Journal of the American statistical Association **100**(470), 680–701 (2005)

[9] Gilovich, T., Griffin, D., Kahneman, D.: Heuristics and biases: The psychology of intuitive judgment. Cambridge university press (2002)

[10] Jansen, C., Blocher, H., Augustin, T., Schollmeyer, G.: Information efficient learning of complexly structured preferences: Elicitation procedures and their application to decision making under uncertainty. International Journal of Approximate Reasoning **144**, 69–91 (2022). https://doi.org/10.1016/j.ijar.2022.01.016

[11] Kadane, J.B., Winkler, R.L.: Separating probability elicitation from utilities. Journal of the American Statistical Association **83**(402), 357–363 (1988)

[12] Maron, O., Moore, A.W.: The racing algorithm: Model selection for lazy learners. Artificial Intelligence Review **11**, 193–225 (1997)

[13] Miranda, E., Destercke, S.: Extreme points of the credal sets generated by comparative probabilities. Journal of Mathematical Psychology **64**, 44–57 (2015)

[14] Oakley, J.E., O'Hagan, A.: Uncertainty in prior elicitations: a nonparametric approach. Biometrika **94**(2), 427–441 (2007)

[15] O'Hagan, A., Buck, C.E., Daneshkhah, A., Eiser, J.R., Garthwaite, P.H., Jenkinson, D.J., Oakley, J.E., Rakow, T.: The Elicitation of Probabilities, chap. 4, pp. 61–96. John Wiley & Sons, Ltd (2006). https://doi.org/10.1002/0470033312.ch4

[16] O'Hagan, A., Buck, C.E., Daneshkhah, A., Eiser, J.R., Garthwaite, P.H., Jenkinson, D.J., Oakley, J.E., Rakow, T.: Uncertain judgements: eliciting experts' probabilities. John Wiley & Sons (2006)
[17] Quaeghebeur, E., Wesseling, C., Beauxis-Aussalet, E., Piovesan, T., Sterkenburg, T.: The cwi world cup competition: Eliciting sets of acceptable gambles. In: Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications. pp. 277–288. PMLR (2017)
[18] Renooij, S.: Probability elicitation for belief networks: issues to consider. The Knowledge Engineering Review **16**(3), 255–269 (2001). https://doi.org/10.1017/S0269888901000145
[19] Rinderknecht, S.L., Borsuk, M.E., Reichert, P.: Bridging uncertain and ambiguous knowledge with imprecise probabilities. Environmental Modelling & Software **36**, 122–130 (2012)
[20] Sandri, S.A., Dubois, D., Kalfsbeek, H.W.: Elicitation, assessment, and pooling of expert judgments using possibility theory. IEEE transactions on fuzzy systems **3**(3), 313–335 (1995)
[21] Silva, A.A., de Souza, F.C.: A protocol for the elicitation of imprecise probabilities. In: ISIPTA. vol. 1, pp. 315–321 (2005)
[22] Smithson, M.: Elicitation, chap. 15, pp. 318–328. John Wiley & Sons, Ltd (2014). https://doi.org/10.1002/9781118763117.ch15
[23] T'Joens, N.: Active Elicitation of Imprecise Probability Models. Master's thesis, Ghent University (2017), https://lib.ugent.be/catalog/rug01:002367397
[24] Troffaes, M.C.M.: Decision making under uncertainty using imprecise probabilities. International Journal of Approximate Reasoning **45**(1), 17–29 (May 2007). https://doi.org/10.1016/j.ijar.2006.06.001
[25] Troffaes, M.C.M., de Cooman, G.: Lower Previsions. Wiley Series in Probability and Statistics, Wiley (May 2014). https://doi.org/10.1002/9781118762622
[26] Utkin, L.V.: Probabilities of judgments provided by unknown experts by using the imprecise dirichlet model. Risk Decision and Policy **9**(4), 371–389 (2004)
[27] Walley, P.: Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London (1991)
[28] Williams, P.M.: Notes on conditional previsions. International Journal of Approximate Reasoning **44**(3), 366–383 (2007)
[29] Yaghlane, A.B., Denoeux, T., Mellouli, K.: Constructing belief functions from qualitative expert opinions. In: 2006 2nd International Conference on Information & Communication Technologies. vol. 1, pp. 1363–1368. IEEE (2006)