

Decision Making under Interval (and More General) Uncertainty: Monetary vs. Utility Approaches

Vladik Kreinovich

Department of Computer Science
University of Texas at El Paso, El Paso, TX 79968, USA
vladik@utep.edu

Abstract. In many situations, e.g., in financial and economic decision making, the decision results either in a money gain (or loss) and/or in the gain of goods that can be exchanged for money or for other goods. In such situations, interval uncertainty means that we do not know the exact amount of money that we will get for each possible decision, we only know lower and upper bounds on this amount. In this case, a natural idea is to assign a fair price to different alternatives, and then to use these fair prices to select the best alternative. In the talk, we show how to assign a fair price under interval uncertainty. We also explain how to assign a fair price in the case of more general types of uncertainty such as p-boxes (bounds on cumulative distribution function), twin intervals (when we only know approximate bounds), fuzzy values (when we have imprecise expert estimates of the gains), etc.

In other situations, e.g., when buying a house to live in or selecting a movie to watch, the result of the decision is the decision maker's own satisfaction. In such situations, a more adequate approach is to use utilities - a quantitative way of describing user's preferences. In this talk, after a brief introduction describing what are utilities, how to evaluate them, and how to make decisions based on utilities, we explain how to make decisions in situations with user uncertainty - a realistic situation when a decision maker cannot always decide which alternative is better for him or her.

Keywords: decision making, interval uncertainty

1 Need for Decision Making under Uncertainty

Need for decision making. In many practical situations, we have several alternatives, and we need to select one of these alternatives. For example:

- a person saving for retirement needs to find the best way to invest money;
- a company needs to select a location for its new plant;
- a designer must select one of several possible designs for a new airplane;
- a medical doctor needs to select a treatment for a patient.

Need for decision making under uncertainty. Decision making is easier if we know the exact consequences of each alternative selection. Often, however, we only have an incomplete information about consequences of different alternative, and we need to select an alternative under this uncertainty.

What we do in this paper. In this paper, we explain two approaches to decision making under uncertainty: monetary and utility approaches, we explain when each of these approaches is appropriate, how to justify the corresponding heuristic techniques, and how to go beyond these heuristic techniques.

Comments. Most results about the monetary approach are mentioned in [3], most results about the utility approach are described in [2], the combined description – with a clear separation of the two approaches – is new.

Decision making under uncertainty is a vast research area, with thousands of relevant books and papers; we refer readers to papers cited in [2, 3] and to the edited book [1].

2 Decision Making under Uncertainty: Monetary Approach and Its Heuristics

2.1 When monetary approach is appropriate

In many situations, e.g., in financial and economic decision making, the decision results:

- either in a money gain (or loss) and/or
- in the gain of goods that can be exchanged for money or for other goods.

In this case, we select an alternative which the highest exchange value, i.e., the highest price u .

Uncertainty means that we do not know the exact prices. The simplest case is when we only know lower and upper bounds on the price, i.e., we only know that $u \in [\underline{u}, \bar{u}]$ for given bounds \underline{u} and \bar{u} .

2.2 Hurwicz optimism-pessimism heuristic for decision making under interval uncertainty

In the early 1950s, the future Nobelist L. Hurwicz's proposed to select an alternative for which

$$\alpha_H \cdot \bar{u} + (1 - \alpha_H) \cdot \underline{u} \rightarrow \max.$$

Here, $\alpha_H \in [0, 1]$ described the optimism level of a decision maker:

- $\alpha_H = 1$ means optimism;
- $\alpha_H = 0$ means pessimism;
- $0 < \alpha_H < 1$ combines optimism and pessimism.

This approach works well in practice. However, this is a semi-heuristic idea.

It is desirable to come up with an approach which can be uniquely determined based first principles.

3 Monetary Approach to Decision Making under Uncertainty: How to Justify the Existing Heuristics and How to Move Beyond these Heuristics

3.1 Fair price approach: an idea

When we have a full information about an object, then we can express our desirability of each possible situation by declaring a price that we are willing to pay to get involved in this situation. Once these prices are set, we simply select the alternative for which the participation price is the highest.

In decision making under uncertainty, it is not easy to come up with a fair price. A natural idea is to develop techniques for producing such fair prices. These prices can then be used in decision making, to select an appropriate alternative.

3.2 Case of interval uncertainty

Analysis of the problem. In the ideal case, we know the exact gain u of selecting an alternative. A more realistic case is when we only know the lower bound \underline{u} and the upper bound \bar{u} on this gain – and we do not know which values $u \in [\underline{u}, \bar{u}]$ are more probable or less probable. This situation is known as *interval uncertainty*.

We want to assign, to each interval $[\underline{u}, \bar{u}]$, a number $P([\underline{u}, \bar{u}])$ describing the fair price of this interval.

