
Incremental Preference Elicitation for Decision Making Under Risk with the Rank-Dependent Utility Model

Patrice Perny

Sorbonne Universités
UPMC Univ Paris 06
CNRS, LIP6 UMR 7606
Paris, France

Paolo Viappiani

Sorbonne Universités
UPMC Univ Paris 06
CNRS, LIP6 UMR 7606
Paris, France

Abdellah Boukhatem

Sorbonne Universités
UPMC Univ Paris 06
CNRS, LIP6 UMR 7606
Paris, France

Abstract

This work concerns decision making under risk with the rank-dependent utility model (RDU), a generalization of expected utility providing enhanced descriptive possibilities. We introduce a new incremental decision procedure, involving monotone regression spline functions to model both components of RDU, namely the probability weighting function and the utility function. First, assuming the utility function is known, we propose an elicitation procedure that incrementally collects preference information in order to progressively specify the probability weighting function until the optimal choice can be identified. Then, we present two elicitation procedures for the construction of a utility function as a monotone spline. Finally, numerical tests are provided to show the practical efficiency of the proposed methods.

1 INTRODUCTION

Uncertainty is pervasive in human activities and represents an important source of complexity in individual and collective decision making. As soon as intelligent systems are used for supporting human decision making or simulating realistic human decision behaviors, preference modeling and preference elicitation becomes particularly important. In decision making under uncertainty and risk, preference models are used to represent uncertain outcomes and to provide normative decision rules for choice problems, but also to describe, predict or simulate observed human behaviors. Given a mathematical model used to compare the alternatives of a choice problem, preference elicitation consists in fitting the parameters of the model to a specific Decision Maker (DM), to capture her attitude towards risk. Then the model can be used to support her choices in complex situations involving a large number of alternatives or to integrate her value system into an autonomous decision

agent. Considering this context, our paper aims at providing new tools for interactive decision support under risk.

Decision under risk is a standard formal framework for handling uncertainty in decision making, characterized by a probabilistic representation of uncertainty. In this framework, risky prospects are represented by probability distributions with a finite support, namely lotteries. In the seminal work of Bernoulli (1738; refer to [1954] for an English translation) and in the theory of von Neumann and Morgenstern (vNM) [1947], the values of lotteries are measured in terms of *expected utility* (EU). This well-known decision criterion is linear in probabilities and characterized by a utility function encoding the subjective value of any possible consequence for the DM; EU is used to compare lotteries and choice problems are solved by EU maximization. This choice model has been axiomatically justified in the context of risk by vNM [1947], but also in the more general context of uncertainty introduced by Savage [1954], where probabilities are not assumed to exist a priori.

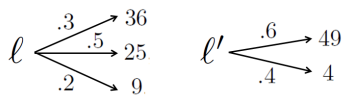
From EU to RDU. Despite these axiomatic results and the intuitive appeal of EU theory, the model suffers from well known limitations, from a descriptive viewpoint. There exist situations where individuals may exhibit behaviors that are not consistent with EU theory, as illustrated by the so-called Allais's paradox. For example, people often prefer \$3000 with certainty (choice A) to \$4000 with probability .8 (choice B), but they prefer \$4000 with probability .2 (choice D) to \$3000 with probability .25 (choice C). Remarking that $C = .25A + .75O$ and $D = .25B + .75O$ where O is the choice to win nothing with probability 1, it turns out that such preferences violate the vNM independence axiom. This example, used by Kahneman and Tversky [1979] in their experiments, is frequently observed and known as the *certainty effect*: a reduction in probability of winning a reward creates a larger (negative) psychological effect when it is done from certainty than from uncertainty. Similar observations have been made in actual decision contexts, for instance in route-choice problems where evidence was found of violation of EU theory [Avineri and Prashker, 2004].

These experiments and others in the same spirit suggest that probabilities are distorted in the decision making process. More precisely, small probabilities (greater than 0) tend to be overestimated while large ones (smaller than 1) underestimated. This has motivated the introduction of a probabilistic transformation model to generalize EU, leading to evaluate the lottery $(x_1, p_1; \dots, x_n; p_n)$ that yields outcome x_i with probability p_i by $\sum_i w(p_i)u(x_i)$. This generalization of EU dates back to [Edwards, 1955] and appears also in prospect theory [Kahneman and Tversky, 1979]. However, it can be easily shown that it violates the principle of stochastic dominance and could lead to select dominated alternatives; this is often considered as a serious weakness from a normative viewpoint. To overcome this problem, a solution proposed by Quiggin [1982] and Yaari [1987] consists in applying the probability transformation to the decumulative probability distribution function, and not to the probabilities of individual outcomes. This solution leads to the *rank-dependent utility* model (RDU) that consists in evaluating the lottery $\ell = (x_1, p_1; \dots, x_n; p_n)$ such that $x_1 \geq \dots \geq x_n \geq 0^1$ by:

$$V(\ell) = \sum_{i=1}^n w\left(\sum_{k=1}^i p_k\right) [u(x_i) - u(x_{i+1})]$$

$$= w(p_1)u(x_1) + \sum_{i=2}^n \left[w\left(\sum_{k=1}^i p_k\right) - w\left(\sum_{k=1}^{i-1} p_k\right) \right] u(x_i)$$

where $x_{n+1} = 0$ and $w : [0, 1] \rightarrow [0, 1]$ is a non-decreasing function such that $w(0) = 0$ and $w(1) = 1$. Note that the coefficient of $u(x_i)$ in RDU depends on the cumulative probability of the outcomes greater or equal to x_i (since outcomes are indexed to be sorted by decreasing order); this shows that this coefficient depends on the rank of x_i in the set of possible outcomes. To illustrate the use of RDU, we compare the two following lotteries ℓ and ℓ' with $w(p) = p^2$ and $u(x) = \sqrt{x}$:



$V(\ell) = (.3)^2(6 - 5) + (.8)^2(5 - 3) + 1^2 \cdot 3 = 4.37$ and $V(\ell') = (.6)^2(7 - 2) + (1)^2 \cdot 2 = 3.8$, hence ℓ is preferred to ℓ' . To give another example, let us remark that, for any lottery of type $\ell = (x^+, p; x^-, 1 - p)$, $x^+ > x^-$, we have:

$$V(\ell) = w(p)u(x^+) + (1 - w(p))u(x^-)$$

Function w is referred to as the *probability weighting function*. It is worth noting that, when $w(p) = p$, RDU boils down to EU as we can see from the second formulation. On the other hand, RDU is also known to be an instance of Choquet Expected Utility [Schmeidler, 1989]. More pre-

¹As suggested by Gilboa [2008], it is more natural to think of the object of choice in Quiggin and Yaari's models as lotteries over final wealth (with positive outcomes) rather as prospects of gains and losses, because they did not emphasize gain-loss asymmetry in their theory.

cisely, it turns out to be the natural instance for decision making under risk due its compatibility with stochastic dominance [Wakker, 1990]. For more details about RDU theory see [Quiggin, 2012, Diecidue and Wakker, 2001].

