The Use of Utility

Utility functions in a Bayesian Decision Analysis framework

Cecina Babich Morrow

COMPASS
Computational Statistics and Data Science
University of Bristol

7 December 2023

Overview

Decision theory

COMPASS

How can we make decisions under uncertainty?

University of BRISTOL

How do we make decisions in a mathematically formalized way while accounting for uncertainty?

Goal: Create a decision rule that is optimal given the information we have available

- Our decision rule will determine what decision we make given what we observe
- We will use the observations to infer an uncertain state of nature (Bayesian inference)
- We create our decision rule using some definition of ”optimal”, depending on our context-specific preferences

Notation

- Θ: space of all possible states of nature θ
- $\mathcal{X}$: space of observations
- $\mathcal{R}$: space of all possible rewards r
- $\mathcal{D}$: space of all possible decisions d
- $R(\theta, d) : \Theta \times \mathcal{D} \to \mathcal{R}$: reward function giving the reward for making decision d if the true state of nature is θ
  - Alternatively, $L(\theta, d) = -R(\theta, d)$: loss function

Utility function

$U(R(\theta, d)) : \mathcal{R} \to \mathbb{R}$: utility function mapping rewards to utility

But why use utility?

Life without utility

EMV Strategy

What if we just made decisions to maximize our expected reward?

Expected Monetary Value strategy

Select the decision d^* such that

$$d^* = \arg\max_{d} \sum_{\theta \in \Theta} R(\theta, d) p(\theta)$$
$$= \arg\max_{d} \bar{R}(d)$$

St. Petersburg Paradox

University of BRISTOL

The game:

Start with an amount of money S_1.

At each stage of the game $r \geq 1$, you can either take the money...

End with S_r

...or keep playing $\rightarrow$ flip a fair coin

Heads: your new stake is $4S_r$

Tails: lose everything

St. Petersburg Paradox

University of BRISTOL

The game:

Start with an amount of money S_1.

At each stage of the game $r \geq 1$, you can either take the money...

End with S_r

...or keep playing $\rightarrow$ flip a fair coin

Heads: your new stake is $4S_r$

Tails: lose everything

With an EMV strategy:

For each stage:

- $\bar{R}(\text{quit}) = S_r$
- $\bar{R}(\text{play another round}) = \frac{1}{2} \cdot 4S_r + \frac{1}{2} \cdot 0 = 2S_r$

So we should play indefinitely!

St. Petersburg Paradox

With an EMV strategy:
$P(\text{infinite number of heads}) = 0$, so
we will lose our money with probability 1.

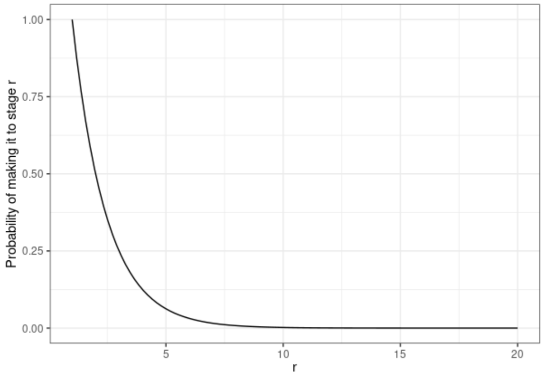

Figure: Probability of making it to stage *r*

Von Neumann-Morgenstern Utility Theorem

COMPASS

A little more notation

- $\mathcal{R}$: space of all possible rewards r
- $\mathcal{P}$: space of ***lotteries*** on $\mathcal{R}$
  - Probability distributions on $\mathcal{R}$
  - $\mathcal{P} = \{p : \mathcal{R} \to [0,1] | \sum_{r \in \mathcal{R}} p(r) = 1\}$
  - For a given $L \in \mathcal{P}$, $L = \sum_i p_i r_i$
- $\preceq$: representing preferences on $\mathcal{P}$

Example

$$\mathcal{R} = \{-£20, £0, £80\}$$

$$L = 0.8r_1 + 0.1r_2 + 0.1r_3 \quad M = 0.1r_1 + 0.3r_2 + 0.6r_3 \quad N = 0r_1 + 0r_2 + 1r_3$$