Since we know that $u \leq \bar{u}$, we have $P([\underline{u}, \bar{u}]) \leq \bar{u}$. Similarly, since we know that $u \geq \underline{u}$, we have $\underline{u} \leq P([\underline{u}, \bar{u}])$.

Monotonicity. Let us first consider the case when we keep the lower endpoint \underline{u} intact but increase the upper bound. This means that we:

- keep all the previous possibilities, but
- we allow new possibilities, with a higher gain.

In this case, it is reasonable to require that the corresponding price not decrease:

$$\text{if } \underline{u} = \underline{v} \text{ and } \bar{u} < \bar{v} \text{ then } P([\underline{u}, \bar{u}]) \leq P([\underline{v}, \bar{v}]).$$

Let us now consider another case, when we dismiss some low-gain alternatives. This should increase (or at least not decrease) the fair price:

$$\text{if } \underline{u} < \underline{v} \text{ and } \bar{u} = \bar{v} \text{ then } P([\underline{u}, \bar{u}]) \leq P([\underline{v}, \bar{v}]).$$

Additivity: idea. Let us consider another requirement on the fair price. This requirement is related to the fact that we can consider two decision processes separately. Alternatively, we can also consider a single decision process in which we select a pair of alternatives:

- the 1st alternative corresponding to the 1st decision, and
- the 2nd alternative corresponding to the 2nd decision.

If we are willing to pay:

- the amount u to participate in the first process, and
- the amount v to participate in the second decision process,

then we should be willing to pay $u + v$ to participate in both decision processes.

Additivity: case of interval uncertainty. Let us describe what this requirement will look like in the case of interval uncertainty.

In this case, about the gain u from the first alternative, we only know that this (unknown) gain is in $[\underline{u}, \bar{u}]$. About the gain v from the second alternative, we only know that this gain belongs to the interval $[\underline{v}, \bar{v}]$. The overall gain $u + v$ can thus take any value from the interval

$$[\underline{u}, \bar{u}] + [\underline{v}, \bar{v}] \stackrel{\text{def}}{=} \{u + v : u \in [\underline{u}, \bar{u}], v \in [\underline{v}, \bar{v}]\}.$$

It is easy to check that $[\underline{u}, \bar{u}] + [\underline{v}, \bar{v}] = [\underline{u} + \underline{v}, \bar{u} + \bar{v}]$. Thus, the additivity requirement about the fair prices takes the form

$$P([\underline{u} + \underline{v}, \bar{u} + \bar{v}]) = P([\underline{u}, \bar{u}]) + P([\underline{v}, \bar{v}]).$$

Fair price under interval uncertainty. Let us see what all these requirement lead to.

Definition 1. By a fair price under interval uncertainty, we mean a function $P([\underline{u}, \bar{u}])$ for which:

- $\underline{u} \leq P([\underline{u}, \bar{u}]) \leq \bar{u}$ for all \underline{u} and \bar{u} (conservativeness);
- if $\underline{u} = \underline{v}$ and $\bar{u} < \bar{v}$, then $P([\underline{u}, \bar{u}]) \leq P([\underline{v}, \bar{v}])$ (monotonicity);
- (additivity) for all $\underline{u}, \bar{u}, \underline{v}, \bar{v}$, we have

$$P([\underline{u} + \underline{v}, \bar{u} + \bar{v}]) = P([\underline{u}, \bar{u}]) + P([\underline{v}, \bar{v}]).$$

Theorem 1. Each fair price under interval uncertainty has the form

$$P([\underline{u}, \bar{u}]) = \alpha_H \cdot \bar{u} + (1 - \alpha_H) \cdot \underline{u} \text{ for some } \alpha_H \in [0, 1].$$

Comment: we thus get a new justification of the Hurwicz optimism-pessimism criterion.

Proof: main ideas.

- Due to monotonicity, $P([u, u]) = u$.
- Due to monotonicity, $\alpha_H \stackrel{\text{def}}{=} P([0, 1]) \in [0, 1]$.
- For $[0, 1] = [0, 1/n] + \dots + [0, 1/n]$ (n times), additivity implies that $\alpha_H = n \cdot P([0, 1/n])$, so $P([0, 1/n]) = \alpha_H \cdot (1/n)$.
- For $[0, m/n] = [0, 1/n] + \dots + [0, 1/n]$ (m times), additivity implies $P([0, m/n]) = \alpha_H \cdot (m/n)$.
- For each real number r , for each n , there is an m s.t. $m/n \leq r \leq (m+1)/n$.
- Monotonicity implies

$$\alpha_H \cdot (m/n) = P([0, m/n]) \leq P([0, r]) \leq P([0, (m+1)/n]) = \alpha_H \cdot ((m+1)/n).$$

- When $n \rightarrow \infty$, $\alpha_H \cdot (m/n) \rightarrow \alpha_H \cdot r$ and $\alpha_H \cdot ((m+1)/n) \rightarrow \alpha_H \cdot r$, hence $P([0, r]) = \alpha_H \cdot r$.
- For $[\underline{u}, \bar{u}] = [\underline{u}, \underline{u}] + [0, \bar{u} - \underline{u}]$, additivity implies $P([\underline{u}, \bar{u}]) = \underline{u} + \alpha_H \cdot (\bar{u} - \underline{u})$. Q.E.D.

