

Two Families of Scale-Invariant Regret, and their Monotonicity Properties

Nic Wilson^{a,*}

^aInsight Centre for Data Analytics, School of Computer Science and IT, University College Cork, Ireland
ORCID (Nic Wilson): <https://orcid.org/0000-0003-1874-8255>

Abstract. A weakness of max regret for multi-objective optimisation is its lack of scale-invariance: it is not robust to changes of units of the objectives. This paper defines two families of variations of regret that are scale-invariant as well as translation-invariant, and analyses their monotonicity properties. The first family is based on dividing regret by an aggregated regret (over all the alternatives), and the second family can be viewed as being like the regret of an aggregated alternative. Both families are shown to have good monotonicity properties.

1 Introduction

In a setting where there is uncertainty over the user preference model in an optimisation problem, max regret is used as a measure of how close an alternative is to being guaranteed to be optimal; the regret of an alternative with respect to a preference model is the difference in utility between that alternative and an optimal one.

The max regret of an alternative is always non-negative, and is only zero if the alternative is necessarily optimal, i.e., optimal with respect to all the considered preference models. Max regret is valuable, in particular, in interactive optimisation, for determining when the dialogue can be terminated, because an alternative can be recommended as it is sufficiently close to being optimal; and is used for generating informative user queries.

A common kind of preference model for multi-objective optimisation is a weighted sum of objective values, and there may be only partial information about the weights vector, leading to a set \mathcal{W} of compatible weights vectors. The set \mathcal{W} may be based on previous user inputs in the form of comparison queries, that one alternative is at least as good as another.

Max regret is translation-invariant, i.e., is not affected by translating all the multi-objective alternatives (corresponding to changing the zeros of the objective scales). However, as shown in [20], it is not scale-invariant: it is affected by changes of scales (i.e., the units) of the objectives. The reason for this is connected with the fact that in order for max regret to be finite, one needs to bound the set \mathcal{W} of weights vectors, for instance, by assuming each vector of weights sums to one. The chosen units can make a big difference to relative values of max regret, and can lead to a different alternative minimising max regret (see [20]). This indicates that max regret is somewhat unsatisfactory in this respect, since the choice of units for objectives can be rather arbitrary.

One way of attempting to deal with this problem is to reason based on families of rescalings; methods for this were developed in [20]. However, in [21] a neater and more fundamental approach is derived by changing the definition of max regret to make it scale-invariant whilst maintaining translation-invariance and other fundamental properties. This is done by dividing regret by an appropriate function D to form D -relative regret, with four appropriate functions D being defined there.

In this paper we generalise this idea by defining two kinds (or families) of function D that lead to translation- and scale-invariant D -relative [max] regret, both of which involve aggregating over the alternatives.

The first kind is based on defining the function D as an aggregation of the values of regret of the set of alternatives, where the aggregation function could, for instance, be a max, mean, or other Ordered Weighted Average.

The second kind can be viewed as defining D to be the regret of an artificial alternative that is a pointwise (i.e., co-ordinate-wise) aggregation of the set of alternatives, e.g., the min or mean of them.

A very important kind of property that was not considered in [21] relates to monotonicity; such a property could be lacking if a change that would seem to favour an alternative α could actually make its D -relative regret worse (i.e., larger). We formalise three different kinds of monotonicity, based on different types of change: (i) when α is improved, (ii) when other alternatives are worsened, but the maximum utility value is unchanged, and (iii) when the maximum utility is reduced.

We see that the third kind of monotonicity can only hold when there is an additional restriction, but that the first two kinds of monotonicity can hold more generally.

Summary of Contributions: we generate two families of scale-invariant and translation-invariant forms of regret: see Theorems 4 and 8; an instance of either of these two families can be used instead of regret, for instance in interactive user optimisation. We also show that both families have positive monotonicity properties, especially for smaller values of relative regret.

The structure of the paper is as follows. Section 3 gives the formal background, especially based on [21]. Section 4 defines certain kinds of aggregation operation that we will make use of in the new kinds of D -relative regret. The first family of D -relative regret, based on aggregating the values of regret of each alternative, is defined in Section 5, along with the invariance properties; and Section 6 does similarly for the second family based on pointwise aggregation. The three kinds of monotonicity property are defined in Section 7, and the

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

analysis for the two families is given in Sections 8 and 9. Section 10 considers monotonicity properties for another specific D -relative regret defined in [21], and Section 11 concludes.

2 Related Work

Max regret [16] is frequently used for decision problems with partial knowledge of the user preference model, for recommending alternatives, in generating queries, and in deciding when to terminate the user interaction, see e.g., [19, 4, 5, 12, 2, 18]. Simple linear preference models are a commonly applied special case of Multi-Attribute Utility Theory (MAUT) [14, 10, 7] for multi-objective reasoning. As a model for unknown user preferences, using a weighted average with unknown weights, they have been used for instance in [8, 6, 17, 15, 13, 9, 3].

As mentioned above, this paper extends the work in [21], which was motivated by the work in [20], which pointed out the lack of scale-invariance of max regret, and developed approaches for reasoning with families of rescalings, as well as characterising situations in which the ratio of the max regret of two alternatives can be arbitrarily large (for extreme scalings).

3 Some Background

We give some background, in particular, using the notation and definitions from [21]. We are considering multi-objective decision making problems involving a choice between a finite set A of alternatives. Each alternative is assumed to be scored according to each of p objectives (with higher scores being better). Thus, each alternative is associated with a (multi-objective utility) vector α in \mathbb{R}^p ; to simplify notation we identify the alternative with α .

We assume that the user preference model is one of a family parameterised by a set of scenarios \mathcal{W} ; associated with $w \in \mathcal{W}$ is a real-valued utility function f_w over the space \mathbb{R}^p of multi-objective vectors. In this paper we focus on the commonly used simple linear model of the user preferences based on (non-negative) weighted sums of objective values: for all $\alpha \in \mathbb{R}^p$, $f_w(\alpha) = w \cdot \alpha = \sum_{i=1}^p w(i)\alpha(i)$. Here, weights vector w is an element of \mathbb{R}_{\geq}^p , where \mathbb{R}_{\geq}^p is the set of vectors in \mathbb{R}^p , such that each co-ordinate is non-negative i.e., $w(i) \geq 0$ for all $i \in \{1, \dots, p\}$.

We have partial knowledge of the user preferences, expressed by a subset \mathcal{W} of all possible weights vectors \mathbb{R}_{\geq}^p . For a given weights vector w let $O_w(A)$ be the set of all alternatives α in A that are optimal in A with respect to w , i.e., such that $w \cdot \alpha \geq w \cdot \beta$ for all alternatives $\beta \in A$. We say that alternative $\alpha \in A$ is *necessarily optimal*, written $\alpha \in \text{NO}_{\mathcal{W}}(A)$, if α is optimal in A with respect to every weights vector in \mathcal{W} . If there exists a necessarily optimal alternative α then we can recommend it to the user, since it is optimal with respect to the user's (unknown) weights vector $w \in \mathcal{W}$. However, frequently, there will not be a necessarily optimal alternative; we could then recommend an alternative, for example, one minimising max regret; or we could elicit more information from the user.