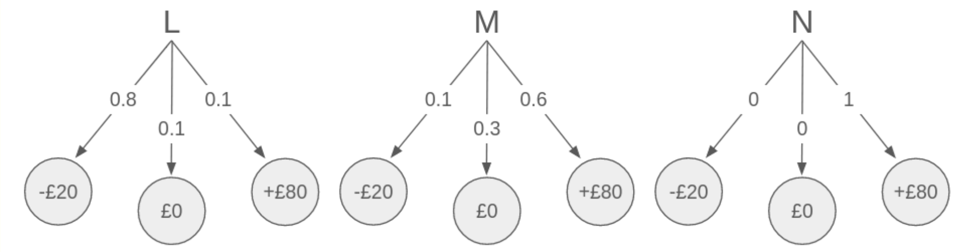

Figure: Three possible lotteries $L, M, N \in \mathcal{P}$

Example

We can also have ***mixtures*** of lotteries:

$$O = \alpha L + (1 - \alpha)N \in \mathcal{P}$$

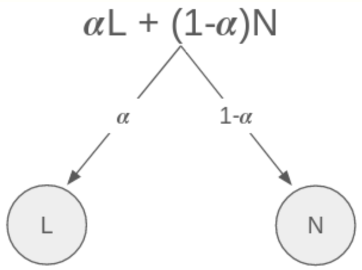

Figure: O is a mixture of L and N

Example

We can also have ***mixtures*** of lotteries:

$$O = \alpha L + (1 - \alpha) N \in \mathcal{P}$$

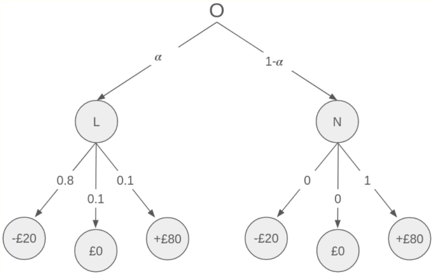

Figure: O is a mixture of L and N

Example

We can express our preferences between lotteries using $\preceq$:

$$N \succ M \succ L$$

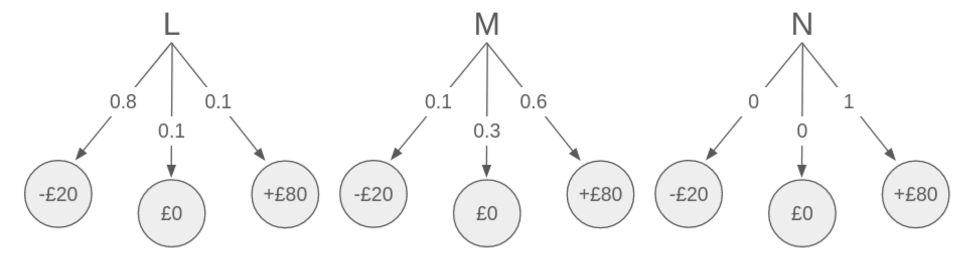

Figure: Three possible lotteries $L, M, N \in \mathcal{P}$

Von-Neumann Morgenstern Utility Theorem

Given a set of axioms of ”rational behavior” governing a decision-maker’s preferences between outcomes, their decisions will act to maximize the expected value of some utility function.

Theorem (Von Neumann-Morgenstern Utility Theorem)

There exists $U : \mathcal{R} \to \mathbb{R}$ *such that for all* $L, M \in \mathcal{P}$,

$$\begin{aligned} L \succ M &\iff \mathbb{E}^L[U(r)] > \mathbb{E}^M[U(r)] \\ &\iff \sum_i \ell_i U(r_i) > \sum_i m_i U(r_i) \end{aligned}$$

Axioms

- **Axiom 1 – Completeness**: $\forall r_1, r_2 \in \mathcal{R}$, $r_1 \succeq r_2$ or $r_2 \succeq r_1$.
- **Axiom 2 – Transitivity**: If $r_1 \succeq r_2, r_2 \succeq r_3$, then $r_1 \succeq r_3$.
- **Axiom 3 – Continuity**: For $L, M, N \in \mathcal{P}$ such that $N \succeq M \succeq L$, there exists $p \in [0, 1]$ such that $pL + (1 - p)N \sim M$.
- **Axiom 4 – Independence**: For all $P, Q, R \in \mathcal{P}$, $a \in (0, 1]$, $P \succ Q \implies aP + (1 - a)R \succ aQ + (1 - a)R$.

