Bayesian decision theory
A simple toy problem

Van Erp, H. R N; Linger, R. O.; Van Gelder, P. H A J M

DOI
10.1063/1.4959058

Publication date
2016

Document Version
Final published version

Published in
Proceedings of Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 35th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, MaxEnt 2015

Citation (APA)
Van Erp, H. R. N., Linger, R. O., & Van Gelder, P. H. A. J. M. (2016). Bayesian decision theory: A simple toy problem. In *Proceedings of Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 35th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, MaxEnt 2015* (Vol. 1757). Article 050002 American Institute of Physics. https://doi.org/10.1063/1.4959058

Bayesian Decision Theory: A Simple Toy Problem

H.R.N. van Erp[1,a)], R.O. Linger[1] and P.H.A.J.M. van Gelder[1,2]

[1]*Safety and Security Science Group, TU Delft, The Netherlands.*
[2]*Director of the TU Delft Safety and Security Institute.*

[a)]Corresponding author: h.r.n.vanerp@tudelft.nl

Abstract. We give here a comparison of the expected outcome theory, the expected utility theory, and the Bayesian decision theory, by way of a simple numerical toy problem in which we look at the investment willingness to avert a high impact low probability event. It will be found that for this toy problem the modeled investment willingness under the Bayesian decision theory is minimally three times higher compared to the investment willingness under either the expected outcome or the expected utility theories, where it is noted that the estimates of the latter two theories seem to be unrealistically low.

INTRODUCTION

The Bayesian decision theory is very simple in structure. Its algorithmic steps are the following:

1. Use the product and sum rules of Bayesian probability theory to construct outcome probability distributions.
2. If our outcomes are monetary in nature, then by way of the Bernoulli utility function we may map utilities to the monetary outcomes of our outcome probability distributions.
3. Maximize a scalar multiple of the sum of lower bound, expectation value, and upper bound of the resulting utility probability distributions.

We will give in this paper a comparison of the expected outcome theory, the expected utility theory, and the Bayesian decision theory, by way of a simple toy-problem in which we look at the investment willingness to avert a high impact low probability event.

A SIMPLE SCENARIO

We now apply our Bayesian framework to a scenario in which a decision maker must decide on how it is willing to invest in a further improvement of its flood defenses. The two decisions under consideration in our simple scenario are

$$D_1 = \text{keep status quo},$$
$$D_2 = \text{improve flood defenses}.$$

The possible outcomes in our risk scenario remain the same under either decision, and as such are not dependent upon the particular decision taken. These outcomes for a given year are

$$O_1 = \text{flooding},$$
$$O_2 = \text{no flooding}.$$

The hypothetical damages associated with these outcomes are, respectively,

$$C_1 = x \text{ euros},$$
$$C_2 = 0 \text{ euro}, \tag{1}$$

Bayesian Inference and Maximum Entropy Methods in Science and Engineering
AIP Conf. Proc. 1757, 050002-1–050002-7; doi: 10.1063/1.4959058
Published by AIP Publishing. 978-0-7354-1415-0/$30.00

and the investment costs associated with the additional flood defenses are expressed by the parameter

$$I = \text{investment costs}. \tag{2}$$

Note that if we were to do an actual cost-benefit analysis, rather than a demonstration of the here proposed decision theoretical framework, then the cost of money itself, in the form of a potential loss of interest on the investment I and the outcomes C_i, should also be taken into account.

The decision whether to improve the flood defenses or not is of influence on the probabilities of the respective outcomes. Under the decision to make no additional investments in flood defenses and keep the status quo, D_1, the probabilities of the outcomes will be, say,

$$\begin{aligned} P(O_1|D_1) &= \theta, \\ P(O_2|D_1) &= 1-\theta. \end{aligned} \tag{3}$$

Under the decision to improve the flood defenses, D_2, the probabilities of the flood outcomes will be decreased, leaving us with hypothetical outcome probabilities, say,

$$\begin{aligned} P(O_1|D_2) &= \phi, \\ P(O_2|D_2) &= 1-\phi, \end{aligned} \tag{4}$$

where $\phi < \theta$; that is, the proposed flood defenses will reduce the chances of a flooding by a factor $c = \theta/\phi$, where $c > 1$.

In what follows we will give the solution of this problem of choice by way of the expected outcome theory, the Bayesian decision theory without utility transformations, expected utility theory, and the Bayesian decision theory with utility transformations.

THE EXPECTED OUTCOME SOLUTION

The notion of 'expectation of profit' was very intuitive to the first workers in probability theory. It seemed obvious to many that a person acting in pure self-interest should always behave so as to maximize his expected profit. The prosperous merchants in 17th century Amsterdam bought and sold mathematical expectations as if they were tangible goods [1].

We may combine (1), (2), (3), and (4) to construct the outcome probability distributions under the decisions D_1 and D_2:

$$p(C_i|D_1) = \begin{cases} \theta, & C_1 = -x, \\ 1-\theta, & C_2 = 0, \end{cases} \tag{5}$$

and

$$p(C_i|I, D_2) = \begin{cases} \phi, & C_1 = -x-I, \\ 1-\phi, & C_2 = -I, \end{cases} \tag{6}$$

where we explicitly conditionalize on the investment parameter I, which is to be to estimated.

