ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230571

On the Variation of Max Regret with Respect to the Scaling of the Objectives

Nic Wilson ;*

Insight Centre, School CS and IT, University College Cork, Ireland
ORCiD ID: Nic Wilson https://orcid.org/0000-0003-1874-8255

Abstract. In a multi-objective optimisation problem, when there is uncertainty regarding the correct user preference model, max regret is a natural measure for how far an alternative is from being necessarily optimal (i.e., optimal with respect to every candidate preference model). It can be used for recommending a relatively safe choice to the user, or used in the generation of an informative query, and in the decision to terminate the user interaction, because an alternative is sufficiently close to being necessarily optimal. We consider a common and simple form of user preference model: a weighted average over the objectives (with unknown weights). However, changing the scale of an objective by a linear factor leads to an essentially different set of preference models, and this changes the max regret values (and potentially their relative ordering), sometimes very considerably. Since the scaling of the objectives is often partly subjective and somewhat arbitrary, it is important to be aware of how sensitive the max regret values are to the choices of scaling of the objectives. We give mathematical results that characterise and enable computation of this variability, along with an asymptotic analysis.

1 Introduction

With a multi-objective optimisation problem, in a situation where there is partial knowledge about the user preference model, it is common to use max regret as a measure of how far an alternative is from being necessarily optimal, i.e., optimal with respect to each compatible preference model. An alternative that minimises max regret can seem like a natural alternative to recommend to the user, given the current partial state of information. Also, the max regret measure is valuable in the generation of an informative query to ask the user, one that will improve the decision analyst system's state of knowledge about the user preferences, whichever way the user answers.

However, the value of max regret depends on the scaling (in particular, units) chosen for each objective. For example, an objective representing monetary cost might be expressed in euros, or alternatively in cents; as we will see in the example in the next section, changing the units can make a very large difference to max regret, and can change which alternative minimises max regret. Furthermore, the choices of the units/scalings of the objectives is often partly subjective and somewhat arbitrary.

In this paper we explore how max regret values vary with the objectives scaling, and we give mathematical results that characterise and enable computation of this variability, along with an asymptotic analysis.

* Email: nic.wilson@insight-centre.org.

Contributions: The methods developed in this paper enable a max regret approach for weighted average user preference models that is robust with respect to choices of scales of the objectives. We characterise the effect of rescaling the objectives, and this enables computation of the bounds on the max regret, given a range of rescaling vectors, and the determination of when the ratio of max regret to minimax regret can become arbitrarily large.

Related work: The max regret and minimax regret measures [18] have often been used for decision problems under uncertainty, as non-Bayesian methods for reasoning about an unknown user preference model; in particular, for recommending alternatives, in generating queries, and in deciding when to terminate the user interaction, see e.g., [23, 3, 5, 12, 1, 22]. One natural way of computing max regret is using extreme points [19, 20]. Weighted average and other linear preference models are a very commonly used special case of Multi-Attribute Utility Theory (MAUT) [16, 11, 7]. As a model for unknown user preferences they have been used for instance in [9, 6, 21, 17, 13, 10, 2]. The effect of rescaling objectives in the context of uncertain preferences has previously been analysed for a support vector machine-based approach to learning preferences in [14, 15].

Structure of the paper: Section 2 gives a motivating example, and discusses the issue of choosing the scaling for an objective. Section 3 gives some technical background regarding max regret, and Section 4 determines the effect of rescaling the objectives. Section 5 characterises when ratios of max regrets can become arbitrarily large. Section 6 derives absolute upper and lower bounds. Section 7 considers the effect of rescaling the objectives when one makes a different assumption on the user preference model: instead of being normalised, it is restricted so that each component is in the interval $[0, 1]$. Section 8 concludes.

2 Motivating Example

Consider a simple example with $p = 2$ objectives, the second objective representing gain in euros (for example, expected profit per unit sold), and the first objective being a measure of reputation gain/loss, where alternatives are different versions of a product to be manufactured. We have a single elicited preference that alternative $\alpha = (20, 2)$ is at least as good as alternative $\delta = (24, 1)$. A non-negative vector $w \in I\!R^p$ (called a *weights vector*) is used to represent a candidate user preference model, and the utility of an alternative $\gamma \in I\!R^p$ with respect to a weights vector w is defined to be $\gamma \cdot w = \sum_{i=1}^p \gamma(i)w(i)$. Thus, our initial information tells us that $w \cdot \alpha \geq w \cdot \delta$, i.e., $w \cdot \lambda \geq 0$, where $\lambda = \alpha - \delta = (-4, 1)$.

The max regret measure requires that the weights vectors are bounded, and it seems very natural to assume a normalisation condition $\sum_{i=1}^p w(i) = 1$. The set $\mathcal{W}_\Lambda^{\mathbf{1}}$ of feasible normalised preference models (or scenarios) based on input preferences $\Lambda = \{\lambda\}$ then consists of all normalised non-negative weights vectors w such that $w \cdot \lambda \geq 0$; thus, $w \in \mathcal{W}_\Lambda^{\mathbf{1}}$ if and only if $w(1) + w(2) = 1$ and $w(2) \geq 4w(1) \geq 0$. Hence, $\mathcal{W}_\Lambda^{\mathbf{1}}$ is a line with endpoints $u = (0, 1)$ and $v = (\frac{1}{5}, \frac{4}{5})$.

Consider set of alternatives $A = \{\alpha, \beta\}$, where $\beta = (44, 1)$. α is optimal in scenario u (since $\alpha \cdot u = 2 \geq 1 = \beta \cdot u$), and β is optimal in scenario v. The max regret of α is equal to the regret of α in scenario v (i.e., the difference between the utility of the best alternative and α) which equals $(\beta - \alpha) \cdot v = (24, -1) \cdot (\frac{1}{5}, \frac{4}{5}) = \frac{20}{5} = 4$. Similarly, the max regret of β equals $(\alpha - \beta) \cdot u = 1$. Thus, α has 4 times the max regret as β, so one might be inclined to say that β is much more promising, and to recommend β.

However, we might have used different units for the monetary gain, e.g., using cents rather than euros. The new representation α' of α is equal to $(20, 200)$, and $\beta' = (44, 100)$ and $\delta' = (24, 100)$, and thus $\lambda' = (-4, 100)$. The extreme points of $\mathcal{W}_{\Lambda'}^{\mathbf{1}}$ are $u' = (0, 1)$ and $v' = (\frac{25}{26}, \frac{1}{26})$ leading to a max regret for α' equalling $(\beta' - \alpha') \cdot v' = (24, -100) \cdot (\frac{25}{26}, \frac{1}{26})$, i.e., around 19, and a max regret of β' of 100. But α' and α represent exactly the same alternative, and similarly, β and β'. So, now the alternative represented by α looks much more promising than β.

The example shows that the scaling of the objectives can make a major difference to the max regret values, and to which alternative minimises max regret. Some kind of normalisation of the objective scales is typically required; for instance, if the variation of values of one objective are orders of magnitude greater than that of another objective, then the first objective will tend to be much more significant for the value of max regret.

However, there are many different reasonable ways one could choose to normalise. We can take a dataset including a set of multi-objective alternatives, and we can scale it to make the difference between the max and min value for each criterion equal to 1; or we can rescale so that the standard deviation of the values of each criterion is equal to 1. Outliers can make a big difference, so there is a choice of excluding certain alternatives in this normalisation computation; but there can be subjective elements in choosing which outliers should be excluded. Even more significantly, there can be many ways of choosing the dataset for rescaling. It can be a dataset for the current problem; or it can be a historical dataset, where one needs to choose how far to go back. A further issue is that adding a very poor alternative (which might seem as if it should be irrelevant) can significantly change the scaling, since it affects the normalisation. So one might perhaps focus only on the better alternatives.

It is thus clear that there are many reasonable ways of normalising a scale, and many subjective choices; we will see that these choices can significantly affect the max regret score. This leads to a form of imprecision in the max regret. In this paper we derive results that allow one to reason about this imprecision, hence leading to a more robust approach.

3 Preference Models and Regret

We consider a system assisting in solving a multi-objective maximisation problem, where p is the number of objectives, and in which (as is common) the system is not aware of the precise user preference model, and considers a set of plausible models. Each alternative is associated with a (multi-objective utility) vector in $I\!R^p$. We consider a finite set $A \subseteq I\!R^p$ of alternatives, and we would ideally like to identify an alternative that the user regards as optimal. We abbreviate the set $\{1, \ldots, p\}$ of indices of the objectives to $[p]$.

