Incremental elicitation of preferences: optimist or pessimist?

Loïc Adam, Sébastien Destercke

Incremental elicitation of preferences: optimist or pessimist?

Loïc Adam[1] and Sébastien Destercke[1]

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

Abstract In robust incremental elicitation, it is quite common to make recommendations and to select queries by using a minimax regret criterion, which corresponds to a pessimistic attitude. In this paper, we explore its optimistic counterpart, showing this new approach enjoys the same convergence properties. While this optimistic approach does not offer the same kind of guarantees than minimax approaches, it still offers some other interesting properties. Finally, we illustrate with some experiments that the best approach amongst the two approaches heavily depends on the underlying setting.

Keywords: preferences elicitation, incremental, minimax regret, maximax utility

1 Introduction

Preference elicitation by interacting with an agent or a user is a crucial step to identify and formalise her preferences. While there are different ways to interact with a user, incremental elicitation [2] is a very interesting approach since each new question takes into account the preferential information provided previously. In the literature, one of the main approaches of incremental elicitation is the robust approach, based on a Minimax regret optimisation [4,3]. Provided their underlying hypotheses[1] are satisfied, the interest of using such approaches is that, due to their pessimistic stance (minimising the regret in the worst situation), they come with strong guarantees about the recommended alternative. They also converge in a reasonable number of steps to a good recommendation, as the space of possible models is guaranteed to shrink after each question. In this paper, we will work under the same hypotheses as the robust approaches to simplify our exposure, as we could easily adapt our proposal to extensions of the robust approach [10] which are able to deal with errors.

Minimax regret approach [14] is a popular choice for making decisions under uncertainty, as it minimises the worst-case regret. Such a decision rule provides

[1] The user is an oracle, and the chosen family of preference model includes the right model.

rather safe recommendations, corresponding to a situation where the agent is rather pessimistic on the outcomes of a decision, fearing a possibly rare but disastrous worst-case scenario. However, Minimax regret is only one amongst many other decision rules under uncertainty (see, e.g., [16] for an account of those), and as all other rules, it has drawbacks one may not appreciate. For instance, it is sensitive to the addition of irrelevant alternatives, and does not guarantee potentially optimal recommendations, i.e. recommendations that are the best for at least one particular model within the set of possible models.

In this paper, we consider a somewhat opposed view, using a Maximax optimist approach to make recommendations. There are multiple reasons to investigate such an alternative: one is that such optimistic approaches in presence of uncertainty are often used in learning under uncertainty, for example to deal with missing data [6] or to identify optimal models [13], hence adopting such an optimistic view in preference elicitation that shares many similarities with the aforementioned learning setting seems relevant; another is that optimistic recommendations do not suffer the same drawbacks we have mentioned for the Minimax regret[2], hence may be acceptable in situations where the Minimax regret is not. Last but not least, such an approach can be computationally more efficient than regret based ones, since regret typically involves comparing pairs of alternatives, whereas Maximax decision typically involves alternative-wise computations.

We introduce in Section 2 all the necessary elements in preference elicitation to understand our work. We then present in Section 3 an optimist approach that we call Maximax gain, adapting the Current Solution Strategy (CSS) heuristic. We also discuss some of its interesting properties. Lastly, Section 4 shows some simulated experiments whose goal are to investigate whether there are situations in which an optimist approach also increases recommendation performances.

2 Preliminaries

In this section, we introduce the various elements necessary to understand the rest of our work. We also introduce a running example that we will repeatedly use to illustrate the introduced notions.

2.1 Notations

Alternatives We define $\mathbb{X}$ as the finite set of available alternatives. Alternatives within $\mathbb{X}$ are denoted $x_1, x_2, ..., x_k$. We assume that alternatives are summarised by q real values, the criteria, such that $x \in \mathbb{R}^q$. The *ith* criterion value of an alternative $x \in \mathbb{X}$ is denoted x^i. Given two alternatives $x, y \in \mathbb{X}$, we denote:

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

[2] It does not, however, offer the same robust guarantees.