3.3 Case of set-valued uncertainty

Need for set-valued uncertainty. In some cases, in addition to knowing that the actual gain belongs to the interval $[\underline{u}, \bar{u}]$, we also know that some values from this interval cannot be possible values of this gain.

For example, if we buy an obscure lottery ticket for a simple prize-or-no-prize lottery from a remote country, we either get the prize or lose the money. In this case, the set of possible values of the gain consists of two values.

Instead of a (bounded) *interval* of possible values, we can therefore consider a general bounded *set* of possible values.

Fair price under set-valued uncertainty. We want a function P that assigns, to every bounded closed set S , a real number $P(S)$, for which:

- $P([\underline{u}, \bar{u}]) = \alpha_H \cdot \bar{u} + (1 - \alpha_H) \cdot \underline{u}$ (*conservativeness*);
- $P(S + S') = P(S) + P(S')$, where $S + S' \stackrel{\text{def}}{=} \{s + s' : s \in S, s' \in S'\}$ (*additivity*).

Theorem 2. *Each fair price under set uncertainty has the form*

$$P(S) = \alpha_H \cdot \sup S + (1 - \alpha_H) \cdot \inf S.$$

Proof: idea.

- $\{\underline{s}, \bar{s}\} \subseteq S \subseteq [\underline{s}, \bar{s}]$, where $\underline{s} \stackrel{\text{def}}{=} \inf S$ and $\bar{s} \stackrel{\text{def}}{=} \sup S$;
- thus, $[2\underline{s}, 2\bar{s}] = \{\underline{s}, \bar{s}\} + [\underline{s}, \bar{s}] \subseteq S + [\underline{s}, \bar{s}] \subseteq [\underline{s}, \bar{s}] + \{\underline{s}, \bar{s}\} = [2\underline{s}, 2\bar{s}]$;
- so $S + [\underline{s}, \bar{s}] = [2\underline{s}, 2\bar{s}]$, hence $P(S) + P([\underline{s}, \bar{s}]) = P([2\underline{s}, 2\bar{s}])$, and

$$P(S) = (\alpha_H \cdot (2\bar{s}) + (1 - \alpha_H) \cdot (2\underline{s})) - (\alpha_H \cdot \bar{s} + (1 - \alpha_H) \cdot \underline{s}).$$

3.4 Case of probabilistic uncertainty

Suppose that for some financial instrument, we know a probability distribution $\rho(x)$ on the set of possible gains x . What is the fair price P for this instrument?

Due to additivity, the fair price for n copies of this instrument is $n \cdot P$. According to the Large Numbers Theorem, for large n , the average gain tends to the mean value

$$\mu = \int x \cdot \rho(x) dx.$$

Thus, the fair price for n copies of the instrument is close to $n \cdot \mu$: $n \cdot P \approx n \cdot \mu$. The larger n , the closer the averages. So, in the limit, we get $P = \mu$, i.e., the fair price is the mean value.

3.5 Case of p-box uncertainty

Probabilistic uncertainty means that for every x , we know the value of the cdf $F(x) = \text{Prob}(\eta \leq x)$. In practice, we often only have partial information about these values.

In this case, for each x , we only know an interval $[\underline{F}(x), \overline{F}(x)]$ containing the actual (unknown) value $F(x)$. The interval-valued function $[\underline{F}(x), \overline{F}(x)]$ is known as a *p-box*.

What is the fair price of a p-box? The only information that we have about the cdf is that $F(x) \in [\underline{F}(x), \overline{F}(x)]$. For each possible $F(x)$, for large n , n copies of the instrument are \approx equivalent to $n \cdot \mu$, with $\mu = \int x dF(x)$.

For different $F(x)$ from the p-box, values of μ for an interval $[\underline{\mu}, \overline{\mu}]$, where $\underline{\mu} = \int x d\underline{F}(x)$ and $\overline{\mu} = \int x d\overline{F}(x)$. Thus, the price of a p-box is equal to the price of an interval $[\underline{\mu}, \overline{\mu}]$.

We already know that this price is equal to $\alpha_H \cdot \overline{\mu} + (1 - \alpha_H) \cdot \underline{\mu}$. So, this is a fair price of a p-box.