Let us mention also some well known variants of RDU theory. First, Yaari [1987] introduced and axiomatized a dual model to expected utility where transformations are applied to probabilities rather than to outcomes. This is a special case of RDU obtained when the utility function is linear in the outcomes. Moreover, the idea of rank-dependent weightings was also incorporated by Kahneman and Tversky [1992] into prospect theory, leading to *cumulative prospect theory*. This theory extends the RDU theory to model different risk attitudes towards gains depending on whether above or below a reference level.

In this paper, we focus on the RDU model, and we address the problem of identifying, within a given (possibly large) set of lotteries, the preferred solution for a given DM consistent with RDU theory. This solution could be used for predicting a choice of the DM or even to make a recommendation. This requires to elicit, at least partially, the probability weighting function and the utility function. Note that the determination of RDU-optimal solutions has motivated various contributions in the past, e.g. [Nielsen and Jaffray, 2006, Jeantet *et al.*, 2012].

Incremental decision making. A first approach to elicitation, standard in mathematical economics, aims to obtain a full description of preferences by the decision model, assuming that DM's preferences are observable on all possible pairs of alternatives. The elicitation process is justified by the axioms of the underlying theory and the components of the models are revealed point by point using systematic sequences of queries to obtain a precise specification of the model. To simplify the process, a frequent option used in economics (but also in artificial intelligence) consists in postulating a parametric form for each component of the decision model; queries are then used to fit the parameters. Refer to [Wakker and Deneffe, 1996, Gonzalez and Wu, 1999, Abdellaoui, 2000] for full elicitation procedures proposed in economics for RDU, and to [Fürnkranz and Hüllermeier, 2003, Torra, 2010, Tehrani *et al.*, 2012] for model-based preference learning.

Recently a number of works have tackled the elicitation of the components of the decision model in an adaptive way [Wang and Boutilier, 2003, Boutilier *et al.*, 2006, Braziunas and Boutilier, 2007, Benabbou *et al.*, 2014]; the aims is that of focusing on learning the "important" part of the utility, allowing to recommend near-optimal decisions with only partial information about the decision model. These incremental approaches consider all possible instances of preference model parameters consistent with the currently known information about the user; the user's responses allow to infer constraints on these parameters.

In this work, we consider the problem of incrementally eliciting the components of the RDU model by interactively asking queries to the DM; this aims to support the identification or approximation of an RDU-optimal choice. We first address the problem of eliciting the weighting function incrementally (Section 2) assuming that the utility is known. Afterwards, we will present utility elicitation procedures adapted to the rank-dependent utility model (Section 3). Finally, we provide numerical tests (Section 4).

2 ELICITATION OF THE PROBABILITY WEIGHTING FUNCTION

Decision makers exhibit various decision behaviors when they are confronted to choices involving risky prospects. Within RDU theory, this diversity of attitudes towards risk can be modeled by controlling the definition of the two components of the model, namely the utility function and the probability weighting function. Since these two components are strongly interlaced in RDU, their joint elicitation is quite challenging. Fortunately, these two components concern two separate and well identified risk-components, on the one hand marginal utilities induced by the shape of the cardinal utility function [Chateauneuf and Cohen, 1994], and on the other hand the probabilistic risk-attitude towards probabilistic mixtures induced by the shape of the probability weighting function [Wakker, 1994]. They can be observed separately using proper preference queries as shown in [Wakker and Deneffe, 1996, Abdellaoui, 2000]. In this section, the utility is assumed to be known (already elicited or known to be linear (Yaari’s model), and we focus on the elicitation of probability weights.

2.1 ON PROBABILITY WEIGHTING FUNCTIONS

The probability weighting function is necessarily non-decreasing but can take different forms depending on the attitude of the DM towards risk. For instance, in Yaari’s model, *weak risk aversion* which consists, for any lottery ℓ , in preferring $E(\ell)$ for sure to ℓ (where $E(\ell)$ is the expected value of ℓ) is equivalent to $w(p) \leq p$ for all $p \in [0, 1]$ [Chateauneuf and Cohen, 1994]. Moreover, in the RDU model, *strong risk aversion* which consists in preferring ℓ to ℓ' whenever ℓ stochastically dominates ℓ' at second order, is equivalent to using a convex weighting function and a concave utility [Hong *et al.*, 1987]. Finally, several experiments including those of Kahneman and Tversky [1979] lead to propose an inverse S-shaped function. For example, the use of the weighting function given in Figure 1 in Yaari’s model allows to explain the preferences observed on lotteries A, B, C, D presented in the introduction. We have indeed $V(A) = 3000 w(1)$, $V(B) = 4000 w(0.8)$, $V(C) = 3000 w(0.25)$ and $V(D) = 4000w(0.2)$. Since $w(0.2) \approx 0.26$, $w(0.25) \approx$

0.29, $w(0.8) \approx 0.60$ and $w(1) = 1$ we obtain $V(A) > V(B)$ while $V(C) < V(D)$.

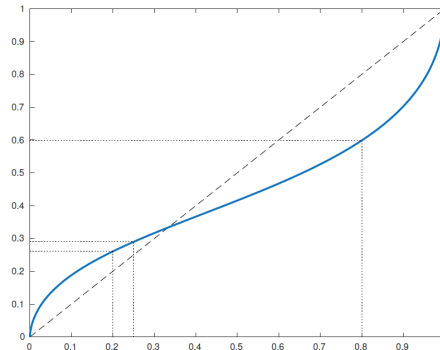


Figure 1: Inverse S-shaped function w .

The following parametric expressions for $w(p)$ have been proposed in the literature [Cavagnaro *et al.*, 2013], among several others:

$$\begin{aligned} \text{TK} : \quad w(p) &= \frac{p^r}{(p^r + (1-p)^r)^{1/r}} && \text{with } 0.28 < r \leq 1 \\ \text{Prelec} : \quad w(p) &= e^{-s(-\ln p)^r} && \text{with } 0 < r \leq 1; 0 < s \\ \text{LinLog} : \quad w(p) &= \frac{s p^r}{s p^r + (1-p)^r} && \text{with } 0 < r, s \end{aligned}$$

where TK denotes the parametric form proposed by Tversky and Kahneman (the w plotted in Figure 1 is of this type, with $r = 0.6$). Note that Prelec and LinLog make use of two parameters (r and s) while TK is based on a single parameter. Karmarkar [1978] proposes a parametric form that is a special case of LinLog with $s = 1$.

Choosing a priori a specific parametric form among these different options and others, with the idea of learning the parameters, may generate errors due to not taking into account the specificity of the DM that may be revealed during the elicitation process. Instead, we propose a new approach based on a much more flexible parametric construction based on the definition of w as a piecewise polynomial spline function. Moreover, we will adopt an incremental elicitation approach in order to reason with the possible functions w that are consistent with the current preference information of the DM. In the spirit of incremental decision making, by considering additional preferences (obtained by asking queries), we will be able to progressively narrow down our model of the DM until we are able to identify the preferred alternative.