The expected outcomes of these probability distributions are, respectively [2],

$$E(C|D_1) = -\theta\, x \tag{7}$$

and

$$E(C|I, D_2) = -\phi\, x - I. \tag{8}$$

The decision theoretical equality

$$E(C|D_1) = E(C|I, D_2) \tag{9}$$

represents the equilibrium situation, where we will be undecided between the decision to keep the status quo D_1 and the decision to invest in additional flood defenses. Now, if we solve for the unknown I in (9), by way of (7) and (8):

$$I = (\theta - \phi)\, x, \tag{10}$$

then we find that investment where we will be undecided between both decisions.

Stated differently, any investment smaller than (10) will turn (9) into an inequality, where D_2 becomes more attractive than D_1. If we assume that we are only motivated by monetary costs, then the equilibrium investment (10) is the maximal investment we will be willing to make to improve our flood defenses.

THE BAYESIAN DECISION THEORY SOLUTION WITHOUT UTILITY TRANSFORMATIONS

In the Bayesian decision theory the mean of the lower confidence bound, expectation value, and upper confidence bound is taken as the position measure of the underlying outcome probability distribution which is to be maximized [3]:

$$R(D_i) = \frac{LB(C|D_i) + E(C|D_i) + UB(C|D_i)}{3}, \tag{11}$$

where the k-sigma lower confidence bound is corrected for undershoot of the worst possible outcome $a = \min(C_i)$, giving

$$LB(C|D_i) = \begin{cases} a, & E(C|D_i) - k \operatorname{std}(C|D_i) < a, \\ E(C|D_i) - k \operatorname{std}(C|D_i), & E(C|D_i) - k \operatorname{std}(C|D_i) \geq a, \end{cases} \tag{12}$$

and the k-sigma upper confidence bound is corrected for overshoot of the best possible outcome $b = \max(C_i)$, giving

$$UB(C|D_i) = \begin{cases} E(C|D_i) + k \operatorname{std}(C|D_i), & E(C|D_i) + k \operatorname{std}(C|D_i) \leq b, \\ b, & E(C|D_i) + k \operatorname{std}(C|D_i) > b. \end{cases} \tag{13}$$

Substituting (12) and (13) into (11), we obtain the k-sigma risk index:

$$R(D_i) = \begin{cases} E(C|D_i), & \text{neither undershoot nor overshoot,} \\ \frac{a+2E(C|D_i)+k \operatorname{std}(C|D_i)}{3}, & \text{undershoot and no overshoot,} \\ \frac{2E(C|D_i)-k \operatorname{std}(C|D_i)+b}{3}, & \text{overshoot and no undershoot,} \\ \frac{a+E(C|D_i)+b}{3}, & \text{both undershoot and overshoot,} \end{cases} \tag{14}$$

we note that the first row of (14) corresponds with the expected outcome theory criterion of choice [1].

In the toy problem under consideration we have a high impact low probability scenario; that is, both large monetary costs and small probabilities for the high-impact event, or, equivalently, $x \gg 0$ and $\theta, \phi << 0.5$. Stated differently, the outcome probability distributions (5) and (6) under consideration will both be highly skewed to the left and, as a consequence, will lead to the third condition in (14): "(upper confidence bound) overshoot and no (lower confidence bound) undershoot."

It follows that the operating criterion of choice will be (14)

$$R(D_i) = \frac{2E(C|D_i) - k \operatorname{std}(C|D_i) + b}{3}. \tag{15}$$

The best possible outcome under decision D_1 is (5)

$$b = \max(-x, 0) = 0, \tag{16}$$

and the standard deviation of (5) is [2]

$$\operatorname{std}(C|D_1) = \sqrt{\theta(1-\theta)}x. \tag{17}$$

So using (7), (15), (16), and (17), the risk index under the decision to keep the status quo is:

$$R(D_1) = -\frac{x\left(2\theta + k\sqrt{\theta(1-\theta)}\right)}{3}. \tag{18}$$

The best possible outcome under decision D_2 is (6)

$$b = \max(-x - I, -I) = -I, \tag{19}$$

and the standard deviation of (6) is [2]

$$\operatorname{std}(C|I, D_2) = \sqrt{\phi(1-\phi)}x. \tag{20}$$

So using (8), (15), (19), and (20), the risk index under the decision invest in additional flood defenses is

$$R(I, D_2) = -\frac{x\left(2\phi + k\sqrt{\phi(1-\phi)}\right)}{3} - I. \tag{21}$$

The decision theoretical equality

$$R(D_1) = R(I, D_2) \tag{22}$$

represents the equilibrium situation, where we will be undecided between the decision to keep the status quo D_1 and the decision to invest in additional flood defenses. Now, if we solve for the unknown I in (22), by way of (18) and (21), we get

$$I = \frac{1}{3}\left[\left(2\theta + k\sqrt{\theta(1-\theta)}\right) - \left(2\phi + k\sqrt{\phi(1-\phi)}\right)\right] x, \tag{23}$$

which is the investment where we will be undecided between both decisions.