3.6 Case of twin intervals

What are twin intervals. Sometimes, in addition to the interval $[\underline{x}, \overline{x}]$, we also have a “most probable” subinterval $[\underline{m}, \overline{m}] \subseteq [\underline{x}, \overline{x}]$. For such “twin intervals”, addition is naturally defined component-wise:

$$([\underline{x}, \overline{x}], [\underline{m}, \overline{m}]) + ([\underline{y}, \overline{y}], [\underline{n}, \overline{n}]) = ([\underline{x} + \underline{y}, \overline{x} + \overline{y}], [\underline{m} + \underline{n}, \overline{m} + \overline{n}]).$$

Thus, the additivity for additivity requirement about the fair prices takes the form

$$P([\underline{x} + \underline{y}, \overline{x} + \overline{y}], [\underline{m} + \underline{n}, \overline{m} + \overline{n}]) = P([\underline{x}, \overline{x}], [\underline{m}, \overline{m}]) + P([\underline{y}, \overline{y}], [\underline{n}, \overline{n}]).$$

Fair price under twin interval uncertainty. What is the fair price under such uncertainty?

Definition 2. By a fair price under twin uncertainty, we mean a function $P([\underline{u}, \overline{u}], [\underline{m}, \overline{m}])$ for which:

- $\underline{u} \leq P([\underline{u}, \overline{u}], [\underline{m}, \overline{m}]) \leq \overline{u}$ for all $\underline{u} \leq \underline{m} \leq \overline{m} \leq \overline{u}$ (conservativeness);
- if $\underline{u} \leq \underline{v}$, $\underline{m} \leq \underline{n}$, $\overline{m} \leq \overline{n}$, and $\overline{u} \leq \overline{v}$, then $P([\underline{u}, \overline{u}], [\underline{m}, \overline{m}]) \leq P([\underline{v}, \overline{v}], [\underline{n}, \overline{n}])$ (monotonicity);
- for all $\underline{u} \leq \underline{m} \leq \overline{m} \leq \overline{u}$ and $\underline{v} \leq \underline{n} \leq \overline{n} \leq \overline{v}$, we have additivity:

$$P([\underline{u} + \underline{v}, \overline{u} + \overline{v}], [\underline{m} + \underline{n}, \overline{m} + \overline{n}]) = P([\underline{u}, \overline{u}], [\underline{m}, \overline{m}]) + P([\underline{v}, \overline{v}], [\underline{n}, \overline{n}]).$$

Theorem 3. Each fair price under twin uncertainty has the following form, for some $\alpha_L, \alpha_u, \alpha_U \in [0, 1]$:

$$P([\underline{u}, \overline{u}], [\underline{m}, \overline{m}]) = \underline{m} + \alpha_u \cdot (\overline{m} - \underline{m}) + \alpha_U \cdot (\overline{U} - \overline{m}) + \alpha_L \cdot (\underline{u} - \underline{m}).$$

3.7 Case of fuzzy numbers

Need for fuzzy uncertainty. An expert is often imprecise (“fuzzy”) about the possible values. For example, an expert may say that the gain is small. To describe such information, L. Zadeh introduced the notion of *fuzzy numbers*.

For fuzzy numbers, different values u are possible with different degrees $\mu(u) \in [0, 1]$. The value w is a possible value of $u + v$ if:

- for some values u and v for which $u + v = w$,
- u is a possible value of 1st gain, and
- v is a possible value of 2nd gain.

If we interpret “and” as min and “or” (“for some”) as max, we get *Zadeh’s extension principle*:

$$\mu(w) = \max_{u,v: u+v=w} \min(\mu_1(u), \mu_2(v)).$$

This operation is easiest to describe in terms of α -cuts

$$\mathbf{u}(\alpha) = [u^-(\alpha), u^+(\alpha)] \stackrel{\text{def}}{=} \{u : \mu(u) \geq \alpha\}.$$

Namely, $\mathbf{w}(\alpha) = \mathbf{u}(\alpha) + \mathbf{v}(\alpha)$, i.e.,

$$w^-(\alpha) = u^-(\alpha) + v^-(\alpha) \text{ and } w^+(\alpha) = u^+(\alpha) + v^+(\alpha).$$

For product (of probabilities), we similarly get

$$\mu(w) = \max_{u,v: u \cdot v = w} \min(\mu_1(u), \mu_2(v)).$$

In terms of α -cuts, we have $\mathbf{w}(\alpha) = \mathbf{u}(\alpha) \cdot \mathbf{v}(\alpha)$, i.e.,

$$w^-(\alpha) = u^-(\alpha) \cdot v^-(\alpha) \text{ and } w^+(\alpha) = u^+(\alpha) \cdot v^+(\alpha).$$

What is the fair price under such fuzzy uncertainty?