Preference Cone \mathcal{V}_Λ : For set $\Lambda \subseteq \mathbb{R}^p$ of utility vectors the preference cone $\mathcal{V}_\Lambda \subseteq \mathbb{R}_{\geq}^p$ is defined to be the set of weights vectors $w \in \mathbb{R}_{\geq}^p$ such that $0 \leq w \cdot \lambda (= \sum_{i=1}^p w(i)\lambda(i))$ for all $\lambda \in \Lambda$. In particular, each λ may come from the answer to a comparison query, preferring an alternative α_λ over an alternative β_λ , implying $w \cdot (\alpha_\lambda - \beta_\lambda) \geq 0$, where $\alpha_\lambda - \beta_\lambda = \lambda$.

Cone and Conic Closure: More generally, a (linear) cone (in \mathbb{R}_{\geq}^p) \mathcal{V} is defined to be a subset of \mathbb{R}_{\geq}^p such that if $w \in \mathcal{V}$ then $rw \in \mathcal{V}$ for all real $r > 0$. For any $\mathcal{W} \subseteq \mathbb{R}_{\geq}^p$, the conic closure $\mathcal{C}(\mathcal{W})$ is defined to be $\{rw : r > 0, w \in \mathcal{W}\}$; this is the unique smallest cone containing \mathcal{W} .

3.1 Max regret

The regret $R_A^w(\alpha)$ and the utility/value $\text{Val}_A(w)$ of A , w.r.t. preference model w : Let $w \in \mathbb{R}_{\geq}^p$, i.e., a vector with all components being non-negative. For finite set A of alternatives, the value (or utility) $\text{Val}_A(w)$ of A with respect to weights vector w is defined to be maximum utility over all elements of A , i.e., $\max_{\beta \in A} w \cdot \beta$. The regret $R_A^w(\alpha)$ for alternative $\alpha (\in A)$ in scenario w is defined by $R_A^w(\alpha) = \text{Val}_A(w) - w \cdot \alpha$; this is thus non-negative and measures how far short α is from being optimal in A with respect to w . Hence, α is optimal with respect to user preference model w (i.e., $\alpha \in O_w(A)$) if and only if $R_A^w(\alpha) = 0$.

Max regret: For arbitrary $\mathcal{W} \subseteq \mathbb{R}_{\geq}^p$, and alternative α in set of alternatives A , we can define the max regret $MR_A(\alpha; \mathcal{W})$ of α to be equal to $\sup_{w \in \mathcal{W}} R_A^w(\alpha)$. It is thus in $[0, \infty]$, i.e., a non-negative real or equal to infinity. We have $MR_A(\alpha; \mathcal{W}) = 0$ if and only if α is necessarily optimal, i.e., α is optimal in A with respect to every user model in \mathcal{W} .

The basic result below considers the effect on regret of different operations on the inputs. Part (i) shows the effect of a simple scaling of the weights vector, where for $r > 0$, the vector $rw \in \mathbb{R}_{\geq}^p$ is given by $(rw)(i) = r(w(i))$ for $i \in \{1, \dots, p\}$. (ii) relates to translating the objective values by a vector $\delta \in \mathbb{R}^p$, which corresponds with redefining the zeros of the objective values. $A + \delta$ is defined to be the set $\{\gamma + \delta : \gamma \in A\}$. Part (iii) relates with changing the scales of the objectives, i.e., changing their units. For $\alpha, \tau \in \mathbb{R}^p$, their pointwise product $\alpha \odot \tau (\in \mathbb{R}^p)$ is defined by $(\alpha \odot \tau)(i) = \alpha(i)\tau(i)$ for $i \in \{1, \dots, p\}$. Also, $A \odot \tau$ is defined to be the set $\{\gamma \odot \tau : \gamma \in A\}$.

Proposition 1. Consider any finite $A \subseteq \mathbb{R}^p$, $\alpha \in A$, and $w \in \mathbb{R}_{\geq}^p$ (i.e., each co-ordinate non-negative), and $r > 0$ and $\delta \in \mathbb{R}^p$ and $\tau \in \mathbb{R}_{+}^p$ (i.e., each co-ordinate strictly positive).

- (i) $R_A^{rw} = r \times R_A^w(\alpha)$;
- (ii) $R_{A+\delta}^w(\alpha + \delta) = R_A^w(\alpha)$; and
- (iii) $R_{A \odot \tau}^{w \odot \tau^{-1}}(\alpha \odot \tau) = R_A^w(\alpha)$.

Part (iii) follows using the fact that for any $\gamma \in \mathbb{R}^p$, $(w \odot \tau^{-1}) \cdot (\gamma \odot \tau) = w \cdot \gamma$. Part (ii) implies that regret and max regret are translation-invariant: i.e., if we add a vector δ to each alternative then the regret and max regret are unchanged. Essentially because of part (i), the set \mathcal{W} will need to be bounded, or else the values of max regret will be zero or infinity. For instance we might consider just weights vectors whose components sum to one. However, as shown in [20, 21], this causes max regret to not be scale-invariant.

3.2 Max relative regret:

A variation of max regret is max relative regret [11, 1] which can be written in our (maximising utility) setting as $\sup_{w \in \mathcal{W}} \frac{\text{Val}_A(w) - w \cdot \alpha}{\text{Val}_A(w)}$. This is scale-invariant but not translation-invariant.

3.3 D -relative regret and some basic properties

We describe here a formalism from [21] where a more general form of [max] relative regret is defined, that generates measures that are

translation- and scale-invariant. Instead of dividing the regret by $Val_A(w)$ (as in done for relative regret), it is divided by a term $D_A(w)$ that depends on the set A of alternatives, but not on α in particular. A key property is positive homogeneity with respect to w .

Sets \mathcal{D} and \mathcal{D}_n of Denominator families: functions in \mathcal{D} are families D , of non-negative real-valued functions D_A on \mathbb{R}_{\geq}^p for all finite $A \subseteq \mathbb{R}^p$, satisfying the positive homogeneity condition $D_A(rw) = rD_A(w)$ for all real $r > 0$ and all $w \in \mathbb{R}_{\geq}^p$.

We define \mathcal{A} to be the set of finite subsets of \mathbb{R}^p , and for each positive integer n we define \mathcal{A}_n to be the set of all subsets of \mathbb{R}^p of cardinality n .

In the formalisms developed in this paper based on aggregation functions it is convenient to focus on input sets A of alternatives of a given cardinality $n \geq 1$, i.e., $A \in \mathcal{A}_n$, and the associated families D^n (or ‘slices’ of a denominator function D). Thus, a family/function in \mathcal{D}_n takes as input a set $A \in \mathcal{A}_n$ and a weights vector w and returns a non-negative real value, whilst satisfying the positive homogeneity property.

Let $\alpha \in A$ and $w \in \mathbb{R}_{\geq}^p$; recall that $R_A^w(\alpha) = Val_A(w) - w \cdot \alpha = \max_{\beta \in A} w \cdot \beta - w \cdot \alpha$. Define the D -relative regret as follows. If $D_A(w) > 0$ we define

$$RR_A^D(\alpha; w) = \frac{R_A^w(\alpha)}{D_A(w)}. \quad (1)$$

We define $RR_A^D(\alpha; w) = 0$ if $R_A^w(\alpha) = D_A(w) = 0$; and $RR_A^D(\alpha; w) = \infty$ if $R_A^w(\alpha) > 0$ and $D_A(w) = 0$.

We define, for $\mathcal{W} \subseteq \mathbb{R}_{\geq}^p$, the *max D -relative regret*

$$MRR_A^D(\alpha; \mathcal{W}) = \sup_{w \in \mathcal{W}} RR_A^D(\alpha; w).$$

So, the max D -relative regret is the supremum, over weights vectors $w \in \mathcal{W}$, of the ratio between the regret $R_A^w(\alpha)$ and the value of the denominator function $D_A(w)$. It takes values in $[0, \infty]$.