Example 1 (Choosing the best sandwich (running example)) *We imagine that a user wants to choose the best possible sandwich among multiple ones (the alternatives). Each sandwich is characterised by two criteria: flavour and price. Each criterion is valued between 0 (worst) and 10 (best). Table 1 lists the available sandwiches. We can see for instance that cheese sandwich is very cheap but with a mediocre flavour, while duck sandwich is full of flavour yet overpriced.*

Table 1. Grades of sandwiches

	Flavour	1/price
Cheese	5	9
Duck	10	0
Fish	8	4
Ham	7	7

Aggregation models We consider that each alternative x is valuated by its utility, and that the utility depends on the preferences of the agent. We also suppose that such an utility is modelled by a function $f_\omega(x)$, parameterised by $\omega \in \Omega$, aggregating the different criteria. ω is also known as the preferential model. Given this evaluation function f_ω, it is possible to compare two alternatives $x, y \in \mathbb{X}$:

$$x \succeq_\omega y \iff f_\omega(x) \geq f_\omega(y). \tag{1}$$

A first model we consider is the weighted sum (WS) model. This is a very simple model, which can be considered as the basic building block of decision theory and is still widely used in multi-criteria decision-making. Given a vector of weights $\omega = \{\omega^1, ..., \omega^q\} \in \mathbb{R}^q$, we have:

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

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

We also consider the Ordered weighted averaging (OWA) model [17]. This model generalises aggregation operators such as the arithmetic mean, median, min or max. With an OWA model, criteria values are ordered increasingly. Given a vector of weights $\omega \in \mathbb{R}^q$ and the ordered criteria values $x^{(1)} \leq ... \leq x^{(q)}$, we have:

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

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

These two models are simple and can be used in most situations. More complex models like Choquet integrals [7,9] exist that can be used, e.g., to model interactions between criteria. Provided the models are linear in ωonce x is fixed, Equation (1) is equivalent to a linear constraint, meaning that we can use them within a linear program [2]. All mentioned models so far are linear in ω.

Example 2 (Application of models) *Given the sandwiches we presented in Table 1, we assume the user evaluates each sandwich with a WS model such that* $\omega = (0.8, 0.2)$*. This means she values the flavour over the price. She then prefers the duck sandwich, as it scores 8, over the fish sandwich, with a score of 7.2.*

If she evaluates with an OWA model such that $\omega = (0.8, 0.2)$*, meaning she penalises a sandwich that is bad on at least one criterion, she will prefer the more balanced fish sandwich:* $f_\omega(\mathit{fish}) = 4 \times 0.8 + 8 \times 0.2 = 4.8$ *while* $f_\omega(\mathit{duck}) = 0 \times 0.8 + 10 \times 0.2 = 2$.

2.2 Robust elicitation with Minmax regret

Motivation Finding a unique model ω from pairwise comparisons is difficult. However, it is often possible to draw reasonable inferences without complete information. Robust recommendation approaches aim at identifying a subset Ω' of possible models ω from preferential information. We then identify the preferences that hold for every model $\omega \in \Omega'$. This results in a partial preorder over $\mathbb{X}$ where:

$$x \succeq^{\Omega'} y \Longleftrightarrow \forall \omega \in \Omega' \; f_\omega(x) \geq f_\omega(y). \tag{4}$$

A good elicitation strategy needs to reduce Ω' as quickly as possible to make good recommendations without exhausting the budget of questions. Such a strategy should also make good recommendations even if $\succeq^{\Omega'}$ does not have a single maximal element, as in practice information collection may end before that.

Regret based elicitation Minmax regret is a well-known notion for decision problems under uncertainty and set-valued information [14]. It still provides strong guarantees on the recommendation quality, while being less conservative than standard Minmax.

We are now introducing the different measures to compute the Minmax regret. The regret of choosing an alternative x over the alternative y given a model ω is defined by:

$$\mathrm{R}_\omega(x, y) = f_\omega(y) - f_\omega(x). \tag{5}$$

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

$$\mathrm{PMR}(x, y, \Omega') = \max_{\omega \in \Omega'} \mathrm{R}_\omega(x, y), \tag{6}$$

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

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

which is the regret of choosing x in the worst case scenario, i.e., considering the worst model for its strongest opponent.
Finally, the min max regret of a set $\mathbb{X}$ of alternatives given a set Ω' of possible models is:

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

In this approach $x^* = \arg \mathrm{mMR}(\Omega')$ is the alternative giving the minimal regret in a worst-case scenario, and is the current recommendation if no further information can be collected.