Fair price under fuzzy uncertainty. We want to assign, to every fuzzy number s , a real number $P(s)$, so that:

- if a fuzzy number s is located between \underline{u} and \bar{u} , then $\underline{u} \leq P(s) \leq \bar{u}$ (*conservativeness*);
- $P(u + v) = P(u) + P(v)$ (*additivity*);
- if for all α , $s^-(\alpha) \leq t^-(\alpha)$ and $s^+(\alpha) \leq t^+(\alpha)$, then we have $P(s) \leq P(t)$ (*monotonicity*);
- if μ_n uniformly converges to μ , then $P(\mu_n) \rightarrow P(\mu)$ (*continuity*).

Theorem 4. *The fair price under fuzzy uncertainty is equal to*

$$P(s) = s_0 + \int_0^1 k^-(\alpha) ds^-(\alpha) - \int_0^1 k^+(\alpha) ds^+(\alpha) \text{ for some } k^\pm(\alpha).$$

Discussion. Here, $\int f(x) \cdot dg(x) = \int f(x) \cdot g'(x) dx$ for a *generalized function* $g'(x)$, hence for generalized $K^\pm(\alpha)$, we have:

$$P(s) = \int_0^1 K^-(\alpha) \cdot s^-(\alpha) d\alpha + \int_0^1 K^+(\alpha) \cdot s^+(\alpha) d\alpha.$$

Conservativeness means that

$$\int_0^1 K^-(\alpha) d\alpha + \int_0^1 K^+(\alpha) d\alpha = 1.$$

For the interval $[\underline{u}, \bar{u}]$, we get

$$P(s) = \left(\int_0^1 K^-(\alpha) d\alpha \right) \cdot \underline{u} + \left(\int_0^1 K^+(\alpha) d\alpha \right) \cdot \bar{u}.$$

Thus, Hurwicz optimism-pessimism coefficient α_H is equal to $\int_0^1 K^+(\alpha) d\alpha$. In this sense, the above formula is a generalization of Hurwicz's formula to the fuzzy case.

4 Non-Monetary (Utility) Approach to Decision Making under Uncertainty: Main Ideas

4.1 Monetary approach is not always appropriate

In some situations, the result of the decision is the decision maker's own satisfaction; examples include:

- buying a house to live in,
- selecting a movie to watch.

In such situations, monetary approach is not appropriate; for example:

- a small apartment in downtown can be very expensive,
- but many people prefer a cheaper – but more spacious and comfortable – suburban house.

4.2 Non-monetary decision making: traditional utility approach

Analysis of the problem. To make a decision, we must find out the user's preference, and help the user select an alternative which is the best – according to these preferences.

Traditional approach is based on an assumption that for each two alternatives A' and A'' , a user can tell:

- whether the first alternative is better for him/her; we will denote this by $A'' < A'$;
- or the second alternative is better; we will denote this by $A' < A''$;
- or the two given alternatives are of equal value to the user; we will denote this by $A' = A''$.

The notion of utility. Under the above assumption, we can form a natural numerical scale for describing preferences.

Let us select a very bad alternative A_0 and a very good alternative A_1 . Then, most other alternatives are better than A_0 but worse than A_1 . For every probability $p \in [0, 1]$, we can form a lottery $L(p)$ in which we get A_1 with probability p and A_0 with probability $1 - p$.

When $p = 0$, this lottery simply coincides with the alternative A_0 : $L(0) = A_0$. The larger the probability p of the positive outcome increases, the better the result:

$$p' < p'' \text{ implies } L(p') < L(p'').$$

Finally, for $p = 1$, the lottery coincides with the alternative A_1 : $L(1) = A_1$.

Thus, we have a continuous scale of alternatives $L(p)$ that monotonically goes from $L(0) = A_0$ to $L(1) = A_1$. Due to monotonicity, when p increases, we first have $L(p) < A$, then we have $L(p) > A$. The threshold value is called the *utility* of the alternative A :

$$u(A) \stackrel{\text{def}}{=} \sup\{p : L(p) < A\} = \inf\{p : L(p) > A\}.$$

Then, for every $\varepsilon > 0$, we have

$$L(u(A) - \varepsilon) < A < L(u(A) + \varepsilon).$$

We will describe such (almost) equivalence by \equiv , i.e., we will write that $A \equiv L(u(A))$.

Fast iterative process for determining $u(A)$. How can we determine the utility value?

Initially, we know the values $\underline{u} = 0$ and $\bar{u} = 1$ such that $A \equiv L(u(A))$ for some $u(A) \in [\underline{u}, \bar{u}]$.