Max D -relative regret of α is non-negative, equalling zero if and only if α is necessarily optimal (see Proposition 2 of [21]).

For any $w \in \mathbb{R}_{\geq}^p$ and $r > 0$, $R_A^{rw}(\alpha) = rR_A^w(\alpha)$ and $D_A(rw) = rD_A(w)$, and so $RR_A^D(\alpha; rw) = RR_A^D(\alpha; w)$. This implies that $MRR_A^D(\alpha; \mathcal{W}) = MRR_A^D(\alpha; \mathcal{C}(\mathcal{W}))$, where $\mathcal{C}(\mathcal{W})$ is the conic closure of \mathcal{W} . This means that we can directly use a preference cone of the form \mathcal{V}_Λ , and that it is equivalent to considering normalised or bounded versions of it.

3.4 Four choices of function D

Four particular denominator functions, D^k for $k = 0, 1, 2, 3$, are defined in [21], and for each k , D^k -relative regret is translation-invariant and scale-invariant, and hence, so are the associated D^k -relative regret and max D^k -relative regret measures, by Proposition 1 and Equation 1. In addition they satisfy the inequalities $0 \leq D_A^0(w) \leq D_A^1(w) \leq D_A^2(w) \leq D_A^3(w)$, and, for $k = 1, 2, 3$, RR^{D^k} and MRR^{D^k} have an upper bound of 1.

Given finite $A \subseteq \mathbb{R}^p$, we define the pointwise max, $\bar{A} \in \mathbb{R}^p$, by, for each $i \in \{1, \dots, p\}$, $\bar{A}(i) = \max_{\alpha \in A} \alpha(i)$; the pointwise min, $\underline{A} \in \mathbb{R}^p$, by $\underline{A}(i) = \min_{\alpha \in A} \alpha(i)$; and the mean, $\hat{A} \in \mathbb{R}^p$, to be $\frac{1}{|A|} \sum_{\alpha \in A} \alpha$. Recall $Val_A(w) = \max_{\alpha \in A} w \cdot \alpha$; similarly we define $Wor_A(w) = \min_{\alpha \in A} w \cdot \alpha$, which is the worst utility (with respect to w) of any element of A .

For finite $A \subseteq \mathbb{R}^p$ and $w \in \mathbb{R}_{\geq}^p$:

$$D_A^0(w) = Val_A(w) - w \cdot \hat{A} = \text{mean}_{\gamma \in A} R_A^w(\gamma).$$

$$\begin{aligned} D_A^1(w) &= Val_A(w) - Wor_A(w) = \max_{\gamma \in A} R_A^w(\gamma). \\ D_A^2(w) &= Val_A(w) - w \cdot \underline{A}. \\ D_A^3(w) &= w \cdot (\bar{A} - \underline{A}). \end{aligned}$$

Hence D^0 and D^1 involve different ways of aggregating regret over the alternatives, i.e., using mean or max. And D^0 and D^2 can be viewed as the regret of an aggregated alternative, using pointwise mean \hat{A} for D^0 and pointwise min \underline{A} for D^2 . Using D^3 , max D -relative regret corresponds exactly with max regret if one first rescales A to have unit ranges for each objective: see Proposition 15 in [21].

3.5 Translation-invariance

As discussed in Section 3.1, max regret satisfies translation-invariance, which means that the max regret is unchanged if we add a constant vector to each alternative in A . Because of this, the objective scales need only be interval scales (rather than ratio scales). This generalises to D -relative [max] regret.

We say that RR^D is *translation-invariant* if for all finite $A \subseteq \mathbb{R}^p$ and $\alpha \in A$ and all $\delta \in \mathbb{R}^p$ and $w \in \mathbb{R}_{\geq}^p$, $RR_{A+\delta}^D(\alpha + \delta; w) = RR_A^D(\alpha; w)$.

Similarly, MRR^D is *translation-invariant* if for all finite $A \subseteq \mathbb{R}^p$ and $\alpha \in A$ and all $\delta \in \mathbb{R}^p$ and $\mathcal{W} \subseteq \mathbb{R}_{\geq}^p$, $MRR_{A+\delta}^D(\alpha + \delta; \mathcal{W}) = MRR_A^D(\alpha; \mathcal{W})$.

D is *translation-invariant* if for all finite $A \subseteq \mathbb{R}^p$ and all $\delta \in \mathbb{R}^p$, and $w \in \mathbb{R}_{\geq}^p$, $D_{A+\delta}(w) = D_A(w)$. This is a sufficient condition for translation-invariance of RR^D and MRR^D , by Proposition 1 (ii) and Equation 1.

3.6 Scale-invariance

Changing the scales of the objectives involves multiplying the values of the objectives by certain scalars, which can be represented by a pointwise product with a strictly positive vector τ . As well as changing the alternatives with the mapping $\gamma \mapsto \gamma \odot \tau$ this changes the inputs Λ for the preference model set \mathcal{V}_Λ , and, as shown in [20, 21], this amounts to mapping each weights vector w to $w \odot \tau^{-1}$; this ensures that the utility, $w \cdot \gamma$, of γ with respect to preference model w , is unchanged.

It is shown in [21] that, for $D \in \mathcal{D}$, scale-invariance can be written in terms of equalities of the form $MRR_{A \odot \tau}^D(\alpha \odot \tau; \mathcal{V}_\Lambda \odot \tau^{-1}) = MRR_A^D(\alpha; \mathcal{V}_\Lambda)$. A sufficient condition for this is that RR^D is scale-invariant, i.e., for all finite $A \subseteq \mathbb{R}^p$ and $\alpha \in A$ and all $\delta \in \mathbb{R}^p$ and $w \in \mathbb{R}_{\geq}^p$, $RR_{A \odot \tau}^D(\alpha \odot \tau; w \odot \tau^{-1}) = RR_A^D(\alpha; w)$. By Proposition 1 (iii) and Equation 1, a sufficient condition is that D is scale-invariant, i.e., if, for all A and w , and for all $\tau \in \mathbb{R}_{\geq}^p$, $D_{A \odot \tau}(w \odot \tau^{-1}) = D_A(w)$. We then have (see Propositions 10 and 13 of [21]):

Proposition 2. *Let D be any family in \mathcal{D} or \mathcal{D}_n .*

If D is translation-invariant then RR^D and MRR^D are translation-invariant.

If D is scale-invariant then RR^D and MRR^D are scale-invariant.

4 Aggregation Operations

The two forms of denominator families that we define here both involve aggregation: the first family, aggregation of regrets, and the second family, co-ordinate-wise aggregation of alternatives; in particular, we use aggregation functions on multisets of reals of cardinality n where $n = |A|$. This section describes the relevant families of aggregation functions, including Ordered Weighted Averages,

based on the properties we will need in Sections 5-9 to ensure good properties of the denominator families.

4.1 Multisets of reals

Let \mathcal{M}_n be the set of all multisets, of cardinality n , of real numbers. We consider real-valued aggregation operations F on \mathcal{M}_n , such as mean, or minimum, or maximum. Hence, if $S \in \mathcal{M}_n$ then $F(S) \in \mathbb{R}$.

If $t \in \mathbb{R}$ and $S \in \mathcal{M}_n$ we define $tS \in \mathcal{M}_n$ to be $\{ts : s \in S\}$, and $S + t \in \mathcal{M}_n$ to be $\{s + t : s \in S\}$.