Example 3 (Initial choice with a Minmax regret) *We want to pick the alternative which minimises the maximum regret in the worst-case scenario, the preferences being evaluated with a WS model. The evolution of the score of each alternative depending on the parameter $\omega^{1/price}$ is depicted on Figure 1. The alternative which minimises the maximum regret is the ham sandwich, with MR(ham) = 3 when we pick it instead of the duck sandwich for $\omega^{1/price} = 0$ (meaning only the flavour is considered).*

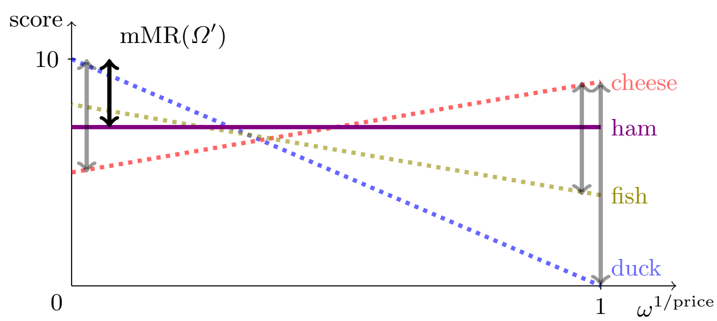

Figure 1. Choice of the best alternative given a Minmax approach

2.3 Elicitation sequence and regret CSS

Preferential information is often collected through pairwise comparison: we present a pair (x, y) to the user, and she tells which one she prefers. We will denote by

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

the subset of models consistent with the assessment $x \succeq y$, and $\omega^{y \succeq x}$ the subset for $y \succeq x$. In an elicitation sequence, we alternatively present a pair to the user, and update the information with the answer. In the robust approach, if Ω^k is the possible subset of models at the kth step, the next step is to present a couple (x, y) to the user, and then compute $\Omega^{k+1} = \Omega^k \cap \omega^{x \succeq y}$ if the user prefers x, $\Omega^{k+1} = \Omega^k \cap \omega^{x \preceq y}$ otherwise.

Choosing a good pair (x, y) is therefore a critical step. We consider for our work the well-known CSS strategy [5], where given a subset Ω', the user compares the current regret-based recommendation $x^* = \arg \mathrm{mMR}(\Omega')$ (so our best option w.r.t this criterion) to its worst opponent:

$$y^* = \arg \max_{y \in \mathcal{X}} \mathrm{PMR}(x^*, y, \Omega'). \tag{10}$$

This heuristic strategy provides good results in general, and guarantees that the updated set will be non-empty.

Example 4 (Updating the model space) *We assume the user decides with a WS model, where* $\omega^* = (0.8, 0.2)$*. In a first question* q_1*, she has to choose her favourite sandwich among the pair* (x_D, x_F)*. We have already shown in Example 2 that she prefers* x_D $(f_\omega(x_D) = 8,\ f_\omega(x_F) = 7.2)$*. We assume she answers correctly that* $x_D \succeq x_F$*. We then have* Ω' *the set of models consistent with her known preferences, such that* $\Omega' = \omega^{x_D \succeq x_F} = \{\omega \in \Omega : \sum_{i=1}^{2} \omega^i.(x_D^i - x_F^i) \geq 0\} = \{\omega \in \Omega : \omega^1 \geq 2\omega^2\}$*, where* ω^1 *corresponds to the flavour, and* ω^2 *to 1/price. The updated model space* Ω^1 *is shown on Figure 2.*

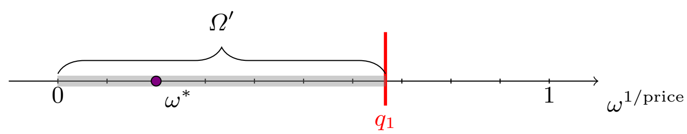

Figure 2. Update of the model space Ω

With more questions, it is possible to further update the subset of possible models Ω' *consistent with the preferences of the user.*

An important property of robust approaches combined with CSS is that, by construction, they guarantee that the elicitation sequence will converge , as we remind here:

Proposition 1 *[1,5]Given* $\Omega^{k+1} \subseteq \Omega^k$*, the sets of possible model at steps* k *and* $k+1$*, we have that:*

$$PMR(x, y, \Omega^k) \geq PMR(x, y, \Omega^{k+1}), \tag{11}$$

$$MR(x, \Omega^k) \geq MR(x, \Omega^{k+1}), \tag{12}$$

$$mMR(\Omega^k) \geq mMR(\Omega^{k+1}). \tag{13}$$

Proof. PMR. Suppose we have a function f and two sets Ω, Ω' such that $\Omega' \subseteq \Omega$. We have $\max_{x \in \Omega} f(x) \geq \max_{x \in \Omega'} f(x)$, the maximum of Ω being either in Ω' or in $\Omega \setminus \Omega'$. We can replace f by the PMR, Ω by Ω^k and Ω' by Ω^{k+1} since $\Omega^{k+1} \subseteq \Omega^k$. (11) is then proved. Proof for MR and mMR directly follows, as they are maximum and minimum taken over decreasing values.