In general, once we know an interval $[\underline{u}, \bar{u}]$ containing $u(A)$, we compute the midpoint u_{mid} of this interval and compare A with $L(u_{\text{mid}})$.

- If $A \leq L(u_{\text{mid}})$, then $u(A) \leq u_{\text{mid}}$, so we know that $u \in [\underline{u}, u_{\text{mid}}]$.
- If $L(u_{\text{mid}}) \leq A$, then $u_{\text{mid}} \leq u(A)$, so $u \in [u_{\text{mid}}, \bar{u}]$.

After each iteration, we decrease the width of the interval $[\underline{u}, \bar{u}]$ by half. After k iterations, we get an interval of width 2^{-k} which contains $u(A)$ – i.e., we get $u(A)$ with accuracy 2^{-k} .

How to make a decision based on utility values. Suppose that we have found the utilities $u(A')$, $u(A'')$, \dots , of the alternatives A' , A'' , \dots . Which of these alternatives should we choose?

By definition of utility, we have $A \equiv L(u(A))$ for every alternative A , and $L(p') < L(p'')$ if and only if $p' < p''$. We can thus conclude that A' is preferable to A'' if and only if $u(A') > u(A'')$. In other words, we should always select an alternative with the largest possible value of utility.

How to estimate utility of an action. For each action, we usually know possible outcomes S_1, \dots, S_n . We can often estimate the probabilities p_1, \dots, p_n of these outcomes. By definition of utility, each situation S_i is equivalent to a lottery $L(u(S_i))$ in which we get:

- A_1 with probability $u(S_i)$ and
- A_0 with the remaining probability $1 - u(S_i)$.

Thus, the action is equivalent to a complex lottery in which:

- first, we select one of the situations S_i with probability p_i : $P(S_i) = p_i$;
- then, depending on S_i , we get A_1 with probability $P(A_1 | S_i) = u(S_i)$ and A_0 with probability $1 - u(S_i)$.

The probability of getting A_1 in this complex lottery is:

$$P(A_1) = \sum_{i=1}^n P(A_1 | S_i) \cdot P(S_i) = \sum_{i=1}^n u(S_i) \cdot p_i.$$

In the complex lottery, we get:

- A_1 with prob. $u = \sum_{i=1}^n p_i \cdot u(S_i)$, and
- A_0 with probability $1 - u$.

So, we should select the action with the largest value of expected utility

$$u = \sum p_i \cdot u(S_i).$$

Non-uniqueness of utility. The above definition of utility u depends on A_0, A_1 . What if we use different alternatives A'_0 and A'_1 ?

Every A is equivalent to a lottery $L(u(A))$ in which we get A_1 with probability $u(A)$ and A_0 with probability $1 - u(A)$. For simplicity, let us assume that $A'_0 < A_0 < A_1 < A'_1$. Then, $A_0 \equiv L'(u'(A_0))$ and $A_1 \equiv L'(u'(A_1))$. So, A is equivalent to a complex lottery in which:

- 1) we select A_1 with probability $u(A)$ and A_0 with probability $1 - u(A)$;
- 2) depending on A_i , we get A'_1 with probability $u'(A_i)$ and A'_0 with probability $1 - u'(A_i)$.

In this complex lottery, we get A'_1 with probability

$$u'(A) = u(A) \cdot (u'(A_1) - u'(A_0)) + u'(A_0).$$

So, in general, utility is defined modulo an (increasing) linear transformation $u' = a \cdot u + b$, with $a > 0$.

Subjective probabilities. In practice, we often do not know the probabilities p_i of different outcomes. For each event E , a natural way to estimate its subjective probability is to fix a prize (e.g., \$1) and compare:

- the lottery ℓ_E in which we get the fixed prize if the event E occurs and 0 if it does not occur, with
- a lottery $\ell(p)$ in which we get the same amount with probability p .

Here, similarly to the utility case, we get a value $ps(E)$ for which, for every $\varepsilon > 0$:

$$\ell(ps(E) - \varepsilon) < \ell_E < \ell(ps(E) + \varepsilon).$$

Then, the utility of an action with possible outcomes S_1, \dots, S_n is equal to $u = \sum_{i=1}^n ps(E_i) \cdot u(S_i)$.

Beyond traditional decision making: towards a more realistic description. Previously, we assumed that a user can always decide which of the two alternatives A' and A'' is better:

- either $A' < A''$,
- or $A'' < A'$,
- or $A' \equiv A''$.

In practice, a user is sometimes unable to meaningfully decide between the two alternatives; we will denote this by $A' \parallel A''$. In mathematical terms, this means that the preference relation is no longer a *total* (linear) order, it can be a *partial* order.

From utility to interval-valued utility. Similarly to the traditional decision making approach:

- we select two alternatives $A_0 < A_1$ and
- we compare each alternative A which is better than A_0 and worse than A_1 with lotteries $L(p)$.