The preference information available to the system is represented in terms of a set of user preference models, parameterised by a set (of *scenarios*) $\mathcal{W}$ where, associated with each scenario $w \in \mathcal{W}$, is a (real-valued) utility function f_w over alternatives. We consider here a simple linear model of the user preferences: for all $\alpha \in I\!R^p$, $f_w(\alpha) = w \cdot \alpha = \sum_{i=1}^p w(i)\alpha(i)$. *Weights vector* w is formally an element of $I\!R^p_{\geq}$, where $I\!R^p_{\geq}$ is the set of vectors in $I\!R^p$, such that each co-ordinate is non-negative ($w(i) \geq 0$ for all $i \in [p]$). Each element w of $\mathcal{W}$ is viewed as a possible model of the user's preferences that is consistent with the preference information we know. If we knew that w were the true scenario, so that f_w represents the user's preferences over alternatives, then we would be able to choose a best element of A with respect to f_w leading to a utility value $Val_A(w) =_{\text{def}} \max_{\alpha \in A} \alpha \cdot w$.

Necessarily Optimal Set $\mathrm{NO}_{\mathcal{W}}(A)$**:** for each $w \in I\!R^p_{\geq}$ and finite $A \subseteq I\!R^p$ we define $\mathrm{O}_w(A)$ to be all elements α of A that are optimal in A in scenario w, i.e., such that $\alpha \cdot w = Val_A(w)$, i.e., such that for all $\beta \in A$, $\alpha \cdot w \geq \beta \cdot w$. For $\mathcal{W} \subseteq I\!R^p_{\geq}$ we define the set of *necessarily optimal* elements $\mathrm{NO}_{\mathcal{W}}(A)$ of A to be $\bigcap_{w \in \mathcal{W}} \mathrm{O}_w(A)$, the set of alternatives that are optimal in every scenario. If there is a necessarily optimal element then we can recommend it to the user, since it is optimal with respect to any of the compatible user preference models. However, very frequently there won't be one. We can elicit more information from the user, to restrict the compatible preference models; or we could recommend an alternative $\alpha \in A$ minimising max regret.

Regret: For finite set A of alternatives, the regret $D_w^A(\alpha)$ for alternative $\alpha(\in A)$ in scenario w is defined by $D_w^A(\alpha) = Val_A(w) - w \cdot \alpha$; this thus measures how far α is from being optimal in A with respect to w. Hence, α is optimal ($\alpha \in \mathrm{O}_w(A)$) $\iff$ $D_w^A(\alpha) = 0$.

Max and minimax regret: Let $\mathcal{W}$ be a compact (i.e., bounded and topologically closed) subset of $I\!R^p_{\geq}$, representing a set of scenarios, i.e., weights vectors. The max regret $MR(\alpha, A; \mathcal{W})$ for alternative $\alpha \in A$ over set of scenarios $\mathcal{W}$ is given by $MR(\alpha, A; \mathcal{W}) = \max_{\beta \in A} \max_{w \in \mathcal{W}} w \cdot (\beta - \alpha) = \max_{w \in \mathcal{W}} (Val_A(w) - w \cdot \alpha)$, which equals $\max_{w \in \mathcal{W}} D_w^A(\alpha)$. The minimax regret is then defined as follows: $MMR(A; \mathcal{W}) = \min_{\alpha \in A} MR(\alpha, A; \mathcal{W})$. Max regret is always non-negative, and the max regret of α is zero if and only α is necessarily optimal in A; max regret can be viewed as a measure of how far α is from being necessarily optimal.

Let $Ext(\mathcal{W})$ be the set of extreme points of $\mathcal{W}$. Using the fact that for compact $\mathcal{W}$, $\max_{w \in \mathcal{W}} w \cdot (\beta - \alpha)$ equals $\max_{w \in CH(\mathcal{W})} w \cdot (\beta - \alpha)$, where $CH(\mathcal{W})$ is the convex hull of $\mathcal{W}$, and that $CH(Ext(\mathcal{W}))$ contains $\mathcal{W}$ and so equals $CH(\mathcal{W})$, we see that $MR(\alpha, A; \mathcal{W}) = MR(\alpha, A; Ext(\mathcal{W}))$. Computing $\{MR(\alpha, A; \mathcal{W})$ for all $\alpha \in A\}$ is $O(|Ext(\mathcal{W})||A|p)$ using $MR(\alpha, A; \mathcal{W}) = \max_{w \in Ext(\mathcal{W})} (Val_A(w) - w \cdot \alpha)$. This is then a feasible method if the number of alternatives and the number of extreme points are not too large.

In this paper we focus on a situation in which the number of extreme points is not huge. This will tend to be the case if the number of input preferences Λ (see below) is not too large; or if the dimension p is small.

Generation of set of user models $\mathcal{W}_\Lambda$: Let us consider, for finite set Λ of utility vectors (each element is in $I\!R^p$), the set $\mathcal{W}_\Lambda$ defined to be the set of weights vectors w such that $w \cdot \lambda \geq 0$ for all $\lambda \in \Lambda$. As in the earlier example, each λ may come from a preference for an alternative α_λ over an alternative β_λ, implying $w \cdot (\alpha_\lambda - \beta_\lambda) \geq 0$, where $\alpha_\lambda - \beta_\lambda = \lambda$.

Consistency assumption: Throughout the paper we will assume that Λ is consistent, i.e., that $\mathcal{W}_\Lambda$ is non-empty.

3.1 Normalisations of weights vectors

Max regret cannot be applied to $\mathcal{W}_\Lambda$, because it's an unbounded set (or else an infinite value would be obtained). It thus seems natural to consider normalised weights vectors, as is commonly done. Let $\mathcal{W}^{\mathbf{1}}_\Lambda$ equal $\{w \in \mathcal{W}_\Lambda : w \cdot \mathbf{1} = 1\}$, which also equals $\{\frac{w}{w \cdot \mathbf{1}} : w \in \mathcal{W}_\Lambda\}$, where $\mathbf{1} \in I\!R^p$ is the vector of ones $(1, \ldots, 1)$. $\mathcal{W}^{\mathbf{1}}_\Lambda$ consists of the elements w of $\mathcal{W}_\Lambda$ that are normalised in the sense that $\sum_{i=1}^p w(i) = 1$. This makes the utility $w \cdot \alpha$ of alternative α a weighted average of its components $\alpha(i)$.

Let σ be an element of $I\!R^p_+ = \{\tau \in I\!R^p : \forall i \in [p], \tau(i) > 0\}$, the set of vectors with strictly positive values. (We will use vectors such as σ to represent rescalings of the objectives.) We define $\mathcal{W}^\sigma_\Lambda$ to be $\{w \in \mathcal{W}_\Lambda : w \cdot \sigma = 1\}$, which also equals $\{\frac{w}{w \cdot \sigma} : w \in \mathcal{W}_\Lambda\}$ and $\{\frac{w}{w \cdot \sigma} : w \in \mathcal{W}^{\mathbf{1}}_\Lambda\}$. $\mathcal{W}^\sigma_\Lambda$ is the set of elements of $\mathcal{W}_\Lambda$ that are normalised with a particular weighted sum.

3.2 Max regret over a differently normalised set of scenarios

Before explicitly focusing on rescaling the objectives we give some mathematical results relating to max regret over $\mathcal{W}^\sigma_\Lambda$, for strictly positive $\sigma \in I\!R^p_+$. In the next section we will show that this is equal to max regret over the rescaled objective space.

For arbitrary non-negative vector $w \in I\!R^p_\geq$ with $w \neq \mathbf{0} = (0, \ldots, 0)$ define $h_\sigma(w) = \frac{w}{w \cdot \sigma}$. The function h_σ changes the magnitude but not the direction of each vector w. It can be shown[1] that h_σ is a bijection from $\mathcal{W}^{\mathbf{1}}_\Lambda$ onto $\mathcal{W}^\sigma_\Lambda$ whose restriction to $Ext(\mathcal{W}^{\mathbf{1}}_\Lambda)$ is a bijection onto $Ext(\mathcal{W}^\sigma_\Lambda)$, leading to Proposition 1.

The following result gives a method for computing the max regret over the set $\mathcal{W}^\sigma_\Lambda$ of weights vectors, using the extreme points $Ext(\mathcal{W}^{\mathbf{1}}_\Lambda)$ of the original set of weights vectors.

Proposition 1 *Let $E = Ext(\mathcal{W}^{\mathbf{1}}_\Lambda)$ be the extreme points of $\mathcal{W}^{\mathbf{1}}_\Lambda$. Then $MR(\alpha, A; \mathcal{W}^\sigma_\Lambda) = \max_{w \in E} \frac{D^A_w(\alpha)}{w \cdot \sigma}$.*

Proof: $\frac{D^A_w(\alpha)}{w \cdot \sigma} = D^A_{h_\sigma(w)}(\alpha)$, and $h_\sigma(E) = Ext(\mathcal{W}^\sigma_\Lambda)$. Then we have $\max_{w \in E} \frac{D^A_w(\alpha)}{w \cdot \sigma} = \max_{w \in E} D^A_{h_\sigma(w)}(\alpha)$, which is equal to $\max_{v \in h_\sigma(E)} D^A_v(\alpha) = \max_{v \in Ext(\mathcal{W}^\sigma_\Lambda)} D^A_v(\alpha) = MR(\alpha, A, Ext(\mathcal{W}^\sigma_\Lambda))$, which equals $MR(\alpha, A; \mathcal{W}^\sigma_\Lambda)$. □

[1] See the longer version of the paper available online, which also contains the proofs of results not included here [24].