3 Optimist approach

Minmax regret is based on a pessimist decision rule: the user wants the alternative that minimises the maximum loss, i.e., in the worst-case scenario. The user is risk-averse and does not mind if the gain is lower on average. However, it is unclear in a preference framework that the user will always be risk-averse, rather than opportunity-seeking. This is why we now consider the Maximax gain approach and its direct CSS adaptation, that considers recommendations based on another decision rule, where the user wants to maximise their gain in the best-case scenario. The choice of the Maximax gain approach can be justified in various ways: we show later in this section that an alternative suggested by a Maximax gain approach is the best possible for at least one situation (one model ω). This is not necessary the case with a Minmax regret, meaning that a risk-adverse user may actually select an option known to be necessarily sub-optimal. It does not mean that one strategy is better than the other, just that they have different properties, as the choice can depend on the willingness of the user for taking risks to maximise their possible gain. In this section, we discuss why such an approach may be an interesting alternative to Minmax regret approaches.

3.1 Robust elicitation with Maximax gain

Given an alternative x, the maximal gain over a set Ω' of possible models is:

$$\mathrm{MG}(x, \Omega') = \max_{\omega \in \Omega'} f_\omega(x), \tag{14}$$

which corresponds to the gain in the best-case scenario. Given a set of alternatives $\mathbb{X}$ and the set Ω', the max maximal gain is:

$$\mathrm{MMG}(\Omega') = \max_{x \in \mathbb{X}} \mathrm{MG}(x, \Omega'), \tag{15}$$

$x^* = \arg \mathrm{MMG}(\Omega')$ is the alternative giving the maximal possible gain in the corresponding best-case scenario. As with Minmax regret, x^* is the current recommendation if no additional information can be collected. When it comes to choosing the question, we still retain the CSS heuristic approach that chooses $y^* = \arg\max_{y \in \mathbb{X}} \mathrm{PMR}(x^*, y, \Omega')$ as an adversary. For convenience, we will refer to the corresponding elicitation as *gain CSS*.

Example 5 (Initial choice with an optimist approach) *In example 3 we have shown how to pick the best alternative based on a Minmax regret, when we have no information on the preferences of a user. We will now find the best alternative based on a Maximax gain, and show it can be different from the one proposed with the Minmax regret, and the preferences are still evaluated with a WS model.*

As shown on figure 3, the maximum gain obtainable for the duck sandwich is 10 when $\omega = (1, 0)$. *We also deduce that* $MG(x_F, \Omega) = 8$ *for* $\omega = (1, 0)$, *and* $MG(x_C, \Omega) = 9$ *for* $\omega = (0, 1)$. *The maximum gain of the ham sandwich is a particular case, since* $MG(x_H, \Omega) = 7\ \forall \omega \in \Omega$.

We then conclude that $MMG(\Omega) = 10$ *and that* $x^* = x_D$. *Pessimist and optimist approaches give different current solutions. The duck sandwich is a great candidate for maximising the gain in the best case scenario* $\omega = (0, 1)$, *but this alternative is the worst if* $\omega^{price} \gtrapprox 0.36$.

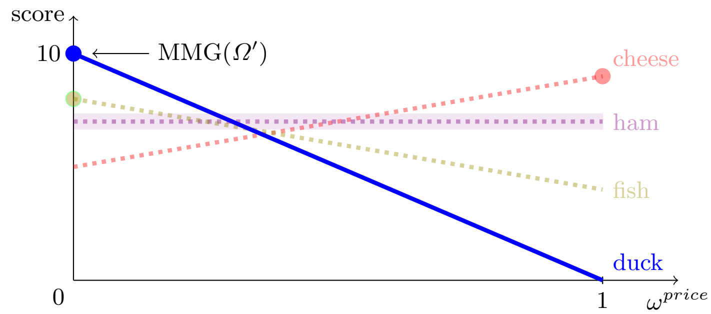

Figure 3. Choice of the best alternative given a Maximax approach

3.2 Optimality

We first introduce the notion of possibly Ω'–optimal solutions [3]. Given the set $\mathbb{X}$ and a subset $\Omega' \subseteq \Omega$ of possible models, the set $PO_{\Omega'}$ is defined by:

$$PO_{\Omega'} = \{x : \exists \omega \in \Omega', x \in \arg\max_{\mathbb{X}} f_w(y)\}. \tag{16}$$

In other words, an alternative $x \in \mathbb{X}$ is possibly Ω'–optimal if x is the best alternative for at least one model $\omega \in \Omega'$. An optimist robust elicitation, based on a Maximax gain, is interesting for the following property, that shows that the recommended item could be the best (not guaranteed by a Minmax approach):

Proposition 2 *The Maximax gain alternative $x^* = \arg MMG(\Omega')$ is possibly Ω'–optimal.*

Proof. Consider the model ω for which is obtained $x^* = \max_{x\in\mathbb{X}} [\max_{\omega\in\Omega'} f_\omega(x)]$. It is clear that for this model which is within Ω', x^* is the best alternative, hence it is possibly Ω'–optimal.