Since preference is a *partial* order, in general:

$$\underline{u}(A) \stackrel{\text{def}}{=} \sup\{p : L(p) < A\} < \bar{u}(A) \stackrel{\text{def}}{=} \inf\{p : L(p) > A\}.$$

For each alternative A , instead of a single value $u(A)$ of the utility, we now have an *interval* $[\underline{u}(A), \bar{u}(A)]$ such that:

- if $p < \underline{u}(A)$, then $L(p) < A$;
- if $p > \bar{u}(A)$, then $A < L(p)$; and
- if $\underline{u}(A) < p < \bar{u}(A)$, then $A \parallel L(p)$.

We will call this interval the *utility* of the alternative A .

Interval-valued utilities and interval-valued subjective probabilities. To feasibly elicit the values $\underline{u}(A)$ and $\bar{u}(A)$, we:

- 1) starting with $[\underline{u}, \bar{u}] = [0, 1]$, bisect an interval such that $L(\underline{u}) < A < L(\bar{u})$ until we find u_0 for which $A \parallel L(u_0)$;
- 2) by bisecting an interval $[\underline{u}, u_0]$ for which $L(\underline{u}) < A \parallel L(u_0)$, we find $\underline{u}(A)$;
- 3) by bisecting an interval $[u_0, \bar{u}]$ for which $L(u_0) \parallel A < L(\bar{u})$, we find $\bar{u}(A)$.

Similarly, when we estimate the probability of an event E , we no longer get a single value $ps(E)$, get an interval $[ps(E), \bar{ps}(E)]$ of possible values of probability. By using bisection, we can feasibly elicit the values $\underline{ps}(E)$ and $\bar{ps}(E)$.

Decision making under interval uncertainty. For each possible decision d , we know the interval $[\underline{u}(d), \bar{u}(d)]$ of possible values of utility. Which decision shall we select?

A natural idea is to select all decisions d_0 that *may* be optimal, i.e., which are optimal for some function $u(d) \in [\underline{u}(d), \bar{u}(d)]$.

Checking all possible functions is not feasible. However, it is easy to show that the above condition equivalent to an easier-to-check one: $\bar{u}(d_0) \geq \max_d \underline{u}(d)$.

The remaining problem is that in practice, we would like to select *one* decision; which one should be select?

Need for definite decision making. At first glance, if $A' \parallel A''$, it does not matter whether we recommend alternative A' or alternative A'' . Let us show that this is *not* a good recommendation.

Let A be an alternative about which we know nothing, i.e., for which $[\underline{u}(A), \bar{u}(A)] = [0, 1]$. In this case, A is indistinguishable both from a “good” lottery $L(0.999)$ and a “bad” lottery $L(0.001)$. Suppose that we recommend, to the user, that A is equivalent both to $L(0.999)$ and to $L(0.001)$. Then this user will feel comfortable:

- first, exchanging $L(0.999)$ with A , and
- then, exchanging A with $L(0.001)$.

So, following our recommendations, the user switches from a very good alternative to a very bad one.

The above argument does not depend on the fact that we assumed complete ignorance about A :

- every time we recommend that the alternative A is “equivalent” both to $L(p)$ and to $L(p')$ ($p < p'$),
- we make the user vulnerable to a similar switch from a better alternative $L(p')$ to a worse one $L(p)$.

Thus, there should be only a single value p for which A can be reasonably exchanged with $L(p)$. In precise terms:

- we start with the utility interval $[\underline{u}(A), \bar{u}(A)]$, and
- we need to select a single $u(A)$ for which it is reasonable to exchange A with a lottery $L(u)$.

How can we find this value $u(A)$?

5 Utility Approach to Decision Making under Uncertainty: How to Justify the Existing Heuristics and How to Move Beyond these Heuristics

Decisions under interval uncertainty: a new justification for the Hurwicz optimism-pessimism criterion. We need to assign, to each interval $[\underline{u}, \bar{u}]$, a utility value $u(\underline{u}, \bar{u}) \in [\underline{u}, \bar{u}]$. Let us denote $\alpha_H \stackrel{\text{def}}{=} u(0, 1)$.

As we have mentioned earlier, utility is determined modulo a linear transformation $u' = a \cdot u + b$. It is therefore reasonable to require that the equivalent utility does not change with re-scaling: for $a > 0$ and b ,

$$u(a \cdot u^- + b, a \cdot u^+ + b) = a \cdot u(u^-, u^+) + b.$$

In particular, for $u^- = 0$, $u^+ = 1$, $a = \bar{u} - \underline{u}$, and $b = \underline{u}$, we get

$$u(\underline{u}, \bar{u}) = \alpha_H \cdot (\bar{u} - \underline{u}) + \underline{u} = \alpha_H \cdot \bar{u} + (1 - \alpha_H) \cdot \underline{u}.$$

This is exactly Hurwicz's *optimism-pessimism criterion*!