2.2 INCREMENTAL ELICITATION OF w

We assume here that the utility function is known and we reason with an incompletely specified weighting function $w \in W$, W being the set of admissible probability weighting functions at a given step of the process. Initially, W can be all non-decreasing functions on the unit interval such that $w(0) = 0$ and $w(1) = 1$ or an approximation of

them representable by a given family of parametric functions. The utility function u being fixed, we will denote $V(\ell; w)$ the RDU value of lottery ℓ for weighting function w . Each time a preference statement of type “ ℓ^+ is at least as good as ℓ^- ” is obtained, this induces a new constraint $V(\ell^+; w) \geq V(\ell^-; w)$ that further restricts W .

In order to make a decision when probability weights are imprecisely known, we will use minimax regret which is a standard decision criterion for robust decision making under uncertainty [Savage, 1954]. It was more recently proposed by Boutilier et al. [2003, 2006] for decision-making under utility uncertainty. Moreover, minimax regret can be used as a driver for preference elicitation and incremental decision making.

Let \mathcal{L} be the set of lotteries representing the alternatives of the decision problem. If we knew the DM’s specific w , then the optimal choice would be any element of $\arg \max_{\ell \in \mathcal{L}} V(\ell; w)$. However, since w is not known precisely, we need to provide a recommendation based on the current information. We are interested in evaluating the loss that can result from choosing a decision ℓ instead of the optimal one. To this end, we use the following notions of regret. The pairwise maximum regret PMR of ℓ against ℓ' is the maximal value that the difference $V_u(\ell'; w) - V_u(\ell; w)$ can take for any admissible function of $w \in W$. The max regret MR of a choice ℓ is defined as the maximum pairwise regret when the adversary is chosen among all lotteries in \mathcal{L} . Finally, the minimax regret MMR is the minimum value of max regret. More formally:

$$\text{PMR}(\ell, \ell'; W) = \max_{w \in W} [V(\ell'; w) - V(\ell; w)] \quad (1)$$

$$\text{MR}(\ell; W) = \max_{\ell' \in \mathcal{L}} \text{PMR}(\ell, \ell'; W) \quad (2)$$

$$\text{MMR}(W) = \min_{\ell \in \mathcal{L}} \text{MR}(\ell; W) \quad (3)$$

Then we recommend one of the decisions associated with the minimum value of max regret, i.e. lottery $\ell^* \in \arg \min_{\ell \in \mathcal{L}} \text{MR}(\ell; W)$. This is a regret-optimal decision given the current preference information. By definition, recommending ℓ^* means to be robust with respect to the possible realizations of $w \in W$.

In addition of being a criterion for decision-making, minimax regret can be used to drive the elicitation of further preference information. We assume an interactive setting where the system (a decision-support agent) can ask additional queries in order to improve the quality of the recommendation. If the current minimax regret value is higher than a given positive threshold, we ask another query to the user. Different types of queries can be asked to the user; among the many possibilities, comparison queries, asking the user to state which choice is best among two presented to him, are particularly natural. It is however important to ask informative queries in order to quickly converge to a recommendation of high value. An effective heuristic to choose the next query is the *current solution* strategy

[Wang and Boutilier, 2003, Boutilier *et al.*, 2006]; it asks the user to compare the regret-optimal lottery ℓ^* with its adversarial challenger $\ell^a = \arg \max_{\ell \in \mathcal{L}} \text{PMR}(\ell^*, \ell)$. This will often reduce minimax regret: if the user states that ℓ^* is preferred to ℓ^a , in the next computation of minimax regret, the adversary will have to choose another lottery, leading to a reduction of regret (unless there were ties in the max regret computations). If, instead ℓ^a is preferred to ℓ^* , the regret-optimal choice in the next computation will necessarily be a choice other than ℓ^* , and, again, we will likely reduce regrets.

Assume \mathcal{P} to be a set of pairs (ℓ^+, ℓ^-) for which we know that the DM considers that ℓ^+ is at least as good as ℓ^- . Let $W = \{w : \forall (\ell^+, \ell^-) \in \mathcal{P}, V(\ell^+; w) \geq V(\ell^-; w)\}$, our goal is to compute $\text{PMR}(\ell, \ell', W)$. This corresponds to solving the following optimization problem:

$$\max_w [V(\ell'; w) - V(\ell; w)] \quad (4)$$

$$\text{s.t. } w(0) = 0, w(1) = 1 \quad (5)$$

$$w(p) \leq w(q), \quad \forall p, q \in [0, 1] : p \leq q \quad (6)$$

$$V(\ell^+; w) \geq V(\ell^-; w), \quad \forall (\ell^+, \ell^-) \in \mathcal{P} \quad (7)$$

Such optimization is not however directly feasible. First, the monotonicity constraint (6) on the unit interval represents implicitly an infinity of constraints. If we only impose monotonicity on probabilities involved in the lotteries of \mathcal{L} , it still represents a large number of constraints. Moreover, constraint (7) is quite difficult to handle. Even if we assume that w belongs to one of the families of parametric curves considered in the previous subsection, the resulting constraints will not be linear in the parameters. In the next subsection we will see how these problems can be overcome by defining w as a monotone spline function.

2.3 A MODEL BASED ON I-SPLINE FUNCTIONS

Spline functions are piecewise polynomials whose pieces connect with a high degree of smoothness. They are very useful in data interpolation and shape approximation due to their capacity to approximate complex shapes through curve fitting and interactive curve design while preserving an important property, missing in many other interpolation methods: they guarantee that smooth curves will be generated from smooth data [Beatty and Barsky, 1995]. The use of piecewise polynomials in non-linear regression extends the advantage of polynomials by providing greater flexibility, local effects of parameter changes and the possibility of imposing constraints on estimated functions [Ramsay, 1988].

One important feature of spline functions is that they can be generated by linear combinations of basis spline functions. A basis of splines particularly appealing for non-linear regression is the M-spline family [Ramsay, 1988]. M-splines of order k are functions $M_i, i = 1, \dots, m$ which are polynomials of degree $k - 1$. They can be used to approximate

any function defined on a given interval $[a, b]$ by a spline of the form $f = \sum_i \lambda_i M_i$. To define M_i precisely, we need to introduce a sequence of knots $t = \{t_1, t_2, \dots, t_l\}$ such that $a = t_1 = \dots = t_k$, $b = t_{l-k+1} = \dots = t_l$ and $\forall i, t_i \leq t_{i+1}$. The basis constructed from sequence t contains $m = |t| - k$ spline functions of order k denoted $M_i(x; k, t)$, $i = 1, \dots, m$, and defined for $k = 1$ by:

$$M_i(x; 1, t) = \begin{cases} \frac{1}{t_{i+1} - t_i} & \text{if } x \in (t_i, t_{i+1}) \\ 0 & \text{otherwise} \end{cases}$$

For $k > 1$ and $t_{i+k} > t_i$, M_i is defined recursively by:

$$M_i(x; k, t) = \frac{k[(x - t_i)M_i(x; k-1, t) + (t_{i+k} - x)M_{i+1}(x; k-1, t)]}{(k-1)(t_{i+k} - t_i)}$$

and otherwise $M_i(x; k, t) = 0$. In particular, we have $M_i(x; k, t) = 0$ whenever $k > 1$ and $t_{i+k} = t_i$.