In Figure 4, x_D, x_C, and a modified version of x_H equals to $(0.6, 0.6)$ noted x_{H^*} are displayed; x_{H^*} being the Minmax regret recommendation. As we can see, such approaches cannot be expected to satisfy Proposition 2. On the other side, one can see in Figure 4 that the Maximax gain recommendation can be a very bad choice in some situations, meaning that checking whether we are in such situations may be of importance.

3.3 Convergence

In the case of regret CSS, the elicitation provably converges to the optimal model as long as no errors are made. It is guaranteed that after the kth update, $\Omega^k \subseteq \Omega^{k-1}$ will never be empty, and that the inclusion will be strict under mild assumptions. In the case of gain CSS, we have the same property:

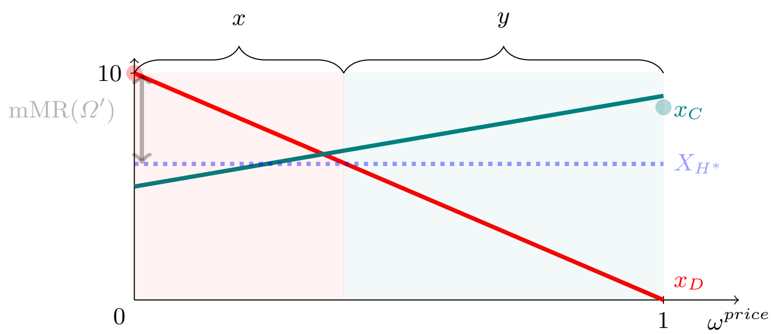

Figure 4. Visualisation of optimality of 3 alternatives

Proposition 3 *Consider the pair* $x^k = \arg MMG(\Omega^k)$, $y^k = \arg\max_{\mathbb{X}} PMR(x^k, y, \Omega^k)$ *chosen at the kth step of gain CSS. Then we have:*

$$\Omega^k \cap \omega^{x^k \succeq y^k} \neq \emptyset,$$

$$\Omega^k \cap \omega^{x^k \preceq y^k} \neq \emptyset.$$

Guaranteeing that $\Omega^k \subseteq \Omega^{k-1}$.

Proof. We prove that whatever the answer, the intersection with Ω^k is non-empty:

- $\Omega^k \cap \omega^{x^k \succeq y^k} \neq \emptyset$: immediate since $x^k \in PO_{\Omega^k}$, by Proposition 2, indicating that there is a model $\omega \in \Omega^k$ such that $f_\omega(x^k) \succeq f_\omega(y^k)$.
- $\Omega^k \cap \omega^{x^k \preceq y^k} \neq \emptyset$: since we picked y^k in accordance with the CSS, Equation (10) tells us that $y^k = \arg\max_{y \in \mathcal{X}} \mathrm{PMR}(x^k, y, \Omega^k)$. Since $\mathrm{PMR}(x^k, y^k, \Omega^k) = \max_{\omega \in \Omega^k} \left[f_\omega(y^k) - f_\omega(x^k)\right] \geq 0$, if follows that there is a $\omega \in \Omega^k$ (e.g., the one for which the PMR is reached) with $f_\omega(y^k) \geq f_\omega(x^k)$, ending the proof.

This proposition tells us that our space of possible models will shrink after each question in a non-degenerate way, guaranteeing us to converge to the true model. Note that we have a strict inclusion, i.e., $\Omega^k \subset \Omega^{k-1}$ relation if the $\arg MMG$ is unique, and if the corresponding PMR is strictly positive.

3.4 Computation complexity

Another advantage of an optimist approach based on a Maximax gain is its lower computational complexity. With a pessimist approach based on a Minmax regret, computing the PMR for all the possible pairs (x, y) such that $x, y \in \mathbb{X}$ and $x \neq y$ is equivalent to solving $n^2 - n$ linear optimisation problems, where $n = |\mathbb{X}|$.

With an optimist approach, computing the MG for all the alternatives $x \in \mathbb{X}$ is equivalent to only solving $2n - 1$ linear optimisation problems (n for the MMG and $n - 1$ for the PMR between x and the other alternatives). The linear optimisation problems in both approaches are similar. This lower complexity cost is very interesting for decision problems with a large set of alternatives.