Which value α_H should we choose? An argument in favor of $\alpha_H = 0.5$. Let us take an event E about which we know nothing. For a lottery L^+ in which we get A_1 if E and A_0 otherwise, the utility interval is $[0, 1]$. Thus, the equivalent utility of L^+ is $\alpha_H \cdot 1 + (1 - \alpha_H) \cdot 0 = \alpha_H$.

For a lottery L^- in which we get A_0 if E and A_1 otherwise, the utility is $[0, 1]$, so equivalent utility is also α_H .

For a complex lottery L in which we select either L^+ or L^- with probability 0.5, the equivalent utility is still α_H . On the other hand, in L , we get A_1 with probability 0.5 and A_0 with probability 0.5. Thus, $L = L(0.5)$ and hence, $u(L) = 0.5$. So, we conclude that $\alpha_H = 0.5$.

Which action should we choose? Suppose that an action has n possible outcomes S_1, \dots, S_n , with utilities $[\underline{u}(S_i), \bar{u}(S_i)]$, and probabilities $[\underline{p}_i, \bar{p}_i]$. We know that each alternative is equivalent to a simple lottery with utility $u_i = \alpha_H \cdot \bar{u}(S_i) + (1 - \alpha_H) \cdot \underline{u}(S_i)$. We know that for each i , the i -th event is equivalent to $p_i = \alpha_H \cdot \bar{p}_i + (1 - \alpha_H) \cdot \underline{p}_i$.

Thus, this action is equivalent to a situation in which we get utility u_i with probability p_i . The utility of such a situation is equal to $\sum_{i=1}^n p_i \cdot u_i$. So, the equivalent utility of the original action is equivalent to

$$\sum_{i=1}^n \left(\alpha_H \cdot \bar{p}_i + (1 - \alpha_H) \cdot \underline{p}_i \right) \cdot \left(\alpha_H \cdot \bar{u}(S_i) + (1 - \alpha_H) \cdot \underline{u}(S_i) \right).$$

Observation: the resulting decision depends on the level of detail. Let us consider a situation in which, with some probability p , we gain a utility u , else we get 0. The expected utility is $p \cdot u + (1 - p) \cdot 0 = p \cdot u$.

Suppose that we only know the intervals $[\underline{u}, \bar{u}]$ and $[\underline{p}, \bar{p}]$. The equivalent utility u_k (k for *know*) is

$$u_k = (\alpha_H \cdot \bar{p} + (1 - \alpha_H) \cdot \underline{p}) \cdot (\alpha_H \cdot \bar{u} + (1 - \alpha_H) \cdot \underline{u}).$$

If we only know that utility is from $[\underline{p} \cdot \underline{u}, \bar{p} \cdot \bar{u}]$, then:

$$u_d = \alpha_H \cdot \bar{p} \cdot \bar{u} + (1 - \alpha_H) \cdot \underline{p} \cdot \underline{u} \text{ (} d \text{ for } \textit{don't know}\text{)}.$$

Here, additional knowledge decreases utility:

$$u_d - u_k = \alpha_H \cdot (1 - \alpha_H) \cdot (\bar{p} - \underline{p}) \cdot (\bar{u} - \underline{u}) > 0.$$

(This is maybe what the Book of Ecclesiastes meant by “For with much wisdom comes much sorrow”?)

Acknowledgments

This work was supported in part by the National Science Foundation grants HRD-0734825 and HRD-1242122 (Cyber-ShARE Center of Excellence) and DUE-0926721,

The author is greatly thankful to the conference organizers for their support.

References

1. V. Kreinovich, “Decision Making under Interval Uncertainty (and beyond)”, In: P. Guo and W. Pedrycz (eds.), *Human-Centric Decision-Making Models for Social Sciences*, Springer Verlag, 2014, pp. 163–193.
2. V. Kreinovich, “Decision Making under Interval Uncertainty (and beyond)”, In: P. Guo and W. Pedrycz (eds.), *Human-Centric Decision-Making Models for Social Sciences*, Springer Verlag, 2014, pp. 163–193.
3. J. Lorkowski and V. Kreinovich, “How Much For an Interval? a Set? a Twin Set? a p-Box? A Kaucher Interval? Towards an Economics-Motivated Approach to Decision Making Under Uncertainty”, *Proceedings of the 16th GAMM-IMACS International Symposium on Scientific Computing, Computer Arithmetic, and Verified Numerical Computation SCAN’2014*, Wuerzburg, Germany, September 21–26, 2014.