Note that $M_i(x; k, t)$ is strictly positive in (t_i, t_{i+k}) and 0 elsewhere, with an integral equal to 1. Moreover, it is a polynomial of degree $k - 1$ so that, for any piecewise polynomial function of the form $f = \sum_i \lambda_i M_i$, adjacent polynomials have matching derivatives up to order $k - 2$. Hence, a good choice for k in practice is $k = 3$ because, in this case, we generate piecewise quadratic f functions with matching first derivatives while preserving a local influence of every component. Choosing a lower k would loose continuity of the first derivative and choosing a higher k would increase the range (t_i, t_{i+k}) of influence of every component M_i , which diminishes controllability of the model. As an example, the family of splines $M_i(x; 3, t)$, $i = 1, \dots, 5$ defined on the unit interval from subdivision $t = (0, 0, 0, \frac{1}{3}, \frac{2}{3}, 1, 1, 1)$ is given in the left part of Figure 2. Finer spline decompositions could be obtained using finer subdivisions. Moreover, it is possible to use non-equally-spaced knot sequences to have a finer control of the shape in some parts of interval $[a, b]$. However, practical tests show that using too many nodes is counterproductive to generate smooth and regular curves.

We may define the probability weighting function w as a weighted combination of M_i -spline functions with probably good fitting possibilities. However, this would not guarantee to obtain a non-decreasing weighting function w . To get rid of the monotonicity constraint (6), we need another basis specifically designed for the generation of non-decreasing spline functions. The solution is given by monotone regression splines that provide very nice descriptive possibilities, as demonstrated by Ramsey [1988]. Such spline functions are non-decreasing because they are generated by conical combinations of basis I-spline functions, i.e., non-decreasing functions defined as the integrals of the M-splines (which are positive). Formally, these functions denoted $I_i(x; k, t)$, $i = 1, \dots, m$ are defined by:

$$I_i(x; k, t) = \int_a^x M_i(y; k, t) dy$$

Let j be the index such that $t_j \leq x < t_{j+1}$, the value of an

I-spline is computed as follows:

$$I_i(x; k, t) = \begin{cases} 1 & \text{if } i < j - k + 1 \\ 0 & \text{if } i > j \\ \sum_{s=i}^j \frac{t_{s+k+1} - t_s}{k+1} M_s(x; k+1, t) & \text{otherwise} \end{cases}$$

As an illustration, the family of splines $I_i(x; 3, t)$, $i = 1, \dots, 5$ for subdivision $t = (0, 0, 0, 0, \frac{1}{3}, \frac{2}{3}, 1, 1, 1, 1)$ is given in the right part of Figure 2.

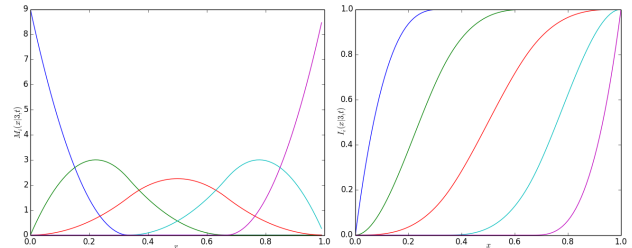


Figure 2: Basis functions $M_i(x; 3, t)$ and $I_i(x; 3, t)$

2.4 REGRET MINIMIZATION WITH I-SPLINES

The computation of PMR is the main building block to compute minimax regret. In order to compute MR and MMR, it is indeed sufficient to perform a quadratic number of PMR optimizations². We now focus the discussion on PMR computations and show that the optimization of regrets is drastically simplified when weighting function w is defined by a conical combination of I_i -splines:

$$w(p) = \sum_{j=1}^m \lambda_j I_j(p; k, t), \quad \lambda_j \geq 0, j = 1, \dots, m \quad (8)$$

Note that, k and t being fixed, function w is completely characterized by vector $\lambda = (\lambda_1, \dots, \lambda_m) \in \mathbb{R}_+^m$. So the elicitation of the entire function boils down to the elicitation of vector λ . In that case the RDU criterion reads, for any lottery $\ell = (x_1, p_1; \dots; x_n, p_n)$, as follows:

$$V(\ell; \lambda) = \sum_{i=1}^n \sum_{j=1}^m \lambda_j I_j \left(\sum_{k=1}^i p_k; k, t \right) [u(x_i) - u(x_{i+1})]$$

By permuting the two first summations in the above formulation, we obtain $V(\ell; \lambda) = \lambda^\top v(\ell, k, t)$ where $v(\ell, k, t) \in \mathbb{R}^m$ is a vector whose j^{th} component is equal to $\sum_{i=1}^n I_j(\sum_{k=1}^i p_k; k, t) [u(x_i) - u(x_{i+1})]$. Hence PMR values can be reformulated using vector λ and a set Λ of admissible weighting vectors replacing W . We obtain:

$$\begin{aligned} \text{PMR}(\ell, \ell'; \Lambda) &= \max_{\lambda \in \Lambda} [V(\ell'; \lambda) - V(\ell; \lambda)] \\ &= \max_{\lambda \in \Lambda} \lambda^\top [v(\ell', k, t) - v(\ell, k, t)] \end{aligned}$$

²In fact, a more efficient strategy consists in implementing an alpha-beta search procedure, in this way it is possible to prune several cases and compute a much lower number of PMRs [Braziunas, 2011].

Moreover, the constraint $V(\ell^+; w) \geq V(\ell^-; w)$ present in (7) takes now the form of a linear constraint $\lambda^\top [v(\ell^+, k, t) - v(\ell^-, k, t)] \geq 0$. Hence the set

$$\Lambda_{\mathcal{P}} = \{\lambda \in \mathbb{R}_+^m : \forall (\ell^+, \ell^-) \in \mathcal{P}, \lambda^\top [v(\ell^+, k, t) - v(\ell^-, k, t)] \geq 0\}$$

is a convex polyhedron. Therefore, under the assumption that I-spline functions are convenient to describe the probability weighting function, the computation of PMR for a given pair of lotteries (ℓ, ℓ') , initially introduced as a difficult optimization problem, see Equations (4-7), can now easily be achieved by solving the following linear program:

$$\begin{aligned} & \max \lambda^\top [v(\ell', k, t) - v(\ell, k, t)] \\ \text{s.t. } & \lambda^\top [v(\ell^+, k, t) - v(\ell^-, k, t)] \geq 0, \forall (\ell^+, \ell^-) \in \mathcal{P} \\ & \lambda_i \geq 0, i = 1, \dots, m. \end{aligned}$$

2.5 THE ELICITATION ALGORITHM

Algorithm 1 details the steps involved in our regret-based approach for the elicitation of w . Given a dataset of decisions (lotteries) \mathcal{L} (for each decision we are given a specification of the numerical outcomes and their associated probabilities) and a set of initial preferences \mathcal{P} (that could be empty), we compute the initial minimax regret value v , the max-regret-optimal decision ℓ^* and the adversarial decision ℓ^a . We then start asking queries.