Example 1 *Continuing the example in Section 2, the set E of extreme points of $\mathcal{W}^{\mathbf{1}}_\Lambda$ equals $\{u, v\}$, where $u = (0, 1)$ and $v = (\frac{1}{5}, \frac{4}{5})$. We have $D^A_w(\alpha) = 0$ and $D^A_v(\alpha) = v \cdot (\beta - \alpha) = (\frac{1}{5}, \frac{4}{5}) \cdot (24, -1) = 4$; and $D^A_v(\beta) = 0$ and $D^A_u(\beta) = u \cdot (\alpha - \beta) = 1$. With $\sigma = (1, \epsilon)$ (where $0 < \epsilon \leq 1$) we obtain $u \cdot \sigma = \epsilon$ and $v \cdot \sigma = \frac{1+4\epsilon}{5}$. Thus, $\max_{w \in E} \frac{D^A_w(\alpha)}{w \cdot \sigma} = \frac{D^A_v(\alpha)}{v \cdot \sigma} = \frac{20}{1+4\epsilon}$; and $\max_{w \in E} \frac{D^A_w(\beta)}{w \cdot \sigma} = \frac{D^A_u(\beta)}{u \cdot \sigma} = \frac{1}{\epsilon}$.*

The set $\mathcal{W}^\sigma_\Lambda$ has extreme points $(0, \frac{1}{\epsilon})$ and $\frac{1}{1+4\epsilon}(1, 4)$. We have $MR(\alpha, A; \mathcal{W}^\sigma_\Lambda) = (\beta - \alpha) \cdot \frac{1}{1+4\epsilon}(1, 4) = (24, -1) \cdot \frac{1}{1+4\epsilon}(1, 4) = \frac{20}{1+4\epsilon}$; and $MR(\beta, A; \mathcal{W}^\sigma_\Lambda) = (\alpha - \beta) \cdot (0, \frac{1}{\epsilon}) = \frac{1}{\epsilon}$.

4 Rescaling and its Effect on Max Regret

In this section we consider the effect on max regret of changing the scales of the objectives, which involves multiplying or dividing the objective values by real numbers. Mathematically this can be expressed in terms of the pointwise product of vectors. For $\alpha, \beta \in I\!R^p$, their pointwise product $\alpha \odot \beta$ ($\in I\!R^p$) is defined by $(\alpha \odot \beta)(i) = \alpha(i)\beta(i)$ for $i \in [p]$. For $\sigma \in I\!R^p_+$, we define σ^{-1} ($\in I\!R^p$) by $\sigma^{-1}(i) = 1/\sigma(i)$ for each $i \in [p]$. Thus, $\sigma \odot \sigma^{-1} = \mathbf{1}$. Pointwise product is associative and commutative, and commutes (in a certain sense) with dot product: for $\alpha, \beta, \gamma \in I\!R^p$, $(\alpha \odot \beta) \cdot \gamma = \alpha \cdot (\beta \odot \gamma) = \sum_{i=1}^p \alpha(i)\beta(i)\gamma(i)$. In the obvious way we extend to sets: if $\alpha \in I\!R^p$ and $X \subseteq I\!R^p$ we define $\alpha \odot X = X \odot \alpha = \{\alpha \odot \beta : \beta \in X\}$.

A *rescaling vector* σ has every co-ordinate strictly greater than zero, so is an element of $I\!R^p_+ = \{\tau \in I\!R^p : \forall i \in [p], \tau(i) > 0\}$. It is associated with a transformation H_σ on multi-objective utility vectors defined by: $H_\sigma(\alpha) = \alpha \odot \sigma^{-1} = (\frac{\alpha(1)}{\sigma(1)}, \ldots, \frac{\alpha(p)}{\sigma(p)})$. In the example in Section 2, changing from euros to cents in the second objective used the transformation H_σ with $\sigma = (1, 0.01)$ (because 1 cent = €0.01). For $\alpha, \beta \in I\!R^p$, we have that $H_\sigma(\alpha - \beta) = (\alpha - \beta) \odot \sigma^{-1} = (\alpha \odot \sigma^{-1}) - (\beta \odot \sigma^{-1}) = H_\sigma(\alpha) - H_\sigma(\beta)$. It is natural then to define H_σ also on the input preference vectors: $H_\sigma(\Lambda) = \Lambda \odot \sigma^{-1} = \{\lambda \odot \sigma^{-1} : \lambda \in \Lambda\}$.

The following result relates the rescaled version of $\mathcal{W}_\Lambda$ and its normalised form to the original versions.

Proposition 2 *For finite $\Lambda \subseteq I\!R^p$ and $\sigma \in I\!R^p_+$, we have $\mathcal{W}_{H_\sigma(\Lambda)} = \mathcal{W}_{\Lambda \odot \sigma^{-1}} = (\mathcal{W}_\Lambda) \odot \sigma$, and $\mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)} = \mathcal{W}^{\mathbf{1}}_{\Lambda \odot \sigma^{-1}} = (\mathcal{W}^\sigma_\Lambda) \odot \sigma$, and $\mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)} \odot \sigma^{-1} = \mathcal{W}^\sigma_\Lambda$.*

Proof: $w \in \mathcal{W}_{\Lambda \odot \sigma^{-1}}$ if and only if for all $\lambda \in \Lambda$ we have $w \cdot (\lambda \odot \sigma^{-1}) \geq 0$, i.e., for all $\lambda \in \Lambda$ we have $(w \odot \sigma^{-1}) \cdot \lambda \geq 0$, which is if and only if $w \odot \sigma^{-1} \in \mathcal{W}_\Lambda$, which is if and only if $w \in (\mathcal{W}_\Lambda) \odot \sigma$.

The normalised set $\mathcal{W}^{\mathbf{1}}_{\Lambda \odot \sigma^{-1}}$ is, by the first part, equal to the set of all elements $w \odot \sigma$ of $(\mathcal{W}_\Lambda) \odot \sigma$ such that $(w \odot \sigma) \cdot \mathbf{1} = 1$, i.e., such that $w \cdot \sigma = 1$. Hence, $\mathcal{W}^{\mathbf{1}}_{\Lambda \odot \sigma^{-1}} = (\mathcal{W}^\sigma_\Lambda) \odot \sigma$. Applying $\odot \sigma^{-1}$ to both sides of the equality gives $\mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)} \odot \sigma^{-1} = \mathcal{W}^\sigma_\Lambda$. □

The user input sets Λ and $H_\sigma(\Lambda)$ ($= \Lambda \odot \sigma^{-1}$) represent exactly the same information: they are just expressed differently, i.e., with different units. Similarly, $\mathcal{W}_\Lambda$ and $\mathcal{W}_{H_\sigma(\Lambda)}$ ($= (\mathcal{W}_\Lambda) \odot \sigma$) represent essentially the same set of user models. However, $\mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)}$ ($= (\mathcal{W}^\sigma_\Lambda) \odot \sigma$) represents an essentially different set of user models from $\mathcal{W}^{\mathbf{1}}_\Lambda$; instead it represents the same set of user models as $\mathcal{W}^\sigma_\Lambda$ (with the latter using the original units). $\mathcal{W}^{\mathbf{1}}_\Lambda$ and $\mathcal{W}^\sigma_\Lambda$ both correspond to $\mathcal{W}_\Lambda$, but they involve a different way of normalising that affects models differently. For example, if $\sigma(1) = 1000$ and otherwise, $\sigma(i) = 1$, then any $w \in \mathcal{W}^{\mathbf{1}}_\Lambda$ with a relatively large value of

$w(1)$, will be scaled right down in $\mathcal{W}_\Lambda^\sigma$; e.g., if $w(1) = 0.5$ then w will become a vector $w' = rw \in \mathcal{W}_\Lambda^\sigma$ with $r = \frac{1}{w \cdot \sigma} \approx 0.002$. In contrast, a vector $v \in \mathcal{W}_\Lambda^{\mathbf{1}}$ with $v(1) = 0$ will be left unchanged. With this rescaling, the new version w' of w is likely to make much less influence on max regret values than it did without the change of scale.

In general, increasing a scaling factor $\sigma(i)$ will tend to reduce the influence of the ith objective on max regret values, especially when not too many user preferences Λ have been elicited.