For $S \in \mathcal{M}_n$, let $S^{(1)}$ be $\min(S)$, i.e., the smallest element of S , let $S^{(n)}$ be $\max(S)$, the largest element, and, more generally, for $j = 1, \dots, n$, let $S^{(j)}$ be the j^{th} smallest element, so that $S = \{S^{(1)}, \dots, S^{(n)}\}$ and $S^{(1)} \leq \dots \leq S^{(n)}$.

There is a natural Pareto-like ordering on \mathcal{M}_n : for $S, T \in \mathcal{M}_n$ we define the relation \leq on \mathcal{M}_n by $S \leq T$ if and only if for all $j = 1, \dots, n$, $S^{(j)} \leq T^{(j)}$. Then \leq is a partial order. $S \leq T$ if and only if one can generate T by increasing elements of S . As usual, we write $S < T$ to mean $S \leq T$ and $S \neq T$.

4.2 Properties of particular aggregation operations

We consider some particular properties that particular aggregation operations may have, and that will be used in later sections.

We say that aggregation function F

—satisfies **Positive Homogeneity** if for all $S \in \mathcal{M}_n$ and for all $r > 0$, $F(rS) = rF(S)$.

—satisfies **Monotonicity** if $F(S) \leq F(T)$ holds for any $S, T \in \mathcal{M}_n$ such that $S \leq T$, i.e., such that T Pareto-dominates S , and satisfies **Strict Monotonicity** if $F(S) < F(T)$ holds for any $S, T \in \mathcal{M}_n$ such that $S < T$.

—satisfies **Translation** if for all $S \in \mathcal{M}_n$ and for all $t \in \mathbb{R}$ we have $F(S + t) = F(S) + t$.

Consider F satisfying positive homogeneity and monotonicity. If no element of S is negative then, since $2S \geq S$, we have $2F(S) = F(2S) \geq F(S)$ and so $F(S) \geq 0$. Similarly, if no element of S is positive then $2S \leq S$ so we have $2F(S) = F(2S) \leq F(S)$ and thus, $F(S) \leq 0$. If F also satisfies Translation then this implies that for any $S \in \mathcal{M}_n$, $\min(S) \leq F(S) \leq \max(S)$.

Definition of \mathcal{G}_n : We let \mathcal{G}_n be the set of all real-valued aggregation functions on \mathcal{M}_n that satisfy Positive Homogeneity, Monotonicity and Translation. Hence, for all $G \in \mathcal{G}_n$ and any $S \in \mathcal{M}_n$, $\min(S) \leq F(S) \leq \max(S)$.

For the pointwise aggregation methods (to be defined in Section 6) we will also assume the following property.

For real-valued aggregation function F on \mathcal{M}_n , we say that F is **mean-dominated** if for all $S \in \mathcal{M}_n$, $F(S) \leq \hat{S}$, i.e., $F(S)$ is never more than the mean of S .

Definition of \mathcal{H}_n : We define \mathcal{H}_n to the set of functions in \mathcal{G}_n that are mean-dominated.

4.3 Ordered Weighted Averages (OWAs)

A classic family of aggregation operations which include mean, min and max are Ordered Weighted Averages [22].

Let \mathcal{I}_n be the set of vectors in \mathbb{R}^n such that each component is non-negative, and the components sum to one. Thus, if $\vec{t} \in \mathcal{I}_n$ then for all $j = 1, \dots, n$, $\vec{t}(j) \geq 0$, and $\sum_{j=1}^n \vec{t}(j) = 1$.

For $\vec{t} \in \mathcal{I}_n$, we define the Ordered Weighted Average (OWA) aggregation operation $F^{\vec{t}}$ as follows: for $S \in \mathcal{M}_n$, $F^{\vec{t}}(S) = \sum_{j=1}^n \vec{t}(j)S^{(j)}$. Thus, $F^{\vec{t}}(S)$ is a weighted mean of the elements of S , but where the weights are dependent on the order of elements of S , with weight $\vec{t}(j)$ being the weight assigned to the j^{th} smallest element.

In particular, if $\vec{t} = (1, 0, \dots, 0)$ then $F^{\vec{t}}(S) = S^{(1)} = \min(S)$. And if \vec{t} is the uniform vector $(\frac{1}{n}, \dots, \frac{1}{n})$ then $F^{\vec{t}}(S) = \frac{1}{n} \sum_{j=1}^n S^{(j)} = \frac{1}{n} \sum_{s \in S} s$, i.e., the mean \hat{S} of S .

OWAs satisfy Positive Homogeneity, Monotonicity and Translation, and there is a simple characterisation of when they are mean-dominated:

Proposition 3. (i) For any $\vec{t} \in \mathcal{I}_n$, the OWA $F^{\vec{t}}$ satisfies Positive Homogeneity, Monotonicity and Translation.

(ii) $F^{\vec{t}}$ is mean-dominated if and only if for all $k \in \{1, \dots, n\}$, $\sum_{j=1}^k t_j \geq \frac{k}{n}$.

For multiset $S \in \mathcal{M}_n$ we say that S' is a *single increase of S of magnitude r* ($r > 0$) if, for some $s \in S$, $S' = S \setminus \{s\} \cup \{s + r\}$, i.e., S is unchanged except for (one copy of) s becoming $s + r$.

Let F be a real-valued function on a multiset of n (not necessarily different) reals. For $t > 0$ we say that F is t -bounded if for any multiset $S \in \mathcal{M}_n$ and any single increase of S of magnitude r we have $F(S') - F(S) \leq tr$, increasing just one element of S by r leads to at most an increase of tr in $F(S)$.

For an OWA $F^{\vec{t}}$, increasing an element of S by a small amount ϵ leads to an increase in $F^{\vec{t}}(S)$ of $t_j \epsilon$ for some $j \in \{1, \dots, n\}$, and more generally, for any single increase of S of magnitude r to from S' , we have $F^{\vec{t}}(S') - F^{\vec{t}}(S) \leq t_{max} r$, where $t_{max} = \max\{t_j : j = 1, \dots, n\}$. Hence, $F^{\vec{t}}$ is t_{max} -bounded.

5 D-Relative Regret Based on Aggregation/Normalisation of Regrets

For our first kind of function, $D_A(w)$ is an aggregation G of the regrets of all the alternatives in A .

Recall that $R_A^w(\gamma) = Val_A(w) - f_w(\gamma)$ is the regret of γ in A with respect to w . We write $R_A^w(A)$ as an abbreviation for the multiset $\{R_A^w(\gamma) : \gamma \in A\} \in \mathcal{M}_n$.

Recall from Section 4.2 that \mathcal{G}_n is the set of aggregation functions satisfying Positive Homogeneity, Monotonicity and Translation; in particular, by Proposition 3(i), this includes all the Ordered Weighted Averages.

Consider any $G \in \mathcal{G}_n$. We define D^G as follows. For any $A \in \mathcal{A}_n$ and $w \in \mathbb{R}_>^n$, define $D_A^G(w)$ to be $G(R_A^w(A))$.

Thus, $D_A^G(w)$ is the aggregated regret over all alternatives in A . For example, if G is the mean then D^G is equal to D^0 defined in Section 3.4, and if G is the maximum then D^G is equal to D^1 .