4 Experiments

While Section 3 did present some interesting properties of adopting a gain CSS approach rather than a regret one, we need to check if this alternative heuristic can provide reasonable, if not better, performances.

This section brings some elements of answer, by comparing regret and gain CSS strategies in different situations.

4.1 Experimental protocol

We performed numerical simulations to compare the performances of regret and gain CSS approaches in different contexts, in which we change the kinds of alternatives we consider, as well as the kind of models.

The first element of comparison we consider is the choice of the alternatives:

- In a first setting, we generated randomly multiple alternatives x_i with 8 criteria from a uniform distribution, such that $x_i \in [0,1]^8$, until we obtained a Pareto front of 100 alternatives. In this case, most alternatives have quite high values on the different criteria, and we consider them to be *good alternatives*
- In a second setting, we generated randomly 100 alternatives x_i with 8 criteria from a Dirichlet distribution, such that $x_i \in [0,1]^8$ and $\sum_{j=1}^{8} x_i^j = 1$. We then have alternatives whose average utility is the same, and on which trade-offs have to be made. Since such alternatives are poorly noted (average utility of 1 when 8 is the best), we consider them to be *bad alternatives*

The second element of comparison is the choice of a function f_ω to estimate the utility of an alternative. We compared both approaches with 4 different functions, all generated randomly from different Dirichlet distributions:

- WSB: a "balanced" weighed sum (WS), where all criteria values are close. Parameters: $\alpha = 1000.(1/8, 1/8, ..., 1/8)$.
- WSU: an "unbalanced" weighed sum (WS), where some criteria can have significantly higher values than the others. Parameters: $\alpha = (1/8, 1/8, ..., 1/8)$.
- OWAU: an "unfair" OWA, which favours the criteria with the higher values. Parameters: $\alpha = 50.(1/36, 2/36, ..., 8/36)$.
- OWAR: a "redistributive" OWA, which favours the criteria with the lowest values. Parameters: $\alpha = 50.(8/36, 7/36, ..., 1/36)$.

We propose two measures for evaluating the prediction quality of each approach. A first measure is the real score of the current recommendation, computed from the hypothetical true preference model of the user. A second measure is the position of the current recommendation compared to the other alternatives, given the real score of each alternative. 0 means we have the best alternative and 99 the worst one.

We also reduce the variability of the two measures by averaging them on 200 simulations, and by computing a confidence interval of 95%.

4.2 Results

This section discusses the results of our experiments. Since displaying all graphs renders the reading difficult, we only display some of them, picturing different behaviours. Some synthetic statistics on all cases are given in Tables 2 and 3.

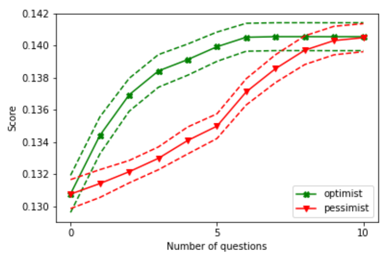

5.a: score

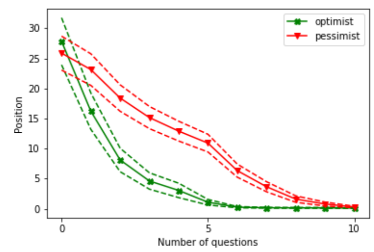

5.b: position

Figure 5. Score and position with poor alternatives on a balanced WS model

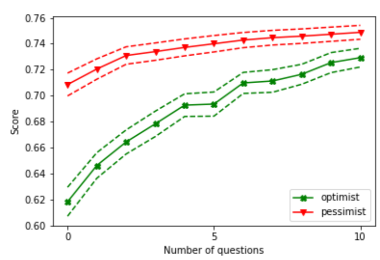

6.a: score

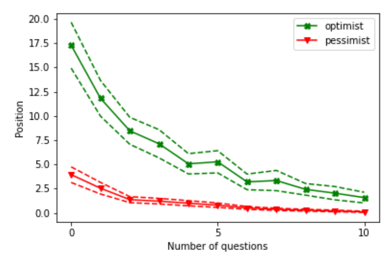

6.b: position

Figure 6. Score and position with good alternatives on a balanced WS model

On one extreme of the spectrum, we can see on Figures 5 and 6 the score and position for the balanced WS. In this case, the superiority of a method highly depends on the kind of available alternatives. On the other end of the spectrum, we can see on Figures 7 and 8 the results for the fair OWA model. While there is a slight advantage for the regret CSS strategy, it is not remarkable, and even not significant in the case of poor alternatives.