The next general property of rescaling max regret follows easily from the definitions. It implies that the effect of rescaling on max regret can be viewed in terms of a transformation on the set of scenarios (as shown more explicitly by Theorem 1 below).

Proposition 3 *For finite $A \subseteq \mathbb{R}^p$, $\alpha \in A$, and $\sigma \in \mathbb{R}^p_+$, and for any compact subset $\mathcal{W}$ of $\mathbb{R}^p_\geq$, we have $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}) = MR(\alpha, A; \mathcal{W} \odot \sigma^{-1})$.*

Proof: $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W})$ is equal to $MR(\alpha \odot \sigma^{-1}, A \odot \sigma^{-1}); \mathcal{W}) = \max_{\beta \in A} \max_{w \in \mathcal{W}} w \cdot (\beta \odot \sigma^{-1} - \alpha \odot \sigma^{-1}) = \max_{\beta \in A} \max_{w \in \mathcal{W}} w \odot \sigma^{-1} \cdot (\beta - \alpha) = \max_{\beta \in A} \max_{v \in \mathcal{W} \odot \sigma^{-1}} v \cdot (\beta - \alpha) = MR(\alpha, A; \mathcal{W} \odot \sigma^{-1})$. □

Consider preference inputs Λ and alternative α in A, which has associated max regret $MR(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ (assuming normalised scenarios). If we rescale with σ then α, A and Λ are all transformed using H_σ, and the associated max regret is $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)})$. Using Proposition 3 and Proposition 2, this equals $MR(\alpha, A; \mathcal{W}_\Lambda^\sigma)$, and Proposition 1 then implies the theorem below, summarising the results so far on the effect of rescaling on max regret.

Theorem 1 *Let A and Λ be finite subsets of $\mathbb{R}^p$ and let $\alpha \in A$, and let $\sigma \in \mathbb{R}^p_+$. Then the max regret $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)})$ of the rescaled space equals $MR(\alpha, A; \mathcal{W}_\Lambda^\sigma) = \max_{w \in E} \frac{D_w^A(\alpha)}{w \cdot \sigma}$, where $E = Ext(\mathcal{W}_\Lambda^{\mathbf{1}})$ is the set of extreme points of $\mathcal{W}_\Lambda^{\mathbf{1}}$.*

Example 2 *Continuing Example 1, $H_\sigma(\lambda) = (-4, 1) \odot \sigma^{-1} = (-4, \frac{1}{\epsilon})$, and so $\mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)}$ consists of all $w \in \mathbb{R}^p_\geq$ such that $w \cdot \mathbf{1} = 1$ and $w \cdot (-4, \frac{1}{\epsilon}) \geq 0$. The extreme points of $\mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)}$ are $(0, 1)$ and $\frac{1}{1+4\epsilon}(1, 4\epsilon)$. It can be seen that $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)})$ equals $H_\sigma(\beta - \alpha) \cdot \frac{1}{1+4\epsilon}(1, 4\epsilon) = (24, \frac{-1}{\epsilon}) \cdot \frac{1}{1+4\epsilon}(1, 4\epsilon) = \frac{20}{1+4\epsilon}$. Also, $MR(H_\sigma(\beta), H_\sigma(A); \mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)})$ equals $H_\sigma(\alpha - \beta) \cdot (0, 1) = (-24, \frac{1}{\epsilon}) \cdot (0, 1) = \frac{1}{\epsilon}$. This tallies with Example 1, thus illustrating Theorem 1.*

Notation: To emphasise the dependence on the rescaling vector σ, we now use the notation $MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ to mean the max regret $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}^{\mathbf{1}}_{H_\sigma(\Lambda)})$ of the rescaled space, which, by Theorem 1, also equals $MR(\alpha, A; \mathcal{W}_\Lambda^\sigma)$.

The result below expresses the rescaled max regret in a convenient form, when we are interested in a range of different rescaling vectors.

Proposition 4 *Let A and Λ be finite subsets of $\mathbb{R}^p$ and let $\sigma \in \mathbb{R}^p_+$. Let $E = Ext(\mathcal{W}_\Lambda^{\mathbf{1}})$ and for each $\alpha \in A$ define E_α to be $\{w \in E : \mathrm{O}_w(A) \not\ni \alpha\}$, the set all elements w of E in which α is not optimal with respect to w, and define for each $w \in E_\alpha$, w_A^α (abbreviated to w^α) to be $\frac{w}{D_w^A(\alpha)}$. Define $T(\alpha)$ to be $\{w^\alpha : w \in E_\alpha\}$. Then, $MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}}) = 0$ if and only if $\alpha \in \mathrm{NO}_{\mathcal{W}_\Lambda^{\mathbf{1}}}(A)$ if and only if $T(\alpha) = \emptyset$. Otherwise, if $\alpha \notin \mathrm{NO}_{\mathcal{W}_\Lambda^{\mathbf{1}}}(A)$, we have $\frac{1}{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})} = \min_{w \in E_\alpha} w^\alpha \cdot \sigma = \min_{v \in T(\alpha)} v \cdot \sigma$.*

For $\alpha, \gamma \in A \setminus \mathrm{NO}_{\mathcal{W}_\Lambda^{\mathbf{1}}}(A)$, we then have

$$\frac{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR_\sigma(\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})} = \frac{\min_{v \in E_\gamma} v^\gamma \cdot \sigma}{\min_{w \in E_\alpha} w^\alpha \cdot \sigma} = \frac{\min_{v \in T(\gamma)} v \cdot \sigma}{\min_{w \in T(\alpha)} w \cdot \sigma}.$$

Example 3 *With the running example we have $E = Ext(\mathcal{W}_\Lambda^{\mathbf{1}}) = \{u, v\}$, where $u = (0, 1)$ and $v = (\frac{1}{5}, \frac{4}{5})$. Then, $E_\alpha = \{v\}$ and $E_\beta = \{u\}$ since $\mathrm{O}_v(\{\alpha, \beta\}) = \{\beta\}$ and $\mathrm{O}_u(\{\alpha, \beta\}) = \{\alpha\}$. We have $D_v^A(\alpha) = v \cdot (\beta - \alpha) = (\frac{1}{5}, \frac{4}{5}) \cdot (24, -1) = 4$, and thus, $v^\alpha = (\frac{1}{20}, \frac{1}{5})$. We have $D_u^A(\beta) = u \cdot (\alpha - \beta) = 1$, and so $u^\beta = (0, 1)$. Then, with $\sigma = (1, \epsilon)$ (where $0 < \epsilon \leq 1$) we have $\frac{1}{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})} = v^\alpha \cdot \sigma = \frac{1}{20} + \frac{\epsilon}{5} = \frac{1+4\epsilon}{20}$. Also, $\frac{1}{MR_\sigma(\beta, A; \mathcal{W}_\Lambda^{\mathbf{1}})} = u^\beta \cdot \sigma = \epsilon$.*

We abbreviate $\frac{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR_\sigma(\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})}$ to $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$, and similarly, $MR(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ is an abbreviation for $\frac{MR(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR(\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})}$, which is equal to $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ when $\sigma = (1, \ldots, 1)$.

The next results give lower and upper bounds on the ratio of $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ and $MR(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$; these show limits to how much the max regret $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ is affected by the rescaling vector σ.

Proposition 5 *Let A and Λ be finite subsets of $\mathbb{R}^p$ and let $\alpha, \gamma \in A$, and let $\sigma \in \mathbb{R}^p_+$. Define $\underline{U}$ and $\overline{U}$ by for each $i \in [p]$, $\underline{U}(i) = \min_{w \in \mathcal{W}_\Lambda^{\mathbf{1}}} w(i)$ (which equals $\min_{w \in Ext(\mathcal{W}_\Lambda^{\mathbf{1}})} w(i)$) and $\overline{U}(i) = \max_{w \in \mathcal{W}_\Lambda^{\mathbf{1}}} w(i)$. Then for all $\sigma \in \mathbb{R}^p_+$,*

$$\frac{MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})} \in \left[\min_{i \in [p]} \frac{\underline{U}(i)}{\overline{U}(i)}, \max_{i \in [p]} \frac{\overline{U}(i)}{\underline{U}(i)}\right].$$

For $\delta > 0$ and $\mathcal{W} \subseteq \mathbb{R}^p$ let us say that $\mathcal{W}$ is *δ-small* if there exists $u \in \mathbb{R}^p_\geq$ such that for all $w \in \mathcal{W}$ and for all $i \in [p]$, $u(i) \leq w(i) \leq (1 + \delta)u(i)$.

The following result shows that once Λ is sufficiently large to make $\mathcal{W}_\Lambda^{\mathbf{1}}$ δ-small for small δ then the ratio of max regrets $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ is not much affected by the choice of rescaling σ because it is always in the relatively narrow interval $[\frac{x}{1+\delta}, (1 + \delta)x]$, where $x = MR(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$.