Function `Query` simulates the question/answer protocol. It takes two lotteries as input and returns the preferred lottery among the two. After each user's response we recompute the minimax regret value, the regret-optimal decision ℓ^* , and its challenger ℓ^a . We loop until the regret drops below a positive threshold ε . Queries are asked according to the *current solution strategy* (analogue to the strategy used in [Boutilier *et al.*, 2006]), that requires the user to compare ℓ^* with ℓ^a . The least preferred among ℓ^* and ℓ^a is removed from \mathcal{L} as it is dominated by the other. This guarantees a strict reduction of $|\mathcal{L}|$ at every step, ensuring a linear convergence in the number of lotteries. The practical efficiency of this algorithm is illustrated in Section 4.

Note that when the lotteries in \mathcal{L} involve a large number of branches, we cannot expect a DM to be able to compare them with confidence. In such cases, we recommend to work with two sets of lotteries: the actual set \mathcal{L} of lotteries (in which the preferred solution must be found), on which pairwise regrets are computed, and a second set of simpler lotteries used only for preference queries. For the latter, we may use lotteries of type $\ell_p = (M, p; 0, 1 - p)$ (where M is the top outcome) and ask the DM to compare ℓ_p , for some probability p , to some certain consequence x in $[0, M]$. Under the assumptions $u(0) = 0$ and $u(M) = 1$, we derive from the response either $w(p) \geq u(x)$ or the reverse inequality, thus reducing uncertainty on function w and, in most cases, the MMR on actual lotteries.

Algorithm 1: Regret-based elicitation of w

Input: $\mathcal{L}, \mathcal{P}, \varepsilon$

Output: ℓ^*

begin

$\ell^* \leftarrow \arg \min_{\ell \in \mathcal{L}} \text{MR}(\ell; \Lambda_{\mathcal{P}});$

$\ell^a \leftarrow \arg \max_{\ell \in \mathcal{L}} \text{PMR}(\ell^*, \ell; \Lambda_{\mathcal{P}});$

$v \leftarrow \text{MMR}(\Lambda_{\mathcal{P}});$

while $v > \varepsilon$ **do**

$p \leftarrow \text{Query}(\ell^*, \ell^a);$

$\mathcal{L} \leftarrow \mathcal{L} \setminus \text{the least preferred in } \{\ell^*, \ell^a\};$

$\mathcal{P} \leftarrow \mathcal{P} \cup \{p\};$

$v \leftarrow \text{MMR}(\Lambda_{\mathcal{P}});$

$\ell^* \leftarrow \arg \min_{\ell \in \mathcal{L}} \text{MR}(\ell; \Lambda_{\mathcal{P}});$

$\ell^a \leftarrow \arg \max_{\ell \in \mathcal{L}} \text{PMR}(\ell^*, \ell; \Lambda_{\mathcal{P}});$

end

end

3 UTILITY ELICITATION IN RDU

In the previous section, we have discussed the elicitation of the weighting probability function in the RDU model, assuming the utility function was known. We discuss now the elicitation of function u when w is not known. The main difficulty to overcome lies in the fact that functions u and w are strongly interlaced in the computation of RDU values, due to products of type $w(\sum_{k=1}^i p_i)u(x_i)$. A given value $u(x_i)$ may impact differently on preferences, depending on the probability weighting function. Fortunately, there is a part of DM preferences that are not impacted by w and that can be used to elicit u . We present two illustrations of this idea in the two next subsections. The first one provides a precise spline function interpolating a sample of points of the utility curve constructed with the *gamble tradeoff method* proposed by Wakker and Deneffe [1996]. The second proposes an alternative approach based on simple queries aiming at obtaining the certainty equivalent of simple lotteries.

3.1 USING THE GAMBLE TRADEOFF METHOD

We first recall the principle of the gamble tradeoff method [Wakker and Deneffe, 1996] to construct points on the utility curve. To start with an example, let us consider the two following lotteries: $\ell = (c, p; b, 1 - p)$ and $\ell' = (d, p; a, 1 - p)$ for a probability $p \in (0, 1)$, where a, b, c, d are four distinct outcomes such that $a < b < c < d$. Then it can be easily checked from the definition of RDU that the DM is indifferent between ℓ and ℓ' (i.e., $V(\ell) = V(\ell')$) if and only if:

$$[u(b) - u(a)](1 - w(p)) = [u(d) - u(c)]w(p) \quad (9)$$

To construct some points of the utility function $u(x)$, x in $[0, M]$, we need two reference values α and β chosen in such a way that $M < \alpha < \beta$. Then, the gamble tradeoff method consists in generating the increasing sequence

defined by $r_0 = 0$ and $r_i = \text{Query}(\ell_i, \ell'_i)$ for $i > 1$ where:

$$\ell_i = (\alpha, p; r, 1 - p) \quad \ell'_i = (\beta, p; r_{i-1}, 1 - p)$$

and $\text{Query}(\ell_i, \ell'_i)$ is a function asking to the DM which value r makes the two lotteries indifferent and returns this value. The sequence is generated until $r \geq M$. By construction we have $V(r_i, 1 - p; \alpha, p) = V(r_{i-1}, 1 - p; \beta, p)$. Similarly we have $V(r_{i+1}, 1 - p; \alpha, p) = V(r_i, 1 - p; \beta, p)$. From these two equalities, we obtain, from equation (9):

$$\begin{aligned} [u(r_i) - u(r_{i-1})](1 - w(p)) &= [u(\beta) - u(\alpha)]w(p) \\ [u(r_{i+1}) - u(r_i)](1 - w(p)) &= [u(\beta) - u(\alpha)]w(p) \end{aligned}$$

Hence $u(r_{i+1}) - u(r_i) = u(r_i) - u(r_{i-1})$, therefore:

$$u(r_{i+1}) = 2u(r_i) - u(r_{i-1}) \quad (10)$$

Since u is defined up to a positive affine transformation, we can set, without loss of generality, $u(r_0) = 0$ and $u(r_1) = 1$. Hence, using (10) we obtain $u(r_i) = i$ for all $i \geq 0$ which means that the utility function should interpolate points (r_i, i) . Moreover, the density of values r_i within a given interval can be increased (resp. decreased) by changing parameters α and β to reduce or increase the difference $\beta - \alpha$.