Tables 2 and 3 provide synthetic information about the different settings. Regarding the case of poor alternatives, the Gain CSS approach seems to give overall either significantly better or similar performances across the different scenarios. However, as indicated on Figures 5 to 8, both methods tend to quickly

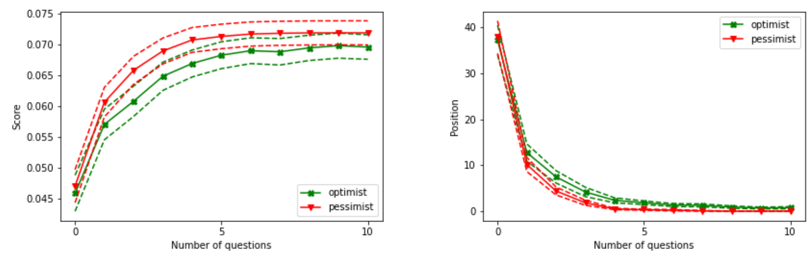

7.a: score 7.b: position

Figure 7. Score and position with poor alternatives on a fair OWA model

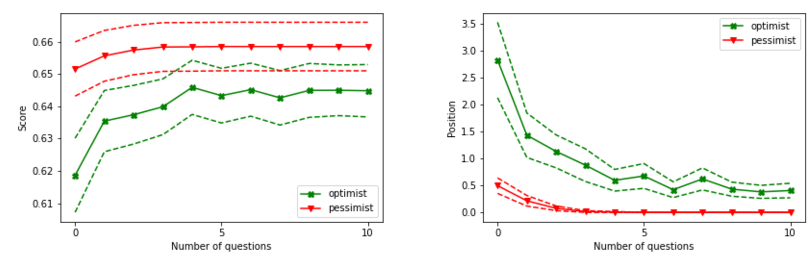

8.a: score 8.b: position

Figure 8. Score and position with good alternatives on a fair OWA model

converge to the same result, and provide essentially the same quality after 15 questions. If we go into more details about the case of poor alternatives:

- On balanced and unbalanced models (WSB and WSU), gain CSS is significantly more interesting than regret CSS.
- On unfair models (OWAU), gain CSS finds the best alternative only after one or two questions, which is very interesting. However, the regret CSS finds it after around 5 questions, and the score difference is non-significant whatever the number of questions is.
- On redistributive models (OWAR), gain CSS is usually a bit less effective than regret CSS. However, the differences are very small.

Regarding the case of good alternatives, the gain CSS appears overall less effective than regret CSS. On all the experiments but one, our optimist approach is slower to find a good solution. Again, both approaches converge to the best alternative after some questions. Let us now give a bit more details about the two approaches in the case of good alternatives:

- On balanced models, gain CSS is significantly worse than regret CSS.
- On unbalanced models, gain CSS is slightly worse than regret CSS, but they quickly converge to the same recommendation (after 7 or 8 questions).

Table 2. Current solution score after 5 questions on different contexts

	Poor alternatives				Good alternatives			
Method	WSB	WSU	OWAU	OWAR	WSB	WSU	OWAU	OWAR
Optimist/Gain	**0.140**	**0.607**	**0.231**	0.068	0.693	0.881	0.846	0.643
Pessimist/Regret	0.135	0.451	**0.231**	**0.071**	**0.740**	**0.887**	**0.853**	**0.659**

Table 3. Current solution position after 5 questions on different contexts

	Poor alternatives				Good alternatives			
Method	WSB	WSU	OWAU	OWAR	WSB	WSU	OWAU	OWAR
Optimist/Gain	**1.075**	**2.165**	**0.81**	1.86	5.27	2.91	0.475	0.675
Pessimist/Regret	10.91	7.505	1.335	**0.34**	**0.795**	**2.065**	**0**	**0**

– On unfair and redistributive models, gain CSS is slightly slower to find the best solution. The difference in score and position is small, yet significant.

4.3 Summary

The performances of gain CSS compared to the performances of regret CSS are qualitatively summarised on Table 4.

The main conclusion to draw from this table is that both the nature of the true underlying model, and the values of the alternatives, may have a huge impact on the results.

Table 4. Performance interest of gain CSS compared to regret CSS (++: quite interesting, −−: quite uninteresting)

	Poor alternatives	Good alternatives
Balanced	++	−−
Unbalanced	++	−
Unfair	∼	−
Redistributive	∼	−

In our opinion, this observation has two important impacts: the first is that how simulations are carried out in the validation of preference elicitation methods can have a huge impact on the results of these simulations, calling both for deeper theoretical studies about the situations in which a given heuristic has chances to work better, and for simulations considering large spectrum of situations.