Positive homogeneity of G , and the fact that $R_A^w(\gamma) = r \times R_A^w(\gamma)$ (see Proposition 1(i)), implies positive homogeneity of D^G . Also,

since $R_A^w(A)$ has minimum element of zero, and G satisfies positive homogeneity and monotonicity, $G(R_A^w(A)) \geq 0$. Hence, for any $G \in \mathcal{G}_n$, $D^G \in \mathcal{D}_n$, i.e., D^G is a denominator function.

The associated D^G relative regret function is then given by, for $\alpha \in A$,

$$RR_A^{D^G}(\alpha; w) = \frac{R_A^w(\alpha)}{G(R_A^w(A))}, \quad (2)$$

with division of zero being treated in the same way as in the definition of RR_A^D in Section 3.3.

Because of the translation property of G , we have that $G(R_A^w(A))$, i.e., $G(\{R_A^w(\gamma) : \gamma \in A\})$, is equal to $Val_A(w) + G(\{-w \cdot \gamma : \gamma \in A\})$, and so, $RR_A^{D^G}(\alpha; w) \leq 1$ if and only if $w \cdot \alpha \geq -G(\{-w \cdot \gamma : \gamma \in A\})$, i.e., the utility of α is at least the aggregate utility.

5.1 Basic results for G -systems

When $G(R_A^w(A)) > 0$, positive homogeneity of G implies $G(\{RR_A^{D^G}(\gamma; w) : \gamma \in A\}) = \frac{1}{G(R_A^w(A))}G(\{R_A^w(\gamma) : \gamma \in A\}) = 1$, so we have the following normalisation property, stating that the G -aggregation of the D^G -relative regret values of all the alternatives is equal to 1:

When $G(R_A^w(A)) > 0$,

$$G(\{RR_A^{D^G}(\gamma; w) : \gamma \in A\}) = 1. \quad (3)$$

Consider any $A \in \mathcal{A}_n$ and $\delta \in \mathbb{R}^p$, and $w \in \mathbb{R}_{\geq}^p$. Mapping A to $A + \delta (= \{\gamma + \delta : \gamma \in A\})$ does not change the values of regret, by Proposition 1(ii). This immediately implies the translation-invariance of D^G , i.e., $D_{A+\delta}^G(w) = D_A^G(w)$, which by Proposition 2 implies that RR^{D^G} and MRR^{D^G} are translation-invariant.

Similarly, by Proposition 1(iii), for any $\tau \in \mathbb{R}_{+}^p$, and any $\gamma \in A$, $R_{A \odot \tau}^{w \odot \tau^{-1}}(\gamma \odot \tau) = R_A^w(\gamma)$ which immediately implies that scale-invariance of D^G : and hence, by Proposition 2, that RR^{D^G} and MRR^{D^G} are scale-invariant.

To summarise we have:

Theorem 4. For any $G \in \mathcal{G}_n$, $D^G \in \mathcal{D}_n$ and RR^{D^G} and MRR^{D^G} are translation-invariant and scale-invariant.

6 D -Relative Regret Based on Regret of Pointwise-Aggregated Alternative

The second kind of function involves pointwise aggregation of the alternatives, and computing the ‘regret’ value of such an aggregated alternative.

For $A \in \mathcal{A}_n$ and $i \in \{1, \dots, p\}$, $A(i)$ is defined to be the multiset $\{\gamma(i) : \gamma \in A\}$. For real-valued function H on \mathcal{M}_n define:

- The vector $\theta_A^H \in \mathbb{R}^p$ by, for $i \in \{1, \dots, p\}$, $\theta_A^H(i) = H(A(i))$.
- The function $D^{(H)}$, by for $A \in \mathcal{M}_n$ and $w \in \mathbb{R}_{\geq}^p$ $D_A^{(H)}(w) = Val_A(w) - w \cdot \theta_A^H$.

If we extend the function R_A^w to apply to any element $\theta \in \mathbb{R}^p$, i.e., $R_A^w(\theta) = Val_A(w) - w \cdot \theta$, then $D_A^{(H)}(w) = R_A^w(\theta_A^H)$, and so we have that the associated $D^{(H)}$ relative regret function is given by, for $\alpha \in A$,

$$RR_A^{D^{(H)}}(\alpha; w) = \frac{R_A^w(\alpha)}{R_A^w(\theta_A^H)} = \frac{Val_A(w) - w \cdot \alpha}{Val_A(w) - w \cdot \theta_A^H}, \quad (4)$$

with division of zero being treated in the same way as in the definition of RR_A^D in Section 3.3.

If we extend the definitions of $RR_A^{D^{(H)}}$ and $MRR_A^{D^{(H)}}$ to apply to an arbitrary element of \mathbb{R}^p then Equation 4 gives the normalisation conditions: $RR_A^{D^{(H)}}(\theta_A^H; w) = 1$ and $MRR_A^{D^{(H)}}(\theta_A^H; \mathcal{W}) = 1$. Also, $MRR_A^{D^{(H)}}(\alpha; \mathcal{W}) \leq 1$ if and only if α always has at least as high utility as θ_A^H ; in particular, this holds if α Pareto-dominates θ_A^H .

6.1 Some basic results for pointwise aggregation

The following result shows that mean-dominated H leads to a valid, i.e., non-negative, function $D^{(H)}$, that is, $w \cdot \theta_A^H \leq Val_A(w)$. Because of mean-dominance, it is sufficient to show this for H being the mean, i.e., $H(S) = \hat{S}$ for $S \in \mathcal{M}_n$, so that $\theta_A^H = \hat{A}$. But then $w \cdot \theta_A^H$ can be seen to be equal to the mean over $\gamma \in A$ of $w \cdot \gamma$, which is at most $Val_A(w)$, the max over $\gamma \in A$ of $w \cdot \gamma$. Hence we have:

Proposition 5. If H is mean-dominated then for all $A \in \mathcal{A}_n$, $D_A^{(H)}$ is non-negative.

Proposition 5 above shows that H being mean-dominated is a sufficient condition for $D_A^{(H)}$ being non-negative (which is required, by the definition of a denominator family). Proposition 6 below shows that in fact, in a certain sense, it is also necessary.

Proposition 6. If H is not mean-dominated then there exists p and finite $A \subseteq \mathbb{R}^p$ of cardinality n and $w \in \mathbb{R}_{\geq}^p$ such that $D_A^{(H)}(w) < 0$.

We also have the desired invariance properties of $D^{(H)}$:

Proposition 7. For any $H \in \mathcal{H}_n$, $D^{(H)}$ is translation-invariant and scale-invariant.

These results lead to the following summary of basic properties.

Theorem 8. : For any $H \in \mathcal{H}_n$, we have $D^{(H)} \in \mathcal{D}_n$, and $RR^{D^{(H)}}$ is translation-invariant and scale-invariant.

7 Monotonicity properties

These properties relate to whether the D -relative regret is non-increasing if we change the inputs A in a way that favours the alternative α . We distinguish three different kinds of changes of A , the first where we increase one or more co-ordinates of α , and the other two relating to reductions in other alternatives.

7.1 Definitions of Types of Monotonicity

The monotonicity properties all involve a condition that the D -relative regret of alternative α is not increased by modifying A in some way to form another set of alternatives A' , where the modification might seem to favour α over other alternatives. We consider three kinds of modifications, type-(I), type-(II) and type-(III).

(I)-monotonicity relates to increasing $\alpha(i)$ for some i and leaving everything else fixed.

(II)-monotonicity relates to decreasing $\gamma(i)$ for some $\gamma \neq \alpha$ and leaving everything else fixed, without changing $Val_A(w)$.