Corollary 1 *Let Λ be a finite subset of $\mathbb{R}^p$ such that $\mathcal{W}_\Lambda^{\mathbf{1}}$ is δ-small. Then for any finite $A \subseteq \mathbb{R}^p$ and for all $\alpha, \gamma \in A$, and all $\sigma \in \mathbb{R}^p_+$,*

$$\frac{MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})} \in \left[\frac{1}{1 + \delta}, 1 + \delta\right].$$

5 Extreme Rescaling

In the example in Section 2 we used rescaling vector $\sigma = (1, \epsilon)$ with $\epsilon = 0.01$. If we instead use a more extreme rescaling with $\epsilon = 0.0001$ (corresponding to a unit of a hundredth of a cent) then it follows from Example 3 that we obtain a max regret for the alternatives (corresponding with) β and α to be 10000 and around 20, respectively. In fact, with varying ϵ, we can make the ratio $MR_\sigma(\beta/\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ $(= \frac{1+4\epsilon}{20\epsilon})$ of these max regrets unbounded, i.e., arbitrarily large. Define the *relative max regret* $RMR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ to be $\frac{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MMR_\sigma(A; \mathcal{W}_\Lambda^{\mathbf{1}})}$, which equals $\max_{\gamma \in A} MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$, where $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ equals $\frac{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR_\sigma(\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})}$. In this section we explore when $RMR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ and $MR_\sigma(\alpha/\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})$ can be made arbitrarily large, giving simple characterisations. In this way we can reveal situations in which max regret is especially vulnerable to changes of scale.

First we formalise extreme relative max regret. Let $g(\sigma)$ be a non-negative real-valued function of scale vector σ, where σ varies over a subset R of $I\!R^p_+$. We say that $g(\sigma)$ is *bounded over* $\sigma \in \mathrm{R}$ if there exists a number K such that for all scale vectors $\sigma \in \mathrm{R}$, $g(\sigma) \leq K$. Otherwise, we say that $g(\sigma)$ is *unbounded over* $\sigma \in \mathrm{R}$; this is the case if and only if there exists a sequence $\{\sigma_k : k = 1, 2, \ldots\}$ of elements in R such that $\lim_{k\to\infty} g(\sigma_k) = \infty$.

In the running example, the functions $MR_\sigma(\beta/\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ and $RMR_\sigma(\beta, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ are unbounded, and $MR_\sigma(\alpha/\beta, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ and $RMR_\sigma(\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ are bounded, over $\sigma \in I\!R^p_+$.

We make use of the following notation.

Definition 1 (ze(w) **and** $E(Y)$) *For non-negative vector $w \in I\!R^p_{\geq}$, we define $\mathrm{ze}(w)$ to be $\{i \in [p] : w(i) = 0\}$, the set of components on which w is zero. For $Y \subseteq [p]$ we define $E(Y)$ to consist of all extreme points $w \in E$ (of $\mathcal{W}^{\mathbf{1}}_\Lambda$) such that $\mathrm{ze}(w) \supseteq Y$, i.e., such that w is zero on every element of Y.*

Typically, for many sets Y, the set $E(Y)$ is empty. For non-empty $E(Y)$ we will be interested in the set $\mathrm{NO}_{E(Y)}(A)$: these are the alternatives that are optimal in A with respect to every extreme point w that is zero on every element of Y.

If $Y' \subseteq Y$ then $E(Y') \supseteq E(Y)$ and $\mathrm{NO}_{E(Y')}(A) \subseteq \mathrm{NO}_{E(Y)}(A)$.

The following simple result connects $\mathrm{NO}_{E(Y)}(A)$ and $\mathrm{ze}(w)$, making use of E_γ as defined in Proposition 4.

Proposition 6 *For $Y \subseteq [p]$ such that $E(Y) \neq \emptyset$ and $\gamma \in A$ we have $\mathrm{NO}_{E(Y)}(A) \ni \gamma$ if and only if $\gamma \in \mathrm{O}_v(A)$ for all $v \in E$ such that $\mathrm{ze}(v) \supseteq Y$, which is if and only if for all $v \in E_\gamma$, $\mathrm{ze}(v) \not\supseteq Y$.*

The next result gives conditions under which a ratio of the form $\frac{v\cdot\sigma}{w\cdot\sigma}$ is unbounded.

Proposition 7 *Let $v, w \in I\!R^p_{\geq}$ and suppose that Y ($\subseteq [p]$) is such that $\mathrm{ze}(w) \supseteq Y$ and $\mathrm{ze}(v) \not\supseteq Y$. Define $\sigma^Y_k \in I\!R^p_+$ for each $k = 1, 2, \ldots$, by $\sigma^Y_k(i) = 1$ for $i \in Y$ and $\sigma^Y_k(i) = 1/k$ for $i \notin Y$. Then $\lim_{k\to\infty} \frac{v\cdot\sigma^Y_k}{w\cdot\sigma^Y_k} = \infty$.*

Proof: There exists some j in $Y \setminus \mathrm{ze}(v)$, and we have, for every k, $v \cdot \sigma^Y_k \geq v(j) > 0$. Also, because $\mathrm{ze}(w) \supseteq Y$ we have $w \cdot \sigma^Y_k = \sum_{i\in[p]\setminus \mathrm{ze}(w)} \frac{w(i)}{k} = \frac{1}{k}\sum_{i\in[p]} w(i)$. Thus, $\frac{v\cdot\sigma^Y_k}{w\cdot\sigma^Y_k} \geq k\frac{v(j)}{\sum_{i\in[p]} w(i)}$, which tends to infinity as k tends to infinity. □

Suppose that $\sigma_1, \sigma_2, \ldots$ is a sequence of rescaling vectors converging to a vector $\sigma_\infty \in I\!R^p_{\geq}$, and let Y be the set of components i such that $\sigma_\infty(i) > 0$, like in the sequence defined in Proposition 7. Then $w \cdot \sigma_k$ tends to zero if and only if $\mathrm{ze}(w) \supseteq Y$; and, if $\mathrm{ze}(w) \supseteq Y$ and $\mathrm{ze}(v) \not\supseteq Y$, then $\frac{v\cdot\sigma_k}{w\cdot\sigma_k}$ tends to infinity. Proposition 4 then implies that $MR_{\sigma_k}(\alpha/\gamma, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ tends to infinity if and only if for all $v \in E^\gamma$, $\mathrm{ze}(v) \not\supseteq Y$ and there exists $w \in E^\alpha$ such that $\mathrm{ze}(w) \supseteq Y$. Proposition 6 implies that this is if and only if $\alpha \notin \mathrm{NO}_{E(Y)}(A) \ni \gamma$. This is the basic idea behind Theorem 2, which gives characterisations of when $MR_\sigma(\alpha/\gamma, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded over $\sigma \in I\!R^p_+$. In particular the characterisation (b) is that there exists an extreme point w (of the set of normalised weights vectors) for which α is non-optimal, and γ is optimal with respect to every extreme point v that is zero whenever w is zero.

Theorem 2 *Let A and Λ be finite subsets of $I\!R^p$ and let $\alpha, \gamma \in A$, and let $E = Ext(\mathcal{W}^{\mathbf{1}}_\Lambda)$. The following four conditions are equivalent.*

(a) $\frac{MR_\sigma(\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)}{MR_\sigma(\gamma, A; \mathcal{W}^{\mathbf{1}}_\Lambda)}$ *is unbounded over $\sigma \in I\!R^p_+$.*
(b) There exists $w \in E$ such that $\alpha \notin \mathrm{O}_w(A)$ and $\gamma \in \mathrm{O}_v(A)$ for all $v \in E$ such that $\mathrm{ze}(v) \supseteq \mathrm{ze}(w)$.
(c) There exists $w \in E$ such that $\alpha \notin \mathrm{NO}_{E(\mathrm{ze}(w))}(A) \ni \gamma$.
(d) There exists $Y \subseteq [p]$ such that $\alpha \notin \mathrm{NO}_{E(Y)}(A) \ni \gamma$.