Those results also show that in absence of strong inductive bias or refined knowledge about the alternatives, choosing one elicitation technique can hardly be based on performance requirements, and should therefore focus on which axioms should be satisfied in a given problem.

5 Conclusion

We studied the use of an optimist approach using a Maximax gain criterion for recommending an alternative in a robust preference elicitation, instead of a pessimist approach using a Minimax regret. We demonstrated that such an optimist approach possesses the same convergence properties as the classical regret-based one, and has interesting optimality and computational properties. Experiments on simulated data have shown that an optimist approach can be more effective in some contexts.

Our work has shown that the choice of the alternatives has some impacts on the performances of both approaches. We believe that it could be interesting to study more precisely the influence of the alternatives on the computation of the regret. This could be useful for determining the best strategy to choose alternatives in future works.

We cannot therefore give a definite answer to the question we asked in the title. Section 3 gives some pros and cons in terms of properties and axioms that are similar to the pros and cons of optimist and pessimist approaches one can find in other settings [12]. However, selecting the strategy using a performance requirement clearly calls for more theoretical studies regarding the situations in which different heuristics will perform better. A more empirical way to solve this issue could be to characterise elicitation problems through various quality measures (see, e.g., [8]), and see if we can predict the optimal/winning strategy from that, taking inspiration from machine learning methods [15]. Finally, let us note that while we looked at the problem with a greedy approach, whose interest is its efficiency and its agnosticity w.r.t. to the remaining number of questions, it may also be interesting (but also far more difficult) to consider the sequential version of our decision problem (see, e.g., [11]).

References

1. Benabbou, N., Gonzales, C., Perny, P., Viappiani, P.: Minimax regret approaches for preference elicitation with rank-dependent aggregators. EURO journal on Decision processes **3**(1), 29–64 (2015)
2. Benabbou, N., Perny, P., Viappiani, P.: Incremental elicitation of choquet capacities for multicriteria choice, ranking and sorting problems. Artificial Intelligence **246**, 152–180 (2017)
3. Bourdache, N., Perny, P.: Anytime algorithms for adaptive robust optimization with owa and wowa. In: International Conference on Algorithmic DecisionTheory. pp. 93–107. Springer (2017)
4. 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)
5. Boutilier, C., Patrascu, R., Poupart, P., Schuurmans, D.: Constraint-based optimization and utility elicitation using the minimax decision criterion. Artificial Intelligence **170**(8-9), 686–713 (2006)
6. Cabannnes, V., Rudi, A., Bach, F.: Structured prediction with partial labelling through the infimum loss. In: International Conference on Machine Learning. pp. 1230–1239. PMLR (2020)

7. Choquet, G.: Theory of capacities. In: Annales de l'institut Fourier. vol. 5, pp. 131–295 (1954)
8. Ciomek, K., Kadziński, M., Tervonen, T.: Heuristics for prioritizing pair-wise elicitation questions with additive multi-attribute value models. Omega **71**, 27–45 (2017)
9. Grabisch, M.: The application of fuzzy integrals in multicriteria decision making. European journal of operational research **89**(3), 445–456 (1996)
10. Guillot, P.L., Destercke, S.: Preference elicitation with uncertainty: Extending regret based methods with belief functions. In: International Conference on Scalable Uncertainty Management. pp. 289–309. Springer (2019)
11. Holloway, H.A., White Iii, C.C.: Question selection for multi-attribute decision-aiding. European Journal of Operational Research **148**(3), 525–533 (2003)
12. Hüllermeier, E., Destercke, S., Couso, I.: Learning from imprecise data: adjustments of optimistic and pessimistic variants. In: International Conference on Scalable Uncertainty Management. pp. 266–279. Springer (2019)
13. Maron, O., Moore, A.W.: The racing algorithm: Model selection for lazy learners. Artificial Intelligence Review **11**(1), 193–225 (1997)
14. Savage, L.J.: The theory of statistical decision. Journal of the American Statistical association **46**(253), 55–67 (1951)
15. Tornede, A., Wever, M., Hüllermeier, E.: Extreme algorithm selection with dyadic feature representation. In: International Conference on Discovery Science. pp. 309–324. Springer (2020)
16. Troffaes, M.C.: Decision making under uncertainty using imprecise probabilities. International journal of approximate reasoning **45**(1), 17–29 (2007)
17. Yager, R.R.: On ordered weighted averaging aggregation operators in multicriteria decisionmaking. IEEE Transactions on systems, Man, and Cybernetics **18**(1), 183–190 (1988)