(III)-monotonicity relates to decreasing $\gamma(i)$ for some $\gamma \in A \setminus \{\alpha\}$ and leaving everything else fixed such that $Val_A(w)$ is reduced (and γ continues as the optimal alternative).

Our monotonicity properties involve considering the change of D -relative regret when we change elements of A . For this purpose we consider 1-1 mappings from A onto a subset A' of \mathbb{R}^p . We call these (A, α) -modifications from A to A' , and, to simplify the notation, the image of each $\gamma \in A$ is written as γ' (which may be equal to γ).

A type-(I) (A, α) -modification is an (A, α) -modification such that α is Pareto-increased, i.e., $\alpha' >_P \alpha$, and no other element of A is changed, i.e., $\gamma' = \gamma$ for all $\gamma \neq \alpha$.

A type-(II) (A, α) -modification given $w \in \mathbb{R}_{\geq}^p$ is an (A, α) -modification of A to A' such that $\alpha' = \alpha$, and there exists $\beta \in O_w(A)$ such that, $\beta' = \beta$, and for all $\gamma \in A$, either $\gamma' = \gamma$ or γ Pareto-dominates γ' .

Hence α is unchanged, and the only elements that are changed are Pareto-reduced, but at least one optimal element is unchanged, so that $Val_{A'}(w) = Val_A(w)$.

A type-(III) (A, α) -modification given w only applies if α is not optimal in A . It is an (A, α) -modification of A to A' such that α is unchanged, and only the optimal elements of A are changed, and these remain optimal in A' , i.e., such that $\alpha' = \alpha$ and $\gamma' = \gamma$ for all $\gamma \in A \setminus O_w(A)$, and for all $\beta \in O_w(A)$, β' is Pareto-dominated by β , $\beta >_P \beta'$ and $\beta' \in O_w(A')$.

For (X) being either (I), (II) or (III), we say that the D -relative regret function RR^D satisfies (X)-monotonicity if for all finite $A \subseteq \mathbb{R}^p$ and $\alpha \in A$ and $w \in \mathbb{R}_{\geq}^p$, and any type-(X) (A, α) -modification [given w] of A to A' we have $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$, i.e., the change does not increase the D -relative regret.

For $t > 0$ we also say that RR^D satisfies $(\leq t)$ -monotonicity if for all finite $A \subseteq \mathbb{R}^p$ and $\alpha \in A$ and $w \in \mathbb{R}_{\geq}^p$, and any type-(X) (A, α) -modification of A to A' we have $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$ whenever $RR_A^D(\alpha; w) \leq t$.

We say that RR^D satisfies restricted (X)-monotonicity if it satisfies (≤ 1) -monotonicity.

The reason for separating (II)-monotonicity and (III)-monotonicity, i.e., distinguishing between when one decreases $Val_A(w)$ and when one doesn't, is that (III)-monotonicity can be seen to be somewhat in conflict with the kind of normalised relative regret that we are considering here; in particular, consider the D^G -relative regret $RR_A^{D^G}$, for some aggregation function $G \in \mathcal{G}_n$, which satisfies the normalisation condition $G(\{RR_A^{D^G}(\gamma; w) : \gamma \in A\}) = 1$ (Equation 3). (III)-monotonicity for $RR_A^{D^G}$ would imply that $RR_A^{D^G}(\alpha; w)$ would not increase; but this can be applied for any non-optimal α , and so for every $\gamma \in A$, $RR_A^{D^G}(\gamma; w)$ would not increase (since if γ is optimal, the value stays at zero). However, despite this Pareto decrease of the multiset $\{RR_A^{D^G}(\gamma; w) : \gamma \in A\}$, the normalisation condition has to be maintained, which would be impossible, for example, if G were strictly monotonic (see also Proposition 12 below).

7.2 Result used for proving monotonicity properties

The following technical result is used for showing monotonicity results for both kinds of family. It uses the notation $E_A(w)$ defined to be $Val_A(w) - D_A(w)$.

Proposition 9. Assume that $R_A^w(\alpha), R_{A'}^w(\alpha'), D_A(w), D_{A'}(w) > 0$.

- (i) If $D_A(w) - D_{A'}(w) \leq \frac{R_A^w(\alpha) - R_{A'}^w(\alpha')}{R_A^w(\alpha)/D_A(w)}$ (i.e., $D_A(w) - D_{A'}(w) \leq \frac{R_A^w(\alpha) - R_{A'}^w(\alpha')}{RR_A^D(\alpha; w)}$) then $\frac{R_A^w(\alpha)}{D_A(w)} \geq \frac{R_{A'}^w(\alpha')}{D_{A'}(w)}$, i.e., $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$.

- (ii) If $R_A^w(\alpha) \geq R_{A'}^w(\alpha')$ and $RR_A^D(\alpha; w) \leq 1$ then a sufficient condition for $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$ is $D_A(w) - D_{A'}(w) \leq R_A^w(\alpha) - R_{A'}^w(\alpha')$, i.e., $E_{A'}(w) - E_A(w) \leq w(\alpha') - w(\alpha)$.

Proof sketch: For any $a, b, a', b' > 0$, if $b - b' \leq \frac{a - a'}{a/b}$ then $\frac{b - b'}{b} \leq \frac{a - a'}{a}$, which implies $1 - \frac{b - b'}{b} \geq 1 - \frac{a - a'}{a}$, i.e., $\frac{b'}{b} \geq \frac{a'}{a}$, and hence, $\frac{a}{b} \geq \frac{a'}{b'}$. We apply this with $a = R_A^w(\alpha)$, $b = D_A(w)$, $a' = R_{A'}^w(\alpha')$ and $b' = D_{A'}(w)$ to prove part (i). If we also assume $R_A^w(\alpha) \geq R_{A'}^w(\alpha')$ and $RR_A^D(\alpha; w) \leq 1$ and $D_A(w) - D_{A'}(w) \leq R_A^w(\alpha) - R_{A'}^w(\alpha')$, then $D_A(w) - D_{A'}(w) \leq \frac{R_A^w(\alpha) - R_{A'}^w(\alpha')}{RR_A^D(\alpha; w)}$, so we can apply part (i) to show part (ii). \square

8 Monotonicity results for G -aggregation

The discussion at the end of Section 7.1 indicates that we can only expect (full) (III)-monotonicity in exceptional cases; this is formalised with Proposition 12 below, showing that it only holds when G is a multiple of max. However, restricted (III)-monotonicity does hold, and there is a simple argument showing (I)- and (II)-monotonicity, based on Equation 5 below.

For $\gamma \in A$ let us abbreviate $R_A^w(\gamma)/R_A^w(\alpha)$ as $R_A^w(\gamma/\alpha)$. Consider the case when $R_A^w(\alpha) > 0$. Then, since $1/R_A^w(\alpha) > 0$, we can apply positive homogeneity of G to obtain $G(\{R_A^w(\gamma/\alpha) : \gamma \in A\}) = \frac{1}{R_A^w(\alpha)} G(\{R_A^w(\gamma) : \gamma \in A\}) = \frac{1}{RR_A^{D^G}(\alpha; w)}$, and thus, we have, when $R_A^w(\alpha) > 0$:

$$\frac{1}{RR_A^{D^G}(\alpha; w)} = G(\{R_A^w(\gamma/\alpha) : \gamma \in A\}). \quad (5)$$

Theorem 10. For any $G \in \mathcal{G}_n$, RR^{D^G} satisfies (I)-monotonicity, (II)-monotonicity and restricted (III)-monotonicity.

