DECISION MAKING UNDER UNCERTAINTY WITH APPLICATIONS TO GEOSCIENCES AND FINANCE

LAXMAN BOKATI

Master's Program in Computational Science

APPROVED:

____
Vladik Kreinovich, Ph.D., Chair

____
Martine Ceberio, Ph.D.

____
Aaron Velasco, Ph.D.

____
Stephen L. Crites, Ph.D.
Dean of the Graduate School

DECISION MAKING UNDER UNCERTAINTY WITH APPLICATIONS TO GEOSCIENCES AND FINANCE

by

LAXMAN BOKATI

THESIS

Presented to the Faculty of the Graduate School of

The University of Texas at El Paso

in Partial Fulfillment

of the Requirements

for the Degree of

MASTER OF SCIENCE

Computational Science Program

THE UNIVERSITY OF TEXAS AT EL PASO

May 2019

Acknowledgements

Research cannot be done alone, it requires support of many people to perform research. Although it is not possible to name every individual that has directly or indirectly helped shape up this thesis, I have to thank certain people that were integral to completion of this thesis.

First and foremost, it would not have been possible for me to complete this thesis without remarkable support of my research advisor Prof. Vladik Kreinovich. Prof. Kreinovich is undoubtedly the best professor I have ever had in my academic career. His constant support and encouragement is the driving force behind success of this thesis work.

Besides my advisor, my sincere gratitude also goes to rest of my thesis committee: Prof. Aaron Velasco and Dr. Martine Ceberio, for their enlightening and insightful suggestions that helped finalize this thesis.

I would also like to thank US National Science Foundation for supporting the work done in this thesis via grant HRD-1242122 (Cyber-ShARE Center of Excellence).

I also have to thank Dr. Olga Kosheleva for all the help and suggestions without which this thesis would not have been complete.

Special thanks to Pawan Koirala, Kamal Nyaupane, Neelam Dumre, Bibek Aryal, and Subharaj Ranabhat for assisting me to organize my thesis presentation.

Finally, I must record a debt to my father Damber Dutta Bokati and mother Ishwori Devi Bokati as well as my brother Prakash Bokati for their constant inspiration, love and support without which I would not have been able to pursue my academic goals.

Abstract

In many practical situations, we need to make a decision. In engineering, we need to decide on the best design of a system, and, for existing systems – on the best control strategy. In financial applications, we need to decide what is the best way to invest money. In geosciences, we need to decide whether we should explore a possible mineral deposit – or whether we should perform more experiments and measurements (and what exactly).

In some cases, we can compute the exact consequences of each decision – e.g., if we are controlling a satellite. However, in many other cases, we do not know the exact consequences. In such situations, we need to make a decision under uncertainty.

In many application areas, uncertainty is small – and can be made even smaller by appropriate measurements. For example, when we control a self-driving car, if there is an uncertainty about the locations and speeds of other objects on the road, we can install more accurate sensors and get a better description of the driving environment.

However, there are applications when it is difficult or even impossible to decrease uncertainty. One such area is anything related to human activities. Humans make individual decisions based on their perceived value of different alternatives. The same alternative – be it a movie or a computer – have drastically different value to different people, so it is very difficult to predict their behavior. Such behavior affects economics and finance, so in decision making in economics and finance, it is very important to take such decision making under uncertainty into account.

Another area where it is difficult to decrease uncertainty is geosciences. In this thesis, we analyze the general problem of decision making under uncertainty and show how our results can be applied to finances and geosciences.

Table of Contents

Chapter 1

Introduction

Need for decision making. In many practical situations, we need to make a decision.

- In engineering, we need to decide on the best design of a system, and, for existing systems – on the best control strategy.
- In financial applications, we need to decide what is the best way to invest money.
- In geosciences, we need to decide whether we should explore a possible mineral deposit – or whether we should perform more experiments and measurements (and what exactly).

Need for decision making under uncertainty. In some cases, we can compute the exact consequences of each decision – e.g., if we are controlling a satellite. However, in many other cases, we do not know the exact consequences. In such situations, we need to make a decision under uncertainty.