Lemma

University of BRISTOL

Lemma

If $L \succ M$ and $0 \leq a < b \leq 1$, then $bL + (1-b)M \succ aL + (1-a)M$.

Intuition: we would rather have a higher probability of playing the lottery we prefer.

Proof sketch:

- Let $a = 0$...
- Let $a > 0$, $N = bL + (1-b)M \sim \frac{a}{b}N + (1 - \frac{a}{b})N$...

Proof of Utility Theorem

University of BRISTOL

Theorem (Utility Theorem)

$L \succ M \iff \sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$

Finite case: Assume there are n rewards $A_1, ..., A_n \in \mathcal{R}$ such that $A_n \succeq A_{n-1} \succeq ... \succeq A_1$.
(Assume $A_n \succ A_1$, or this won't be very interesting)
Defining U:

- Define $U(A_1) := 0$, $U(A_n) := 1$
- By Axiom 3 (continuity), $\forall\, A_i\, \exists\, q_i$ such that $A_i \sim q_i An + (1 - q_i)A_1$:
  - $U(A_i) := q_i$

Proof of Utility Theorem

University of BRISTOL

Theorem (Utility Theorem)

$L \succ M \iff \sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$

($\Longleftarrow$) Assume $\sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$.

$$L = \sum_i \ell_i A_i \sim L' := \sum_i \ell_i [q_i A_n + (1 - q_i) A_1]$$

$$= \left[\sum_i \ell_i U(A_i)\right] A_n + \left[\sum_i \ell_i (1 - U(A_i))\right] A_1$$

Similarly,

$$M \sim M' := \left[\sum_i m_i U(A_i)\right] A_n + \left[\sum_i m_i (1 - U(A_i))\right] A_1$$

Proof of Utility Theorem

University of BRISTOL

Theorem (Utility Theorem)

$L \succ M \iff \sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$

($\Longleftarrow$) Assume $\sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$.
Since $A_n \succ A_1$ and $\sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$, by our lemma:

$$\left[\sum_i \ell_i U(A_i)\right] A_n + \left[\sum_i \ell_i (1 - U(A_i))\right] A_1 \succ \left[\sum_i m_i U(A_i)\right] A_n + \left[\sum_i m_i (1 - U(A_i))\right] A_1$$

$$L \sim L' \succ M' \sim M$$

Proof of Utility Theorem

University of BRISTOL

Theorem (Utility Theorem)

$L \succ M \iff \sum_i \ell_i U(r_i) > \sum_i m_i U(r_i)$

($\implies$) Assume $L \succ M$.

$$L \sim L' \succ M' \sim M$$

$$\left[\sum_i \ell_i U(A_i)\right] A_n + \left[\sum_i \ell_i (1 - U(A_i))\right] A_1 \succ \left[\sum_i m_i U(A_i)\right] A_n + \left[\sum_i m_i (1 - U(A_i))\right] A_1$$

Proof sketch:

Proof by contrapositive: assume $\sum_i \ell_i U(r_i) \leq \sum_i m_i U(r_i) \implies$ contradiction!

Utilizing a utility function

Bayes decision

Instead of maximizing the expected reward directly, we can transform $R(\theta, d)$ into $U(R(\theta, d))$.

Bayes decision under utility U

Select the decision d^* such that

$$
\begin{aligned}
d^* &= \arg\max_d \sum_{\theta \in \Theta} U[R(\theta, d)]p(\theta) \\
&= \arg\max_d \bar{U}(d)
\end{aligned}
$$

A return to St. Petersburg

University of BRISTOL

Recap of the game:
Start with an amount of money S_1.
At each stage of the game $r \geq 1$, you can either take the money and leave with S_r
...or keep playing $\rightarrow$ flip a fair coin

Heads: your new stake is $4S_r$

Tails: lose everything

How can we come up with a strategy other than gambling forever?