Assume that q points (r_i, i) have been constructed with the gamble tradeoff method, these points can be interpolated using I-spline regression. So we define the utility function as a piecewise quadratic monotone spline $u(x) = \sum_k \lambda_k I_k(x, 3, t)$ for a uniform knot sequence t . Then, coefficients λ_k can be determined by minimizing the distance of the constructed points to the spline u . This is achieved by solving the following linear program II:

$$\begin{aligned} \min \quad & \sum_{i=1}^q e_i \\ \text{s.t.} \quad & \begin{cases} e_i \geq i - \sum_k \lambda_k I_k(r_i, 3, t), & i = 1, \dots, q \\ e_i \geq \sum_k \lambda_k I_k(r_i, 3, t) - i & i = 1, \dots, q \\ \lambda_k \geq 0, k = 1, \dots, q. \end{cases} \end{aligned}$$

This approach leads to a precise utility function; however, it is not incremental and does not provide an easy control on the number of points elicited on the utility curve. Moreover, preference queries involved in the process are somewhat more complex than in the certainty equivalent method. In the next subsection we present an incremental approach with simpler preference queries that gives control on which points are constructed.

3.2 USING THE CERTAINTY EQUIVALENT

An incremental utility elicitation procedure is proposed by Hines and Larsen [2010] with the aim of eliciting utilities while removing the effects of probability weighting from users answers. Using pairwise max regret minimization, they propose to elicit preferences in both cumulative prospect theory and expected utility theory by determining a specific probability value p^* that allows to ask *outcome*

queries where the effect of w cancels out. This idea could be used here but can be simplified in our context. Note that their definition of regret is based on the EU model and would not be the most appropriate *to derive robust recommendations with respect to RDU*. In order words they assume that EU is the right model to produce recommendations, but observe preferences biased by distorted probabilities as in cumulative prospect theory and RDU.

We propose here another incremental approach involving simpler preference queries based on the use of certainty equivalent. It consists first in identifying a probability p^* such that $w(p^*) = 1/2$. To this end, we only need to construct the two first elements r_1, r_2 following $r_0 = 0$ in the sequence defined in the previous subsection. Then we ask the DM for which probability p^* she is indifferent between lottery $\ell = (r_2, p^*; r_0, 1 - p^*)$ and lottery $\ell' = (r_1, 1)$ yielding r_1 with certainty. Since by construction, $u(r_0) = 0, u(r_1) = 1$ and $u(r_2) = 2, V(\ell) = u(r_0) + w(p^*)[u(r_2) - u(r_1)] = 2w(p^*)$ and $u(\ell') = 1$. This leads to $2w(p^*) = 1$ and therefore $w(p^*) = 1/2$.

Now the utility curve can be incrementally constructed on any interval $[0, M]$ using simple indifference queries based on the following principle: let r^- and r^+ be any two distinct values in $[0, M]$ such that $u(r^-)$ and $u(r^+)$ are known (initially we choose $r^- = 0$ and $r^+ = M$ and we set $u(0) = 0$ and $u(M) = 1$). Then a new point can be constructed between positions r^- and r^+ by asking the DM for the certainty equivalent of lottery $(r^+, p^*; r^-, 1 - p^*)$. If the answer is r (meaning that she is indifferent between the lottery and winning r for sure) we obtain $u(r) = u(r^-) + w(p^*)(u(r^+) - u(r^-)) = w(p^*)u(r^+) + (1 - w(p^*))u(r^-)$. Since $w(p^*) = 1/2$ we finally obtain $u(r) = (u(r^-) + u(r^+))/2$. When u is defined as a monotone spline, the new point will induce new constraints further restricting the set of possible utilities.

Denoting $(r_i, u(r_i)), i = 1, \dots, q$ the points already defined at iteration q , utility uncertainty within interval $[r_i, r_{i+1}]$ is bounded above by $\delta_i = u(r_{i+1}) - u(r_i)$ due to monotonicity of u . The next query is selected to obtain a new point in the interval $[r_k, r_{k+1}]$ which maximizes δ_k ; we ask for the certainty equivalent of lottery $(r_{k+1}, p^*; r_k, 1 - p^*)$. The process can be iterated to refine progressively the intervals until the desired number of points is reached or the maximal δ_i drops below a given threshold. Finally, a monotone regression spline is generated using the linear program II introduced in subsection 3.1 to approximate the elicited points.

4 NUMERICAL TESTS

Incremental elicitation of w . In order to model function w , we used I-spline defined from the non-uniform knot sequence $t' = (0, 0, 0, 0, .1, 0.9, 1, 1, 1, 1)$. The use of knots at positions .1 and .9 instead of 1/3 and 2/3 induces a small

left shift of $I_1(x, 3, t')$ and a small right shift of $I_5(x, 3, t')$ so as to include extremely concave or convex transformations in the family of splines that can be generated. This allows to model extremely risk-averse or risk-prone attitudes if necessary.

In the first series of experiments, we validate our choice of monotone splines for identifying the DM's probability weighting function w . Defining this function as $w(x) = \sum_{j=1}^5 \lambda_j I_j(x, 3, t')$, we used a linear program similar to Π (see Subsection 3.1) to fit parameters λ_j to different standard weighting functions, presenting various curves from extremely concave to extremely convex, including S-shaped and inverse S-shaped curves as well. The fitting possibilities of model w defined from I-splines are shown in Figure 3. In particular, we are able to approximate, with a single model, different kinds of distortion functions w (concave, convex, S-shaped) that usually require different parametric models (Prelec, TK, LinLog; see Section 2.1).

The instances where the approximation is good but not ideal are those with an extreme steepness (first instance, for very small value of p , and fourth instance, for very high values of p , in Figure 3). These situations (w extremely steep) will be rare in practice and correspond to shapes that could be even better described with a more specialized knot sequence. We conclude that the use of I-splines give us good approximations; splines allow us to learn w without committing to a specific parametric form.

In the second series of experiments, we observe the number of queries needed to determine the winning lotteries within sets of different sizes. The set \mathcal{L} is constructed by repeatedly generating at most k outcomes in $[0, 1000]$ and a probability distribution on these outcomes to build a new lottery (with k branches) at step i ; this lottery ℓ_i is inserted in \mathcal{L} only if no stochastic dominance holds in $\{\ell_i\} \cup \mathcal{L}$. The process is continued until the desired number of lotteries is obtained. We made experiments with lotteries with $k = 2, 3, 5, 10$ branches and a simulated DM. For small sets of alternatives (typically 100 lotteries) the procedure solves the problem after very few questions. To test the scalability of the approach to larger sets we have generated instances including 1000 or 2000 lotteries. Interactions with the DM are simulated by generating answers to preferences queries using RDU with an inverse S-shaped probability weighting function w_0 , unknown from the elicitation procedure, and a linear utility function. For each set of lotteries we observed the reduction of the minimax regret (expressed in percentage of the initial max regret before asking any query) as the number of queries increases. For comparison, we also observed the regret reduction obtained with an heuristic that randomly selects the new pair of lotteries to be compared. The resulting curves are plotted in Figure 4, respectively in green and red. At the same time, we observed the reduction of real regrets defined as the difference $V(\ell_0^*; w_0) - V(\ell^*; w_0)$ where ℓ_0^* is the optimal lottery