Possible application areas. In many application areas, uncertainty is small – and can be made even smaller by appropriate measurements. For example, when we control a self-driving car, if there is an uncertainty about the locations and speeds of other objects on the road, we can install more accurate sensors and get a better description of the driving environment.

However, there are applications when it is difficult or even impossible to decrease uncertainty. One such area is anything related to human activities. Humans make individual decisions based on their perceived value of different alternatives. The same alternative – be it a movie or a computer – have drastically different value to different people, so it is

very difficult to predict their behavior. Such behavior affects economics and finance, so in decision making in economics and finance, it is very important to take such decision making under uncertainty into account.

Another area where it is difficult to decrease uncertainty is geosciences. In the cases of self-driving cars, we can decrease uncertainty by installing more accurate sensors, but to make a decision, e.g., on whether a certain underground area contains oil, there is only so much we can do with measurements. The only way to get a more accurate picture of what is going on beneath the earth surface is to dig a well and see – and the whole purpose of decision making is to decide whether such an expensive procedure is worth doing.

What we do in this thesis. In line with the above, in this thesis, we first, describe what is currently known about decision making under uncertainty. This will be discussed in Chapter 2.

In general, decision making under uncertainty requires that we fix some objective function that reflects the decision maker's interests. It is often difficult to find the expression that exactly captures these interests. In such cases, we have an approximate objective functions. It is well known that the quality of the resulting decisions and the computational complexity of computing this decision depend on the choice of approximating family. In Chapter 3, we formulate the problem of selecting an optimal approximating family as a precise mathematical problem, and we show that, under reasonable assumptions, the optimal approximating functions are polynomials.

Usually, decision makers do not just want to have a solution to their current problems, they also want to learn general techniques that will help them solve future problems. In particular, since it makes sense to consider polynomial objective functions, it is important to teach decision makers how to analyze such functions. A new way of doing this is presented in Chapter 4.

From the purely mathematical viewpoint, it is sufficient to formulate the practical problem in precise mathematical terms. However, from the practical viewpoint, even after this formulation, we still need to solve the corresponding computational problem. For simple

problems, the existing algorithms provide solutions in reasonable time. However, for more complex problems, the existing algorithms take too long a time. It is therefore necessary to explore the possibility of faster computations. One such possibility is the use of quantum computing. There are many existing algorithms for quantum computing, but, as we show in Chapter 5, they often do not provide an adequate representation of generic functions – and objective functions corresponding to decision making under uncertainty can be very complex. In this section, we show how to come up with more adequate representation of generic functions in quantum computing.

The practical usefulness of these theoretical developments is illustrated on the two above application examples: finances and geosciences.

We start with applications to finances. Of course, in finances, decision theory is already heavily used. What we are trying to do is illustrate the effectiveness of decision-making-under-uncertainty techniques in situations which are not well described by the usual decision theory approaches. Specifically, in Chapter 6, we explain why people sell and buy in the first place – it turns out that, from the usual economical viewpoint, it is still largely a mystery. In Chapter 7, we explain why for the same person, buy and sell prices are often different – another phenomenon that puzzles economists. These chapters describe the current decision making. If we take into account events in the past – or possible future consequences – then we need to describe how people perceive the corresponding time intervals. In Chapter 8, we explain that the seemingly counterintuitive empirical data about the human perception of time intervals can be naturally explained by the usual results about decision making under uncertainty.

Chapter 9 contains applications to geosciences. To illustrate the power of decision making techniques, we selected one of the most mysterious and challenging phenomenon – the Bhutan landscape anomaly.

Finally, in Chapter 10, we provide our future work plans.

Chapter 2

Decision Theory: A Brief Reminder

What is traditional decision theory. Traditional decision theory (see, e.g., [13, 22, 25, 30, 36]) describes preferences of rational agents, i.e., e.g., agents that when preferring A to B and B to C would always prefer A to C.