Example 4 *Continuing the running example, where $A = \{\alpha, \beta\}$ we have $\mathrm{O}_{(0,1)}(A) = \{\alpha\}$, i.e., α is optimal with respect to $w = (0, 1)$. We have $\mathrm{ze}((0, 1)) = \{1\}$, since the first component of $(0, 1)$ is zero. We have $E(\mathrm{ze}((0, 1))) = \{(0, 1)\}$ since $(0, 1)$ is the only extreme point that is zero on its first component. $\mathrm{NO}_{E(\mathrm{ze}((0,1)))}(A) = \mathrm{O}_{(0,1)}(A) = \{\alpha\}$.*

With this example, Condition (a) of Theorem 2 holds for β/α because, with $\sigma = (1, \epsilon)$, $MR(\beta/\alpha, A; \mathcal{W}^\sigma_\Lambda)$ equals $\frac{1+4\epsilon}{20\epsilon}$, which tends to infinity as ϵ tends to zero. Condition (b) holds for β/α, since $\alpha \in \mathrm{O}_{(0,1)}(A) \not\ni \beta$. Condition (c) holds because $\alpha \in \mathrm{NO}_{E(\mathrm{ze}((0,1)))}(A) \not\ni \beta$. For Condition (d), putting $Y = \{1\}$ we have $\alpha \in \mathrm{NO}_{E(\{1\})}(A) \not\ni \beta$.

Corollary 2 below characterises when the relative max regret $RMR_\sigma(\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded over $\sigma \in I\!R^p_+$. Part (i) easily follows from Theorem 2, and (ii) follows from (i).

Corollary 2 *Let $\alpha \in A$.*

(i) The relative max regret $RMR_\sigma(\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded over $\sigma \in I\!R^p_+$ if and only if there exists $w \in E$ such that $\mathrm{NO}_{E(\mathrm{ze}(w))}(A)$ is non-empty and does not contain α.
(ii) $RMR_\sigma(\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is bounded over $\sigma \in I\!R^p_+$ if and only if $\alpha \in \mathrm{NO}_{E(\mathrm{ze}(w))}(A)$ for all $w \in E$ such that $\mathrm{NO}_{E(\mathrm{ze}(w))}(A)$ is non-empty.

For all $w \in E$ and all $\alpha \in A$ we can compute if for all $u \in E$ such that $\mathrm{ze}(u) \supseteq \mathrm{ze}(w)$ we have $\alpha \in \mathrm{O}_u(A)$; this computes $\mathrm{NO}_{E(\mathrm{ze}(w))}(A)$ for each extreme point $w \in E$. Corollary 2 and the equivalence of (a) and (c) in Theorem 2 enable one then to efficiently determine for which alternatives α the relative max regret $RMR_\sigma(\alpha, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded, and when $MR_\sigma(\alpha/\gamma, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded.

Example 5 *Suppose that there exists an extreme point w_{12} (of $\mathcal{W}^{\mathbf{1}}_\Lambda$) with $\mathrm{ze}(w_{12}) = \{1, 2\}$, and extreme points w_1 and w_2 such that $\mathrm{ze}(w_1) = \{1\}$ and $\mathrm{ze}(w_2) = \{2\}$, and for all other extreme points u we have $\mathrm{ze}(u) = \emptyset$. Suppose that $\mathrm{O}_{w_{12}}(A) = \{\alpha\}$ and $\mathrm{O}_{w_1}(A) = \mathrm{O}_{w_2}(A) = \{\beta\}$. Then $\mathrm{NO}_{E(\mathrm{ze}(w))}(A)$ is empty unless $w = w_{12}$, and since $\mathrm{NO}_{E(\mathrm{ze}(w_{12}))}(A) = \mathrm{O}_{w_{12}}(A) = \{\alpha\}$, the theorem and corollary imply that $MR_\sigma(x/y, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded if and only if [$y = \alpha$ and $x \in A \setminus \{\alpha\}$], and $RMR_\sigma(x, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ is unbounded unless $x = \alpha$.*

Now assume that there exists, as well as w_{12}, w_1 and w_2, an extreme point v_{12} with $\mathrm{ze}(v_{12}) = \{1, 2\}$ and $\mathrm{O}_{v_{12}} = \{\beta\}$. Then $\mathrm{NO}_{E(\mathrm{ze}(w))}(A)$ is empty for all extreme points w, and so $MR_\sigma(x/y, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ and $RMR_\sigma(x, A; \mathcal{W}^{\mathbf{1}}_\Lambda)$ are bounded for all $x, y \in A$.

6 The Maximum and Minimum of Max Regret Over a Compact Set R of Rescalings

As argued in Section 2, the choice of scale for an objective, and thus, the range of a reasonable rescaling of it, can be fairly subjective.

However, there may be more uncertainty about an appropriate scale for one objective than another, because of less data about that objective, or e.g., because in the dataset there are several clusters of the values of that objective and it may be unclear whether all (or which) of the clusters should be used to normalise the scale. In this way, a decision analyst may choose some set R of reasonable rescaling vectors σ. Section 6.1 considers simple and natural forms of R. It is desirable to be able to compute how large and how small the max regret of an alternative can be, if we vary the rescaling σ over a set R; we consider these tasks in Sections 6.2 and Section 6.3.

6.1 Particular forms of the rescaling range R

In this paper our focus is how max regret varies when one considers a range R of rescalings σ. Although one can consider fairly general forms of R, and their expression in terms of a set of linear inequalities, the most obvious representations are based on a rectangular representation, which involves lower and upper bounds for each objective. We say that R is *rectangular* if for each $i \in [p]$ there exists values $a_i, b_i \in I\!R$ such that $\sigma \in \mathrm{R}$ if and only if for all $i \in [p]$, $a_i \leq \sigma(i) \leq b_i$.

However, a rectangular representation includes elements that are scalar multiples of each other. If $\sigma' = c\sigma$ for some real value c with say, $c > 1$ then, by Proposition 4, $MR_{\sigma'}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda}) = \frac{1}{c} MR_{\sigma}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda})$ for all $\alpha \in A$: the max regret function is just scaled down. This suggests that it is perhaps not so interesting to consider both σ and σ' here. It thus seems natural to consider R in which every element is normalised, i.e., every element has the same value of $\sigma \cdot \mathbf{1} = \sum_{i=1}^{p} \sigma(i)$. Since we will very often want to include the vector $\mathbf{1}$ in R, and $\mathbf{1} \cdot \mathbf{1} = p$, the normalisation condition is that $\sigma \cdot \mathbf{1} = p$. To avoid confusion with normalisation of weights vectors, we call this *sigma-normalisation*. Let $\mathcal{N}$ be the set of all sigma-normalised rescaling vectors. The *sigma-normalisation operation* $\sigma \mapsto \hat{\sigma}$ is defined by $\hat{\sigma} = \frac{p}{\sigma \cdot \mathbf{1}} \sigma$, so that σ is multiplied by the scalar $\frac{p}{\sigma \cdot \mathbf{1}}$ to form the sigma-normalised vector $\hat{\sigma}$. We define $\hat{\mathrm{R}} = \{\hat{\sigma} : \sigma \in \mathrm{R}\}$.

We say that R' is a *sigma-normalised rectangular (SNR)* range set if there exists a rectangular set R such that $\mathrm{R}' = \hat{\mathrm{R}}$. We say that R' is *rectangular intersection sigma-normalised* (RISN) if there exists a rectangular set R such that $\mathrm{R}' = \mathrm{R} \cap \mathcal{N}$, i.e., R' is the set of normalised vectors in R.

6.2 Minimising max regret over rescalings

We consider here the computation of $\min_{\sigma \in \mathrm{R}} MR_{\sigma}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda})$, which we abbreviate to MR_{min}. Using Proposition 4 and its notation we have that $MR_{min} = 0$ if and and only if $\alpha \in \mathrm{NO}_{\mathcal{W}^{\mathbf{1}}_{\Lambda}}(A)$; and if $\alpha \notin \mathrm{NO}_{\mathcal{W}^{\mathbf{1}}_{\Lambda}}(A)$ then the reciprocal of MR_{min} equals $\max_{\sigma \in \mathrm{R}} \min_{v \in T(\alpha)} v \cdot \sigma$, which is equal to the maximum value $x \in I\!R$ such that there exists some $\sigma \in I\!R^p$ with the following set of constraints being satisfiable:

$\sigma \in \mathrm{R}$ and $\{\sum_{i=1}^{p} v(i)\sigma(i) \geq x : \forall v \in T(\alpha)\}$.

Thus, if R is expressed in terms of linear constraints then we can compute MR_{min} using a single call of a linear programming solver on a problem with $p+1$ variables and $|E| + L + 1$ constraints, where L is the number of constraints defining R; e.g., $L = 2p + 1$ if R is a RISN representation.

A somewhat trivial case is if R is rectangular. Let σ_{max} be the unique Pareto-maximal point in R, taking the largest value for each objective. Then, MR_{min} equals $MR_{\sigma_{max}}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda})$.