Potential utility function

$$U(R(d,\theta)) := \frac{R(\theta,d)}{\delta + R(\theta,d)}$$

A return to St. Petersburg

University of BRISTOL

Utility function:

$$U(R(d,\theta)) := \frac{R(\theta,d)}{\delta + R(\theta,d)}$$

Let $S_1 = £1, \delta = 4$:

At stage 1:

$\bar{U}(\text{quitting}) = 1 \cdot \frac{S_1}{\delta+S_1} = \frac{1}{4+1} = \frac{1}{5}$

$\bar{U}(\text{playing another round}) = \frac{1}{2} \cdot \frac{4S_1}{\delta+4S_1} + \frac{1}{2} \cdot 0 = \frac{1}{2} \cdot \frac{4}{4+4} = \frac{1}{4}$

So $d^* =$ Play another round

A return to St. Petersburg

At stage 1:

$$\bar{U}(\text{quit}) = 1 \cdot \frac{S_1}{\delta + S_1} = \frac{1}{4+1} = \frac{1}{5}$$

$$\bar{U}(\text{play another round}) = \frac{1}{2} \cdot \frac{4S_1}{\delta + 4S_1} + \frac{1}{2} \cdot 0 = \frac{1}{2} \cdot \frac{4}{4+4} = \frac{1}{4}$$

So $d^* =$ Play another round

If we got heads, then $S_2 = £4$

At stage 2:

$$\bar{U}(\text{quit}) = 1 \cdot \frac{S_2}{\delta + S_2} = \frac{4}{4+4} = \frac{1}{2}$$

$$\bar{U}(\text{play another round}) = \frac{1}{2} \cdot \frac{4S_2}{\delta + 4S_2} + \frac{1}{2} \cdot 0 = \frac{1}{2} \cdot \frac{16}{4+16} = \frac{2}{5}$$

So $d^* =$ Quit $\rightarrow$ No more playing indefinitely!

A return to St. Petersburg

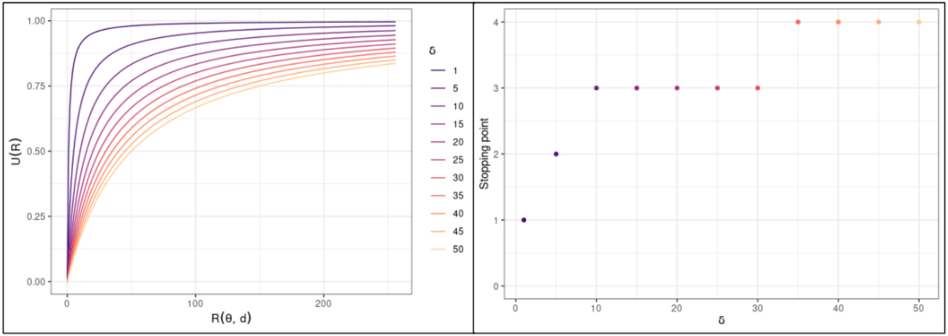

Figure: Utility functions and their corresponding stopping points for different values of δ

Conclusions

Conclusions

University of BRISTOL

- Decision theory formalizes the process of decision making under uncertainty
- Acting to maximize our expected reward can lead to some suboptimal decisions
- If our preferences follow certain axioms of rationality, we can represent them using a utility function
- The shape of our utility function represents our relationship to risk

References

Jim Smith (2010)
Bayesian Decision Analysis: Principles and Practice
Cambridge University Press.

Christian Robert (2007)
The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation
Springer.

Y. Narahari (2012)
Game Theory: Lecture Notes
Indian Institute of Science

Eric Pacuit (2013)
Proof of von Neumann Morgenstern Representation Theorem: Part 1-3
YouTube

Questions?

Cecina Babich Morrow

COMPASS
Computational Statistics and Data Science
University of Bristol

7 December 2023