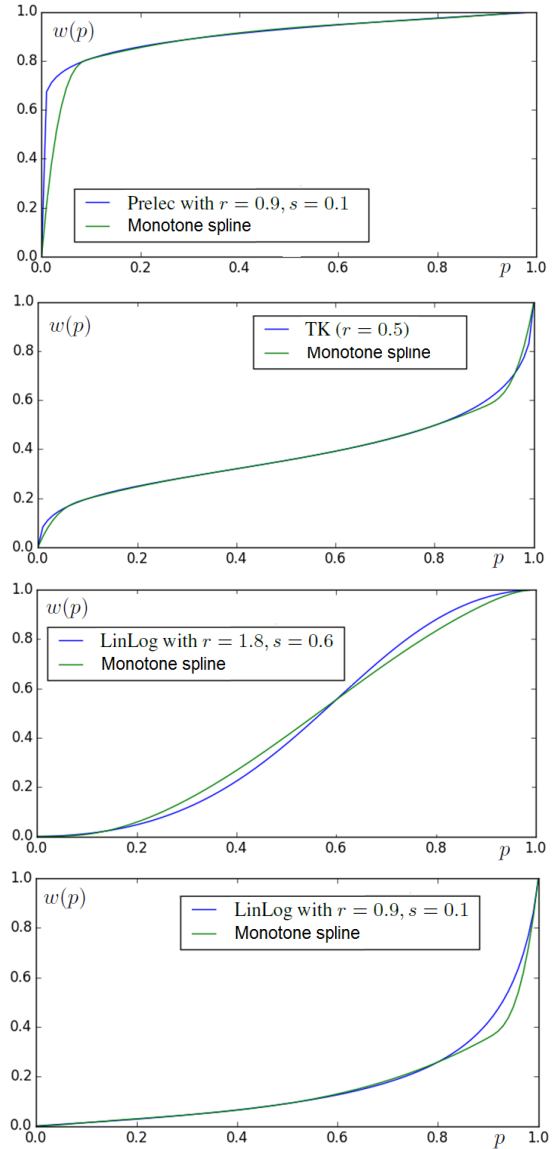


Figure 3: Monotone approximations of $w(p)$.

for the actual weighting function w_0 , and ℓ^* is the solution minimizing the current max regret $MR(\ell; W)$. This third curve is plotted in blue in Figure 4. This experiment has been repeated multiple times on randomly generated instances and the results are very similar.

Considering the green and blue curves showing the diminution of max regret and real regret in Figure 4, it appears that the winner can be determined exactly among 1000 lotteries in about 10 queries (in worst case 15 preference queries). This considerably improves the results obtained with the random strategy. Similar tests have been performed for 2000 alternatives. The curves have similar shape and show that less than 10 additional queries are necessary. In fact, we could stop earlier, when minimax regret is very small (even if not exactly zero). The elicitation procedure can indeed be seen as an anytime algorithm. It can be interrupted

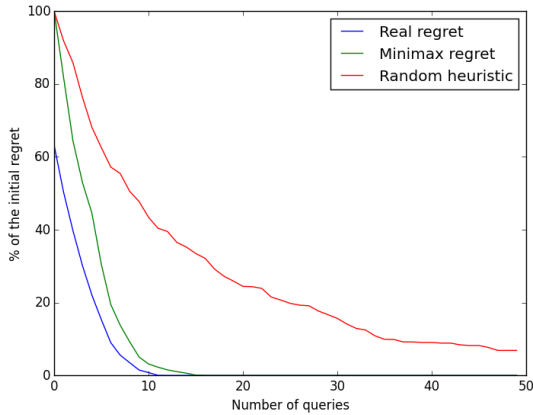


Figure 4: Regret reduction, $|\mathcal{L}| = 1000$, 100 runs.

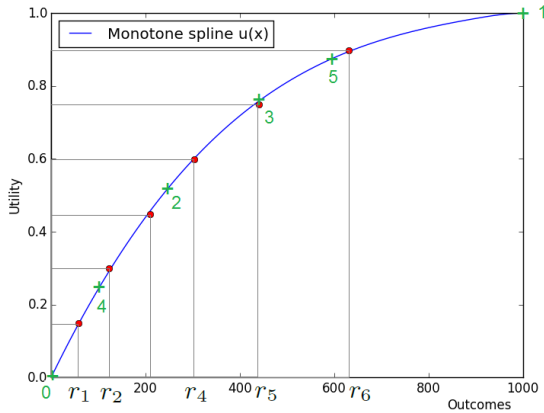


Figure 5: The spline approximating $u(x)$

before reaching a null regret, at any time of the process, and return the current minimax regret solution. We can also provide the associated max regret as a performance guarantee on the quality of the returned solution. Typically, as can be seen from the curves, stopping the elicitation when the minimax regret drops below a given percentage of the initial max regret (or below a given absolute threshold) will save a significant percentage of queries (at least 33% in our tests) without affecting significantly the quality of the recommendation. The efficiency of this approach is due to the fact that in most cases, the parameters of the decision model do not need to be precisely known to be able to determine a necessary winner, i.e., a lottery ℓ such that $V(\ell; w) \geq V(\ell'; w)$ for all $\ell' \in \mathcal{L}$ and all w compatible with the available preference information.

Utility elicitation. We present now an example of construction of a monotone spline on the $[0,1000]$ interval, to reveal and model the implicit utility function of a DM in RDU theory, from observed preferences involving lotteries with outcomes in this interval. The answers of the DM to preference queries have been simulated by a RDU model

with an inverse S-shaped probability weighting function $w(p)$ defined by the TK model introduced in Section 2 with parameter $r = 0.5$, and a concave utility defined by $u(x) = 1 - (1 - x/1600)^4$. We successively use the two methods proposed in Section 3 to elicit points and generate the regression spline. We first use the gamble tradeoff method as described in 3.1 with $\alpha = 1050$ and $\beta = 1600$ to obtain the sequence r_0, \dots, r_6 . We compute by linear programming the parameters λ_k of the monotone regression spline f that fits best points (r_i, i) and define $u(x) = f(x)/f(1000)$ to obtain $u(1000) = 1$. In Figure 5, points $(r_i, u(r_i))$ are represented by round points (in red), and we show the resulting utility function (in blue).

The construction of the monotone spline with the certainty equivalent method is also shown in Figure 5; the elicited points are represented by + (in green), numbered by order of generation, and the resulting utility function is so close to the previous one that they are indiscernible.