Comment. It is well known that real agents are not perfectly rational (see, e.g., [20, 24]), for the simple reason that our ability to process information and select an optimal decision is bounded. However, in many cases, traditional decision theory still provides a very good picture of human behavior.

The notion of utility. To describe the preferences of such an agent, we can select two alternatives:

- a very bad one A_- that is much worse than this agent will actually encounter, and
- a very good one A_+ that is much better than this agent will actually encounter.

For each value p from the interval $[0, 1]$, we can form a lottery $L(p)$ in which we get A_+ with probability p and A_- with the remaining probability $1 - p$.

When p is close to 1, this means that we are almost certainly getting a very good deal. So, for any realistic option A, the corresponding lottery $L(p)$ is better than A: $A < L(p)$. Similarly, when p is close to 0, this means that we are almost certainly getting a very bad deal, so $L(p) < A$. There should be a threshold u at which the preference $L(p) < A$ corresponding to smaller probabilities p is replaced by an opposite preference $A < L(p)$. In other words, we should have:

- $L(p) < A$ for all $p < u$ and

- $A < L(p)$ for all $p > u$.

This threshold value is called the *utility* of the alternative A; it is denoted by $u(A)$.

The above two conditions means that, in a certain reasonable sense, the original alternative A is equivalent to the lottery $L(u(A))$ corresponding to the probability $u(A)$: $A \equiv L(u(A))$.

A rational agent should maximize utility. Of course, the larger the probability of getting a very good outcome A_+, the better. Thus, among several lotteries $L(p)$, we should select the one for which the probability p of getting A_+ is the largest. Since each alternative A is equivalent to the corresponding lottery $L(u(A))$, this implies that we should select the alternative with the largest possible value of utility.

Main conclusion of traditional decision theory: a rational agent must maximize expected utility. In practice, we rarely know the consequences of each action. At best, we know possible outcomes $A_1, \ldots, A_n$, and their probabilities $p_1, \ldots, p_n$. Since each alternative A_i is equivalent to a lottery $L(u(A_i))$ in which we get A_+ with probability $u(A_i)$ and A_- with the remaining probability $1 - u(A_i)$, the whole action is equivalent to a two-stage lottery in which:

- first, we select one of the n alternatives A_i with probability p_i, and
- then, depending on which alternative A_i we selected on the first stage, we select A_+ with probability $u(A_i)$ and A_- with the remaining probability $1 - u(A_i)$.

As a result of this two-stage lottery, we get either A_+ or A_-. The probability u of getting A_+ can be computed by using the formula of complete probability, it is equal to

$$u = p_1 \cdot u(A_1) + \ldots + p_n \cdot u(A_n).$$

This is exactly the formula for the expected value of the utility $u(A_i)$. Thus, the utility of each action to a person is equal to the expected value of utility.

Thus, according to the traditional decision theory, rational agents should select the alternative with the largest possible value of expected utility.

Utility is defined modulo linear transformations. The numerical value of utility depends on the selection of the alternatives A_- and A_+. It can be shown that if we select a different pair (A'_-, A'_+), then the corresponding utility $u'(A)$ is related to the original utility by a linear transformation $u'(A) = a \cdot u(A) + b$ for some $a > 0$ and b; see, e.g., [22, 30].

How utility is related to money. The dependence of utility of money is non-linear: namely, utility u is proportional to the square root of the amount m of money $u = c \cdot \sqrt{m}$; see [20] and references therein.

Comment. This empirical fact can be explained. For example, the non-linear character of this dependence is explained, on a commonsense level, in [21], while the square root formula can also be explained – but it requires more mathematical analysis; see, e.g., [24]. In the current thesis, we simply take this fact as a given.