6.3 Maximising max regret over rescalings

Given a range R of rescalings, let $\mathcal{V} = \bigcup_{\sigma \in \mathrm{R}} \mathcal{W}^{\sigma}_{\Lambda}$, which corresponds with the set of user preference models that are obtained from $\mathcal{W}^{\mathbf{1}}_{\Lambda}$ with any of the rescalings σ in the range R (when translated back into the original units). We consider max regret over $w \in \mathcal{V}$. Theorem 1 implies that this is equal to the maximum max regret over all rescalings σ in R. The following result shows that this can be expressed in a relatively simple way. For each extreme point $v \in E$ we only need to consider the maximum scalar multiple $\frac{1}{s_v}$ that is obtained over all elements of R, and which corresponds with a $\sigma \in \mathrm{R}$ that minimises $v \cdot \sigma$.

Theorem 3 *Let $E = Ext(\mathcal{W}^{\mathbf{1}}_{\Lambda})$, and for $v \in E$, let $s_v = \min_{\sigma \in \mathrm{R}} v \cdot \sigma$, and let E' be the set $\{\frac{v}{s_v} : v \in E\}$. Then $MR(\alpha, A; \mathcal{V})$ is equal to $\max_{\sigma \in \mathrm{R}} MR_{\sigma}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda}) = MR(\alpha, A; E') = \max_{w \in E'} D^A_w(\alpha) = \max_{v \in E} \frac{D^A_v(\alpha)}{s_v}$.*

Proof: By Theorem 1 and Proposition 1, we have $MR(\alpha, A; \mathcal{W}^{\sigma}_{\Lambda})$ equals $MR_{\sigma}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda})$ which is equal to $\max_{v \in E} \frac{D^A_v(\alpha)}{v \cdot \sigma}$. Now, $MR(\alpha, A; \mathcal{V})$ equals $\sup_{\sigma \in \mathrm{R}} MR(\alpha, A; \mathcal{W}^{\sigma}_{\Lambda})$, which thus equals $\sup_{\sigma \in \mathrm{R}} \max_{v \in E} \frac{D^A_v(\alpha)}{v \cdot \sigma} = \max_{v \in E} \sup_{\sigma \in \mathrm{R}} \frac{D^A_v(\alpha)}{v \cdot \sigma}$ which equals $\max_{v \in E} \frac{D^A_v(\alpha)}{s_v} = \max_{w \in E'} D^A_w(\alpha)$, which equals $MR(\alpha, A; E')$. □

If R is expressed in terms of a set of linear inequalities then for each extreme point $v \in E$ we can compute s_v using a linear minimisation. The computation can be made much faster for the cases of rectangular, SNR and RISN sets. For rectangular R we again have a rather trivial case, with $MR(\alpha, A; \mathcal{V}) = MR_{\sigma_{min}}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda})$, where σ_{min} is the unique Pareto-minimal element of R, with minimum value of each co-ordinate.

The following result expresses the property that, for a RISN set R', a $\sigma \in \mathrm{R}'$ minimising $u \cdot \sigma$ has a very simple form, enabling the minimum to be found very efficiently.

Proposition 8 *Let R' be a RISN set, so it can be written as $\mathrm{R} \cap \mathcal{N}$ for rectangular set R, where R' achieves the bounds of R. Let $[a_i, b_i]$ for $i \in [p]$ be the intervals defining R. Consider any vector u. There exists an element $\sigma \in \mathrm{R}'$ that minimises $u \cdot \sigma$ among $\sigma \in \mathrm{R}'$ with the following property: there exists some real value x such that, for all $i \in [p]$ with $u(i) < x$ we have $\sigma(i) = b_i$; and for all $i \in [p]$ with $u(i) > x$ we have $\sigma(i) = a_i$.*

For the case of a SNR set equalling $\hat{\mathrm{R}}$ for some rectangular set R, we have a similar characterisation for σ minimising $v \cdot \sigma$, which again leads to an efficient method for computing each coefficient s_v for $v \in E$. The overall computation for computing $MR(\alpha, A; \mathcal{V}) = \max_{\sigma \in \mathrm{R}} MR_{\sigma}(\alpha, A; \mathcal{W}^{\mathbf{1}}_{\Lambda})$ when R is either a SNR set or a RISN set is $O((|A| + \log p)p|E|)$.

7 Upper Bounded Sets of User Preference Models

In earlier sections we restricted attention to weights vectors that satisfy the normalisation condition that all their components sum to one. An alternative is to consider a set of user preference models defined by an upper bound for each objective; for instance, where the restriction is only that each component is in the interval $[0, 1]$. In this section we show how to extend some of the results to this case.

For topologically closed $\mathcal{W} \subseteq I\!R^p_{\geq}$ and strictly positive vector $\sigma \in I\!R^p_{+}$ we define $\mathcal{W}^{(\sigma)}$ as $\{w \in \mathcal{W} : \forall i \in [p], w(i)\sigma(i) \leq 1\}$.

In particular, $\mathcal{W}^{(\mathbf{1})}$ is the subset of $\mathcal{W}$ such that each co-ordinate is bounded above by 1. For finite subset Λ of $I\!R^p$ we abbreviate $(\mathcal{W}_\Lambda)^{(\sigma)}$ to $\mathcal{W}_\Lambda^{(\sigma)}$.

Instead of the normalised set $\mathcal{W}_\Lambda^{\mathbf{1}}$ we can consider the upper-bounded set $\mathcal{W}_\Lambda^{(\mathbf{1})}$, which consists of all $w \in \mathcal{W}_\Lambda$ such all of its co-ordinates are bounded above by 1. We can thus consider the max regret function $MR(\alpha, A; \mathcal{W}_\Lambda^{(\mathbf{1})})$. Clearly, $\mathcal{W}_\Lambda^{\mathbf{1}} \subseteq \mathcal{W}_\Lambda^{(\mathbf{1})}$, and so, $MR(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}}) \leq MR(\alpha, A; \mathcal{W}_\Lambda^{(\mathbf{1})})$. Of interest then is how this is affected by rescaling of the objectives. That is, we can consider the effect of applying rescaling vector σ, leading to $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})})$.

Proposition 2 shows that $\mathcal{W}_{H_\sigma(\Lambda)}^{\mathbf{1}} = (\mathcal{W}_\Lambda^{\sigma}) \odot \sigma$. We have a similar result with bounding instead of normalisation.

Proposition 9 *For any rescaling vector $\sigma \in I\!R_+^p$, $\mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})} = \mathcal{W}_\Lambda^{(\sigma)} \odot \sigma$, and thus, $\mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})} \odot \sigma^{-1} = \mathcal{W}_\Lambda^{(\sigma)}$.*

Applying Proposition 3 to $\mathcal{W} = \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})}$ and applying Proposition 9 gives the following result. This is analogous to the property that $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{\mathbf{1}})$ equals $MR(\alpha, A; \mathcal{W}_\Lambda^{\sigma})$, from Theorem 1.

Theorem 4 *Let A and Λ be finite subsets of $I\!R^p$ and let $\alpha \in A$, and let $\sigma \in I\!R_+^p$. Then the max regret $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})})$ of the rescaled space is equal to $MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)})$.*

Example 6 *Continuing the running example, again with rescaling vector $\sigma = (1, \epsilon)$, we have that $\mathcal{W}_\Lambda^{(\sigma)}$ consists of all $w \in I\!R_\geq^p$ such that $w(1) \leq 1$, and $w(2) \leq 1/\epsilon$, and $w(2) \geq 4w(1)$. When $0 < \epsilon < 1/4$ we have four extreme points: $(0,0)$, $(0, 1/\epsilon)$, $(1, 4)$, and $(1, 1/\epsilon)$. Recall $\beta - \alpha = (24, -1)$; we obtain $MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)}) = (1,4) \cdot (\beta - \alpha) = 20$, and $MR(\beta, A; \mathcal{W}_\Lambda^{(\sigma)}) = (0, 1/\epsilon) \cdot (\alpha - \beta) = 1/\epsilon$.*

When $\frac{1}{4} \leq \epsilon \leq 1$ we have three extreme points: $(0,0)$, $(0, \frac{1}{\epsilon})$ and $(\frac{1}{4\epsilon}, \frac{1}{\epsilon})$. Then $MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)}) = (\frac{1}{4\epsilon}, \frac{1}{\epsilon}) \cdot (\beta - \alpha) = \frac{5}{\epsilon}$, which equals 5 when $\epsilon = 1$. $MR(\beta, A; \mathcal{W}_\Lambda^{(\sigma)}) = (0, 1/\epsilon) \cdot (\alpha - \beta) = 1/\epsilon$, which equals 1 when $\epsilon = 1$.

Comparing with Example 2 shows that these results for the bounded preference vectors are fairly close to those for the normalised preference vector case, especially when ϵ is small. For instance, we have $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})}) = \min(20, 5/\epsilon)$ for $\epsilon > 0$, and we have $MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{\mathbf{1}}) = \frac{20}{1+4\epsilon}$.