For any type-(I) (A, α) -modification of A to A' , only α changes, with $R_{A'}^w(\alpha') \leq R_A^w(\alpha)$. Hence, for all $\gamma \in A$, $R_{A'}^w(\gamma'/\alpha') \geq R_A^w(\gamma/\alpha)$, and so by monotonicity of G and using Equation 5, $RR_{A'}^{D^G}(\alpha'; w) \leq RR_A^{D^G}(\alpha; w)$, so RR^{D^G} satisfies (I)-monotonicity. The argument for (II)-monotonicity is very similar. Because α does not change, and only non-optimal elements change we again have for all $\gamma \in A$, $R_{A'}^w(\gamma'/\alpha') \geq R_A^w(\gamma/\alpha)$, and thus, using the same argument, $RR_{A'}^{D^G}(\alpha'; w) \leq RR_A^{D^G}(\alpha; w)$, showing (II)-monotonicity of RR^{D^G} .

The argument for restricted-(III)-monotonicity is a little more complex. We have $D_A^G(w) = G(\{Val_A(w) - w(\gamma) : \gamma \in A\})$, which, by the translation property of G is equal to $Val_A(w) + G(\{-w(\gamma) : \gamma \in A\})$. Thus, $E_A(w) = Val_A(w) - D_A(w) = -G(\{-w(\gamma) : \gamma \in A\})$, and similarly, $E_{A'}(w) = Val_{A'}(w) - D_{A'}(w) = -G(\{-w(\gamma') : \gamma \in A\})$. A type-(III) modification only reduces $w(\gamma)$ for optimal elements γ , and so monotonicity of G implies $E_{A'}(w) \leq E_A(w)$. Since we also have $w(\alpha') = w(\alpha)$, under the assumption $RR_A^D(\alpha; w) \leq 1$, we can use Proposition 9(ii) to show $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$, thus proving that RR^{D^G} satisfies restricted-(III)-monotonicity.

If we consider the special case when G is max and so D^G equals D^1 (see Section 3.4), then we always have $RR_A^{D^G}(\alpha; w) \leq 1$ and so restricted (III)-monotonicity is equivalent with (III)-monotonicity, and thus we have the following corollary of Theorem 10.

Corollary 11. RR^{D^1} satisfies (I)-monotonicity, (II)-monotonicity and (III)-monotonicity.

In fact, for continuous G , (full) (III)-monotonicity only holds for D^G if $G = \max$ or a multiple of \max .

Proposition 12. *Suppose that $G \in \mathcal{G}_n$ is continuous and that D^G satisfies (III)-monotonicity. Then there exists $r > 0$ such that for all n -multisets S with minimum element of zero, $G(S) = r \max(S)$.*

9 Monotonicity Results for Pointwise Aggregation

We first give a technical result that we will use for (I)- and (II)-monotonicity. This follows easily from Proposition 9(i) because a summation over all i in the hypothesis gives $w \cdot \theta_{A'}^H - w \cdot \theta_A^H \leq \frac{w \cdot \alpha' - w \cdot \alpha}{RR_A^D(\alpha; w)}$, and because $Val_{A'}(w) = Val_A(w)$, $w \cdot \theta_{A'}^H - w \cdot \theta_A^H = D_A(w) - D_{A'}(w)$ and $w \cdot \alpha' - w \cdot \alpha = R_A^w(\alpha) - R_{A'}^w(\alpha')$.

Proposition 13. *Consider any $H \in \mathcal{H}_n$, $w \in \mathbb{R}_{\geq}^p$, finite $A \subseteq \mathbb{R}^p$, $\alpha \in A$, and any 1-1 mapping A to $A' \subseteq \mathbb{R}^p$ such that $w(\alpha') \geq w(\alpha)$, and $Val_{A'}(w) = Val_A(w)$. Assume also that $Val_A(w) > w(\alpha') \geq w(\alpha)$.*

If for all $i = 1, \dots, p$, $\theta_{A'}^H(i) - \theta_A^H(i) \leq \frac{\alpha'(i) - \alpha(i)}{RR_A^D(\alpha; w)}$ then $RR_{A'}^{D^{(H)}}(\alpha'; w) \leq RR_A^{D^{(H)}}(\alpha; w)$.

The following lemmas relate to the three types of monotonicity.

Lemma 1. *If $H \in \mathcal{H}_n$ and H is t -bounded (see Section 4.3) then $RR^{D^{(H)}}$ satisfies $(\leq 1/t)$ -(I)-monotonicity.*

Proof sketch: For each $i = 1, \dots, p$, since H is t -bounded, we have, for a type-(I) (A, α) -modification, $\theta_{A'}^H(i) - \theta_A^H(i) \leq t(\alpha'(i) - \alpha(i))$. Thus, if $RR_A^D(\alpha; w) \leq 1/t$ Proposition 13 implies that $RR_{A'}^{D^{(H)}}(\alpha'; w) \leq RR_A^{D^{(H)}}(\alpha; w)$. \square

Lemma 2. *For any $H \in \mathcal{H}_n$, $RR^{D^{(H)}}$ satisfies (II)-monotonicity.*

Proof sketch: Consider any type-(II) (A, α) -modification of A to A' . Monotonicity of H implies that for each $i = 1 \dots, p$, $\theta_{A'}^H(i) \leq \theta_A^H(i)$. Since $\alpha' = \alpha$, Proposition 13 implies $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$ showing that $RR^{D^{(H)}}$ satisfies (II)-monotonicity. \square

Lemma 3. *For any $H \in \mathcal{H}_n$, $RR^{D^{(H)}}$ satisfies (≤ 1) -(III)-monotonicity.*

Proof: Assume that $H \in \mathcal{H}_n$ and abbreviate $D^{(H)}$ to D . Consider any type-(III) (A, α) -modification given w to produce A' from A . Then $w(\alpha') = w(\alpha) < Val_A(w)$ and $Val_{A'}(w) \leq Val_A(w)$ and so $R_{A'}^w(\alpha') \geq R_A^w(\alpha)$, and any elements that are changed are Pareto-reduced, so, by the monotonicity property of H we have $\theta_{A'}^H(i) \leq \theta_A^H(i)$ for each $i = 1, \dots, p$, so $E_{A'}(w) = w \cdot \theta_{A'}^H \leq w \cdot \theta_A^H = E_A(w)$. Therefore, Proposition 9 (ii) implies that $RR_{A'}^D(\alpha'; w) \leq RR_A^D(\alpha; w)$ whenever $RR_A^D(\alpha; w) \leq 1$. \square

Putting the three lemmas together we obtain:

Theorem 14. *For any $H \in \mathcal{H}_n$, $RR^{D^{(H)}}$ satisfies (II)-monotonicity, restricted-(III)-monotonicity, and, if H is t -bounded then $RR^{D^{(H)}}$ satisfies $(\leq 1/t)$ -(I)-monotonicity.*

Consider the special case when H is \min , so that $D^{(H)} = D^2$ (see Section 3.4). Viewing \min as an OWA $F^{\bar{t}}$ we have $t_{max} = 1$, and so \min is 1-bounded (see the end of Section 4). The values of RR^{D^2} are always ≤ 1 , and so (≤ 1) -(III)-monotonicity implies (III)-monotonicity, and similarly, for (I)-monotonicity, thus showing the following corollary of Theorem 14.