5 CONCLUSION AND PROSPECTIVES

In this paper we proposed an incremental approach for the elicitation of the RDU model. A first novelty concerns the use, within the RDU theory, of minimax regret in order to incrementally elicit the probability weighting function w and to produce robust recommendations with respect to the uncertainty in probability weights. A second novelty is the use of monotonic regression splines as a model of the probability weighting function, allowing the representation of a wide variety of decision behaviors and the optimization of pairwise max-regrets by linear programming. We also extend the proposed approach to cases where the utility u is also unknown. Our experiments show that, despite the expressivity of the model, the elicitation burden is reasonably low in practice, due to the fixed and limited number of parameters used in splines approximating u and w , and the active learning process implemented for w . There are at least two natural continuations of this work. The first one consists in extending the approach for cumulative prospect theory to model different risk attitudes towards gains and losses. The second would be to jointly learn functions u and w using an incremental approach based on regret minimization, with the aim to save more preference queries.

Acknowledgement

Work supported by the French National Research Agency through the Idex Sorbonne Universit es, ELICIT project under grant ANR-11-IDEX-0004-02. We benefited from discussion with Angelina Vidali and Enrico Diecidue.

References

[Abdellaoui, 2000] M. Abdellaoui. Parameter-free elicitation of utility and probability weighting functions. *Management Science*, 46(11):1497–1512, 2000.

- [Avineri and Prashker, 2004] E. Avineri and J. Prashker. Violations of expected utility theory in route-choice stated preferences: certainty effect and inflation of small probabilities. *Transportation Research Record: J. of the Trans. Res. Board*, 1894:222–229, 2004.
- [Beatty and Barsky, 1995] J. C. Beatty and B. A. Barsky. *An introduction to splines for use in computer graphics and geometric modeling*. Morgan Kaufmann, 1995.
- [Benabbou *et al.*, 2014] N. Benabbou, P. Perny, and P. Viappiani. Incremental elicitation of Choquet capacities for multicriteria decision making. In *ECAI'14*, pages 87–92, 2014.
- [Bernoulli, 1954] D. Bernoulli. Exposition of a new theory on the measurement of risk. *Econometrica*, 22(1):23–36, 1954.
- [Boutilier *et al.*, 2006] C. Boutilier, R. Patrascu, P. Poupard, and D. Schuurmans. Constraint-based Optimization and Utility Elicitation using the Minimax Decision Criterion. *Artificial Intelligence*, 170(8–9):686–713, 2006.
- [Braziunas and Boutilier, 2007] D. Braziunas and C. Boutilier. Minimax Regret-based Elicitation of Generalized Additive Utilities. In *Proceedings of UAI'07*, pages 25–32, Vancouver, 2007.
- [Braziunas, 2011] D. Braziunas. *Decision-theoretic elicitation of generalized additive utilities*. PhD thesis, University of Toronto, 2011.
- [Cavagnaro *et al.*, 2013] D. Cavagnaro, M. Pitt, R. Gonzalez, and J. Myung. Discriminating among probability weighting functions using adaptive design optimization. *Journal of Risk and Uncertainty*, 47(3):255–289, 2013.
- [Chateauneuf and Cohen, 1994] A. Chateauneuf and M. Cohen. Risk seeking with diminishing marginal utility in a non-expected utility model. *Journal of Risk and Uncertainty*, 9(1):77–91, 1994.
- [Diecidue and Wakker, 2001] E. Diecidue and P. P. Wakker. On the intuition of rank-dependent utility. *Journal of Risk and Uncertainty*, 23(3):281–298, 2001.
- [Edwards, 1955] W. Edwards. The prediction of decisions among bets. *Journal of Experimental Psychology*, (50):201–214, 1955.
- [Fürnkranz and Hüllermeier, 2003] J. Fürnkranz and E. Hüllermeier. *Pairwise preference learning and ranking*, pages 145–156. Springer, 2003.
- [Gilboa, 2008] I. Gilboa. *Theory of Decision under Uncertainty*. Cambridge University Press, 2008.
- [Gonzalez and Wu, 1999] R. Gonzalez and G. Wu. On the shape of the probability weighting function. *Cognitive psychology*, 38(1):129–166, 1999.
- [Hines and Larson, 2010] G. Hines and K. Larson. Preference elicitation for risky prospects. In *AAMAS'10*, pages 889–896, 2010.
- [Hong *et al.*, 1987] C. S. Hong, E. Karni, and Z. Safra. Risk aversion in the theory of expected utility with rank dependent probabilities. *Journal of Economic theory*, 42(2):370–381, 1987.
- [Jeantet *et al.*, 2012] G. Jeantet, P. Perny, and O. Spanjaard. Sequential Decision Making with Rank Dependent Utility: a Minimax Regret Approach. In *Proc. of AAAI'12*, pages 1931–1937, 2012.
- [Kahneman and Tversky, 1979] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- [Karmarkar, 1978] U. S. Karmarkar. Subjectively weighted utility: A descriptive extension of the expected utility model. *Organizational Behavior and Human Performance*, 21(1):61–72, 1978.
- [Nielsen and Jaffray, 2006] T. D. Nielsen and J.-Y. Jaffray. Dynamic decision making without expected utility: An operational approach. *European Journal of Operational Research*, 169(1):226–246, 2006.
- [Quiggin, 1982] J. Quiggin. A theory of anticipated utility. *Journal of Economic Behavior and Organization*, (3):225–243, 1982.
- [Quiggin, 2012] J. Quiggin. *Generalized expected utility theory: The rank-dependent model*. Springer Science & Business Media, 2012.
- [Ramsay, 1988] J. O. Ramsay. Monotone regression spline in action. *Statistical Science*, pages 425–441, 1988.
- [Savage, 1954] L. J. Savage. *The Foundations of Statistics*. J. Wiley and Sons, New-York, 1954.
- [Schmeidler, 1989] D. Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.
- [Tehrani *et al.*, 2012] A. F. Tehrani, W. Cheng, K. Dembczynski, and E. Hüllermeier. Learning monotone nonlinear models using the Choquet integral. *Machine Learning*, 89(1-2):183–211, 2012.
- [Torra, 2010] V. Torra. Learning aggregation operators for preference modeling. In J. Fürnkranz and E. Hüllermeier, editors, *Preference learning*, pages 317–333. Springer, 2010.
- [Tversky and Kahneman, 1992] A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.
- [von Neumann and Morgenstern, 1947] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.
- [Wakker and Deneffe, 1996] P. Wakker and D. Deneffe. Eliciting von neumann-morgenstern utilities when probabilities are distorted or unknown. *Management science*, 42(8):1131–1150, 1996.
- [Wakker, 1990] P. Wakker. Under stochastic dominance Choquet-expected utility and anticipated utility are identical. *Theory and Decision*, 29(2):119–132, 1990.
- [Wakker, 1994] P. Wakker. Separating marginal utility and probabilistic risk aversion. *Theory and Decision*, 36(1):1–44, 1994.
- [Wang and Boutilier, 2003] T. Wang and C. Boutilier. Incremental Utility Elicitation with the Minimax Regret Decision Criterion. In *Proceedings of IJCAI'03*, pages 309–316, 2003.
- [Yaari, 1987] M. E. Yaari. The Dual Theory of Choice under Risk. *Econometrica*, 55(1):95–115, January 1987.