How to compare current and future gains: discounting. How can we compare current and future gains? If we have an amount m of money now, then we can place it in a bank and get the same amount plus interest, i.e., get the new amount $m' \stackrel{\text{def}}{=} (1+i) \cdot m$ in a year, where i is the interest rate. Thus, the amount m' in a year is equivalent to the value $m = q \cdot m'$ now, where $q \stackrel{\text{def}}{=} 1/(1+i)$. This reduction of future gains – to make them comparable to current gains – is known as *discounting.*

Discounting: a more detailed description. An event – e.g., a good dinner – a year in the past does not feel as pleasant to a person now as it may have felt a year ago. Similarly, a not-so-pleasant event in the past – e.g., a painful inoculation – does not feel as bad now if it felt a year ago, when it actually happened. Thus, the utility of an event changes with time: positive utility decreases, negative utility increases (i.e., gets closer to 0). If u is the utility of a current event, how can we describe the utility $f(u)$ of remembering the same event that happened 1 year ago?

We can normalize the utility values by assuming that the status quo situation has utility

0. Then the only remaining transformation is re-scaling $u' = a \cdot u$. Similarly to the case of interval uncertainty, it is reasonable to require that the function $f(u)$ is invariant with respect to such a transformation, i.e., that:

- if we have $v = f(u)$,
- then for each a, we should have $v' = f(u')$, where we denoted $v' = a \cdot v$ and $u' = a \cdot u$.

Substituting the expressions for v' and u' into the formula $v' = f(u')$, we conclude that $a \cdot v = f(a \cdot u)$, i.e., $a \cdot f(u) = f(a \cdot u)$. Substituting $u = 1$ into this formula, we conclude that $f(a) = q \cdot a$, where we denoted $q \stackrel{\text{def}}{=} f(1)$. Since $f(u) < u$ for $u > 0$, this would imply that $q < 1$.

So, an event with then-utility u that occurred 1 year ago has the utility $q \cdot u$ now. Similarly, an event with utility u that happened 2 years ago is equivalent to $q \cdot u$ a year ago, and thus, is equivalent to $q \cdot (q \cdot u) = q^2 \cdot u$ now. We can similarly conclude that an event with utility u that occurred t moments in the past is equivalent to utility $q^t \cdot u$ now.

Decision making under interval uncertainty. In real life, we rarely know the exact consequences of each action. As a result, for each alternative A, instead of the exact value of its utility, we often only know the bounds $\underline{u}(A)$ and $\overline{u}(A)$ on this unknown value. In other words, all we know is the interval $[\underline{u}(A), \overline{u}(A)]$ that contains the actual (unknown) value $u(A)$. How can we make a decision under this interval uncertainty?

In particular, for such an interval case, we need to be able to compare the interval-valued alternative with lotteries $L(p)$ for different values p. As a result of such comparison, we will come up with a utility of this interval. So, to make recommendations on decision under interval uncertainty, we need to be able to assign, to each interval $[\underline{u}, \overline{u}]$, a single utility value $u(\underline{u}, \overline{u})$ from this interval that describes this interval's utility.

Since utility is defined modulo a linear transformation $u \to u' = a \cdot u + b$, it is reasonable to require that the corresponding function $u(\underline{u}, \overline{u})$ should also be invariant under such transformations, i.e., that:

- if $u = u(\underline{u}, \overline{u})$,
- then $u' = u(\underline{u}', \overline{u}')$, where we denoted $u' = a \cdot u + b$, $\underline{u}' = a \cdot \underline{u} + b$, and $\overline{u}' = a \cdot \overline{u} + b$.

It turns out that this invariance requirement implies that

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

for some $\alpha_H \in [0, 1]$ [22, 30]. This formula was first proposed by a future Nobelist Leo Hurwicz and is, thus, known as the Hurwicz optimism-pessimism criterion [17, 25].

Theoretically, we can have values $\alpha_H = 0$ and $\alpha_H = 1$. However, in practice, such values do not happen:

- $\alpha_H = 1$ would correspond to a person who only takes into account the best possible outcome, completely ignoring the risk of possible worse situations;
- similarly, the value $\alpha_H = 0$ would correspond to a person who only takes into account the worst possible outcome, completely ignoring the possibility of better outcomes.

In real life, we thus always have $0 < \alpha_H < 1$.