Corollary 15. *RR^{D^2} satisfies (I)-monotonicity, (II)-monotonicity and (III)-monotonicity.*

10 Re-normalised Standard Regret Measure D^3

Of the four variations of relative regret from [21], see Section 3.4, only D^3 remains to be considered, since D^0 and D^1 are members of the family considered in Section 5, and D^2 (and D^0) are members of the family in Section 6. Like the other three, $[max]$ relative regret based on D^3 satisfies translation- and scale-invariance properties, see Theorems 11 and 14 of [21]. We briefly consider monotonicity properties of D^3 .

For $A \in \mathcal{A}_n$, $\alpha \in A$ and $w \in \mathbb{R}^p$ we have $D_A^3(w) = w \cdot (\bar{A} - \underline{A})$, and

$$RR_A^{D^3}(\alpha; w) = \frac{R_A^w(\alpha)}{w \cdot (\bar{A} - \underline{A})} = \frac{Val_A(w) - w \cdot \alpha}{w \cdot \bar{A} - w \cdot \underline{A}}. \quad (6)$$

Since $Val_A(w) \leq w \cdot \bar{A}$ and $w \cdot \alpha \geq w \cdot \underline{A}$, we have $RR_A^{D^3}(\alpha; w) \leq 1$.

It can be seen that RR^{D^3} satisfies (I)-monotonicity. If α' is optimal in A' then we have, $0 = RR_{A'}^{D^3}(\alpha'; w) \leq RR_A^{D^3}(\alpha; w)$. Otherwise, $Val_{A'}(w) = Val_A(w)$ and so the numerator is reduced by $w \cdot (\alpha' - \alpha)$, and the denominator (if it decreases at all) decreases by no more than $w \cdot (\alpha - \alpha')$. Since $RR_A^{D^3}(\alpha; w) \leq 1$, this implies $RR_{A'}^{D^3}(\alpha'; w) \leq RR_A^{D^3}(\alpha; w)$.

It is easy to construct counter-examples to (II)-monotonicity. Choose any $w \in \mathbb{R}_{\geq}^p$ that has no zero components; choose any A and α such that $\bar{A} \notin A$ and $\alpha(i) < \bar{A}(i)$ for $i = 1 \dots, p$ and $\alpha \notin O_w(A)$. Pick any $\beta \in O_w(A)$. Since $\beta \neq \bar{A}$ there exists j such that $\beta(j) < \bar{A}(j)$; and for each $\gamma \in A$ such that $\gamma(j) = \bar{A}(j)$, define $\gamma'(j)$ to be $\frac{1}{2}(\bar{A}(j) - \underline{A}(j))$, and otherwise let $\gamma'(i) = \gamma(i)$ and leave all other elements of A unchanged. This defines a type-(II) (A, α) -modification given w , since α and $Val_A(w)$ are unchanged, and at least one element is reduced. In addition $w \cdot \bar{A}$ is reduced and \underline{A} is unchanged, so by Equation 6, $RR_A^{D^3}(\alpha; w)$ is increased, showing that (II)-monotonicity does not hold for RR^{D^3} . Specifically, $Val_{A'}(w) = Val_A(w) = w \cdot \beta$, and $w \cdot \alpha' = w \cdot \alpha$, and $w \cdot \underline{A}' = w \cdot \underline{A}$ and $w \cdot \bar{A}' < w \cdot \bar{A}$. Hence we have $RR_{A'}^{D^3}(\alpha'; w) > RR_A^{D^3}(\alpha; w)$, giving a counter-example for (II)-monotonicity.

In fact, we can choose α such that $w \cdot \alpha$ is arbitrarily close to $w \cdot \beta = Val_A(w)$ so the same form of example shows that for every $t > 0$, $(\leq t)$ -(II)-monotonicity does not hold for RR^{D^3} .

One can also construct an example to show that (III)-monotonicity does not hold for RR^{D^3} .

11 Discussion

The paper defines two general families of translation- and scale-invariant measures of regret, using aggregation functions in two different ways. Both families can be characterised by normalisation conditions that they satisfy. We analysed their monotonicity properties, and we establish positive results for all instances of both families. We suggest that monotonicity based on just reducing the optimal alternative is not necessarily to be always expected in this context, although an extreme member of each of the two families does satisfy it, i.e., D^1 and D^2 -relative regret.

For simplicity we focused just on linear utility functions. However, at least for the first family, based on aggregated regret, many of the results generalise easily to other utility functions.

Acknowledgements

This publication has emanated from research conducted with the financial support of Research Ireland under Grant number 12/RC/2289-P2 which is co-funded under the European Regional Development Fund. Thanks also to the reviewers for their comments which helped to improve the paper.

References

- [1] I. Averbakh. Computing and minimizing the relative regret in combinatorial optimization with interval data. *Discrete Optimization*, 2(4): 273–287, 2005.
- [2] N. Benabbou and P. Perny. Incremental weight elicitation for multi-objective state space search. In *Proc. AAAI 2015*, pages 1093–1099, 2015.
- [3] N. Benabbou and P. Perny. Adaptive elicitation of preferences under uncertainty in sequential decision making problems. In *Proc. IJCAI 2017*, pages 4566–4572, 2017. doi: 10.24963/ijcai.2017/637.
- [4] C. Boutilier, R. Patrascu, P. Poupart, and D. Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 2006.
- [5] D. Braziunas and C. 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*, pages 65–72. ACM, 2007. doi: 10.1145/1297231.1297243.
- [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] V. S. Iyengar, J. Lee, and M. Campbell. Evaluating multiple attribute items using queries. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 144–153, 2001.
- [9] 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. doi: 10.1016/j.ejor.2017.01.003.
- [10] R. Keeney and H. Raiffa. *Decisions with multiple objectives: preferences and value tradeoffs*. Cambridge University Press, 1993.
- [11] P. Kouvelis and G. Yu. *Robust discrete optimization and its applications*, volume 14. Springer Science & Business Media, 2013.
- [12] T. Lu and C. 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*, pages 497–512, 2013.
- [14] H. Raiffa. *Decision analysis*. Addison-Wesley, 1968.
- [15] A. Salo and R. P. Hämäläinen. Preference programming – multicriteria weighting models under incomplete information. In *Handbook of Multicriteria Analysis*, pages 167–187. Springer Berlin Heidelberg, 2010.
- [16] L. J. Savage. The theory of statistical decision. *Journal of the American Statistical association*, 46(253):55–67, 1951.
- [17] P. Viappiani and C. Boutilier. Regret-based optimal recommendation sets in conversational recommender systems. In *Proc. RecSys-2009*, pages 101–108. ACM, 2009. doi: 10.1145/1639714.1639732.
- [18] P. Viappiani and C. Boutilier. On the equivalence of optimal recommendation sets and myopically optimal query sets. *Artificial Intelligence*, page 103328, 2020.
- [19] T. Wang and C. Boutilier. Incremental utility elicitation with the minimax regret decision criterion. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 3, pages 309–316, 2003.
- [20] N. Wilson. On the variation of max regret with respect to the scaling of the objectives. In *ECAI 2023 – 26th European Conference on Artificial Intelligence*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, pages 2639–2646, 2023.
- [21] N. Wilson. Scale-invariant variations of max regret. In *ECAI 2024 – 27th European Conference on Artificial Intelligence*, 2024.
- [22] R. R. Yager. On ordered weighted averaging aggregation operators in multicriteria decisionmaking. *IEEE Trans. Syst. Man Cybern.*, 18(1): 183–190, 1988.