This example illustrates that, for the bounded case, the number of extreme points of the space of scenarios can vary with the rescaling vector σ, and so we do not have a simple bijection between the extreme points for $\mathcal{W}_\Lambda^{(\sigma)}$ and those of $\mathcal{W}_\Lambda^{(\mathbf{1})}$, in contrast with the normalised case (see Section 3.2).

However, the limiting behaviour for the bounded scenarios case is the same as that for the normalised scenarios case given in Section 5. The result below shows that, for arbitrary rescaling vectors σ, the ratio, $\frac{MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)})}{MR(\alpha, A; \mathcal{W}_\Lambda^{\sigma})}$ is always within the interval $[1, p]$.

Proposition 10 *Let A and Λ be finite subsets of $I\!R^p$ and let $\alpha \in A$, and let $\sigma \in I\!R_+^p$. Then*

$$MR(\alpha, A; \mathcal{W}_\Lambda^{\sigma}) \leq MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)}) \leq p\, MR(\alpha, A; \mathcal{W}_\Lambda^{\sigma}).$$

This then leads to the following result, which means that Theorem 2 in Section 5 can also be used for the bounded preferences case, to determine under which situations the ratios of max regrets can tend to infinity.

Proposition 11 $\frac{MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})})}{MR(H_\sigma(\gamma), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})})}$ *$\left(= \frac{MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)})}{MR(\gamma, A; \mathcal{W}_\Lambda^{(\sigma)})}\right)$ is unbounded over $\sigma \in I\!R_+^p$ if and only if $\frac{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR_\sigma(\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})}$ is unbounded over $\sigma \in I\!R_+^p$.*

Proof: Let $f(\sigma) = \frac{MR(H_\sigma(\alpha), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})})}{MR(H_\sigma(\gamma), H_\sigma(A); \mathcal{W}_{H_\sigma(\Lambda)}^{(\mathbf{1})})} = \frac{MR(\alpha, A; \mathcal{W}_\Lambda^{(\sigma)})}{MR(\gamma, A; \mathcal{W}_\Lambda^{(\sigma)})}$, and let $g(\sigma) = \frac{MR_\sigma(\alpha, A; \mathcal{W}_\Lambda^{\mathbf{1}})}{MR_\sigma(\gamma, A; \mathcal{W}_\Lambda^{\mathbf{1}})} = \frac{MR(\alpha, A; \mathcal{W}_\Lambda^{\sigma})}{MR(\gamma, A; \mathcal{W}_\Lambda^{\sigma})}$. Proposition 10 implies that for any σ, $\frac{f(\sigma)}{g(\sigma)}$ is in the interval $[1/p, p]$, which implies the result. □

8 Summary and Discussion

We pointed out an issue concerning max regret for weighted average preference models, with it being dependent on somewhat arbitrary choices of objective scales. We characterise the extreme cases, when the relative max regret (the ratio of max regret of an alternative to minimax regret) can become arbitrarily large (even for an alternative that initially minimises max regret).

This issue means that the foundations of max regret are weaker than might have been assumed. The situation is somewhat analogous to issues with Bayesian inference, especially when there is weak information about the prior, and where the issues can be ameliorated with robust methods, by considering a range of reasonable prior distributions.

Similarly, for max regret, a principled approach should thus consider reasonable ranges of rescalings. Our methods, in particular, in Section 6, then allow one to calculate the consequent upper and lower bounds on max regret for each alternative under consideration. They also allow one to determine how large the relative max regret can be, which can enable one to show that an alternative is never far from minimising max regret (assuming a given rescaling range).

A natural next step is to adapt max regret methods, such as for generating queries and conditions for terminating a dialogue, to be more robust to reasonable changes of objective scales, using the results developed here. It would also be interesting to explore the application of the technique from Section 6.3, replacing max regret with the maximum max regret over a normalised rectangular rescaling range, so that regret is maximised over both scenarios and rescalings. Like standard max regret, this could be used in the generation of queries (such as the current-solution strategy, e.g., page 3 [22]) or in a stopping condition in interactive preference elicitation, i.e. when the modified max regret is smaller than some threshold.

The issues with choices of scales of objectives analysed here clearly apply also to more general preference models such as MAUT [11] or GAI models [8, 4], and it would be worthwhile considering the extension of our results to those cases.

Acknowledgements

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant number 12/RC/2289-P2 at Insight the SFI Research Centre for Data Analytics at UCC, which is co-funded under the European Regional Development Fund; it has also been supported by the EU H2020 ICT48 project "TAILOR", under contract #952215. I also thank the reviewers for their constructive comments.

References

[1] N. Benabbou and P. Perny, 'Incremental weight elicitation for multiobjective state space search', in *Proc. AAAI 2015*, pp. 1093–1099, (2015).
[2] N. Benabbou and P. Perny, 'Adaptive elicitation of preferences under uncertainty in sequential decision making problems', in *Proc. IJCAI 2017*, pp. 4566–4572, (2017).
[3] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans, 'Constraint-based optimization and utility elicitation using the minimax decision criterion', *Artificial Intelligence*, (2006).
[4] Darius Braziunas and Craig Boutilier, 'Local utility elicitation in GAI models', in *UAI '05, Proceedings of the 21st Conference in Uncertainty in Artificial Intelligence, Edinburgh, Scotland, July 26-29, 2005*, pp. 42–49. AUAI Press, (2005).
[5] Darius Braziunas and Craig Boutilier, 'Minimax regret based elicitation of generalized additive utilities.', in *Proceedings of Uncertainty in Artificial Intelligence (UAI)*, (2007).
[6] D. G. Bridge and F. Ricci, 'Supporting product selection with query editing recommendations', in *Proc. RecSys-2007*, pp. 65–72. ACM, (2007).
[7] J. Figueira, S. Greco, and M. Ehrgott, *Multiple Criteria Decision Analysis—State of the Art Surveys*, Springer International Series in Operations Research and Management Science Volume 76, 2005.
[8] Christophe Gonzales and Patrice Perny, 'Gai networks for utility elicitation.', *KR*, **4**, 224–234, (2004).
[9] Vijay S Iyengar, Jon Lee, and Murray Campbell, 'Evaluating multiple attribute items using queries', in *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pp. 144–153, (2001).
[10] S. Kaddani, D. Vanderpooten, J. M. Vanpeperstraete, and H. Aissi, 'Weighted sum model with partial preference information: Application to multi-objective optimization', *European Journal of Operational Research*, **260**(2), 665–679, (2017).
[11] R. Keeney and H. Raiffa, *Decisions with multiple objectives: preferences and value tradeoffs*, Cambridge University Press, 1993.
[12] Tyler Lu and Craig Boutilier, 'Robust approximation and incremental elicitation in voting protocols', in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, (2011).
[13] R. Marinescu, A. Razak, and N. Wilson, 'Multi-objective constraint optimization with tradeoffs', in *Proc. CP-2013*, pp. 497–512, (2013).
[14] Mojtaba Montazery and Nic Wilson, 'Dominance and optimisation based on scale-invariant maximum margin preference learning', in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, ed., Carles Sierra, pp. 1209–1215. ijcai.org, (2017).
[15] Mojtaba Montazery and Nic Wilson, 'Scaling-invariant maximum margin preference learning', *Int. J. Approx. Reason.*, **128**, 69–101, (2021).
[16] Howard Raiffa, *Decision analysis*, Addison-Wesley, 1968.
[17] A. Salo and R. P. Hämäläinen, 'Preference programming – multicriteria weighting models under incomplete information', in *Handbook of Multicriteria Analysis*, pp. 167–187. Springer Berlin Heidelberg, (2010).
[18] Leonard J Savage, 'The theory of statistical decision', *Journal of the American Statistical association*, **46**(253), 55–67, (1951).
[19] Mikhail Timonin, 'Robust optimization of the Choquet integral', *Fuzzy Sets and Systems*, **213**, 27–46, (2013).
[20] Federico Toffano and Nic Wilson, 'Minimality and comparison of sets of multi-attribute vectors', in *Proceedings of European Conference on Artificial Intelligence (ECAI)*, pp. 913–920, (2020).
[21] P. Viappiani and C. Boutilier, 'Regret-based optimal recommendation sets in conversational recommender systems', in *Proc. RecSys-2009*, pp. 101–108. ACM, (2009).
[22] Paolo Viappiani and Craig Boutilier, 'On the equivalence of optimal recommendation sets and myopically optimal query sets', *Artificial Intelligence*, 103328, (2020).
[23] Tianhan Wang and Craig Boutilier, 'Incremental utility elicitation with the minimax regret decision criterion', in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 3, pp. 309–316, (2003).
[24] Nic Wilson, *The Sensitivity of Max Regret With Respect to Scaling of the Objectives (Longer Version)*, http://ucc.insight-centre.org/nwilson/SenstivityMaxRegretScalingLonger.pdf, 2023.