Stated differently, any investment smaller than (23) will turn (22) into an inequality, where D_2 becomes more attractive than D_1. It follows that the equilibrium investment (23) is also the maximal investment we will be willing to make to improve our flood defenses.

THE EXPECTED UTILITY SOLUTION

The utility of a given outcome is the perceived worth of that outcome. If we take the utilities that monetary outcomes hold for us to be an incentive for our decisions, then we may perceive money to be a stimulus.

For the rich man hundred one hundred euros is an insignificant amount of money. So, the prospect of gaining or losing hundred euros will fail to move the rich man; that is, an increment of hundred euros for him has a utility which tends to zero. For the poor man one hundred euros will be a significant amount of money. So, the prospect of gaining or losing one hundred euros will most likely move the poor man to action. It follows that an increment of one hundred euros for him has a utility significantly greater than zero.

Bernoulli in 1738 derived his utility function for the subjective value of objective monies by way of a variance argument, in which he considered the subjective effect of a given fixed monetary increment c for two persons holding different initial wealths. Based on this variance argument he derived the utility function of going from an initial asset position x to the asset position $x + c$:

$$u(x, x+c) = q \log \frac{x+c}{x} \tag{24}$$

where q is some scaling constant greater than zero [4]. An alternative consistency argument for the derivation of Bernoulli's utility function may be found in [3].

In expected utility theory the expectation values of the utility probability distributions are maximized. Assuming that the decision maker has a total wealth, that is, an actual income and asset portfolio, of

$$M = m \text{ euros}, \tag{25}$$

then, using (24), we may construct from (5) and (6) the utility probability distributions under the decisions D_1 and D_2 as

$$p(U_i | D_1) = \begin{cases} \theta, & U_1 = q \log \frac{m-x}{m}, \\ 1-\theta, & U_2 = q \log \frac{m}{m}, \end{cases} \tag{26}$$

and

$$p(U_i | I, D_2) = \begin{cases} \phi, & U_1 = q \log \frac{m-x-I}{m}, \\ 1-\phi, & U_2 = q \log \frac{m-I}{m}. \end{cases} \tag{27}$$

The expected outcomes of the utility probability distributions are, respectively [2],

$$E(U | D_1) = q\left(\theta \ \log \frac{m-x}{m}\right) \tag{28}$$

and

$$E(U | I, D_2) = q\left(\phi \ \log \frac{m-x-I}{m-I} + \log \frac{m-I}{m}\right). \tag{29}$$

The decision theoretical equality

$$E(U|D_1) = E(U|I, D_2) \tag{30}$$

represents the equilibrium situation, where we will be undecided between decisions D_1, keep the status quo, and D_2, invest in additional flood defenses. Now, if we substitute (28) and (29) into (30), then we obtain the closed expression for that investment value where we will be undecided between both decisions:

$$\log \frac{m - I}{m} = \theta \log \frac{m - x}{m} - \phi \log \frac{m - x - I}{m - I}. \tag{31}$$

Any investment smaller than the numerical solution of I in (31) will turn (31) into an inequality, where D_2 becomes more attractive than D_1. It follows that the investment equilibrium solution of (31) is also the maximal investment we will be willing to make to improve our flood defenses.

THE BAYESIAN DECISION THEORY SOLUTION WITH UTILITY TRANSFORMATIONS

Because of the left skewness of the utility probability distributions (26) and (27), the third condition in (14) remains to be the operating criterion of choice in the Bayesian decision theory with utility transformations:

$$R(D_i) = \frac{2E(U|D_i) - k \operatorname{std}(U|D_i) + b}{3}. \tag{32}$$

The best possible outcome under decision D_1 is (26)

$$b = \max\left(q \log \frac{m - x}{m}, q \log \frac{m}{m}\right) = q \log \frac{m}{m} = 0, \tag{33}$$

and the standard deviation of (26) is [2]

$$\operatorname{std}(U|D_1) = -q \sqrt{\theta(1-\theta)} \log \frac{m - x}{m}. \tag{34}$$

So from (28), (32), (33), and (34), the risk index under the decision to keep the status quo is

$$R(D_1) = \frac{q \log \frac{m-x}{m} \left[2\theta + k \sqrt{\theta(1-\theta)}\right]}{3}. \tag{35}$$

The best possible outcome under decision D_2 is (27)

$$b = \max\left(q \log \frac{m - x - I}{m}, q \log \frac{m - I}{m}\right) = q \log \frac{m - I}{m}, \tag{36}$$

and the standard deviation of (27) is [2]:

$$\operatorname{std}(U|I, D_2) = -q \sqrt{\phi(1-\phi)} \log \frac{m - x - I}{m - I}. \tag{37}$$

So from (29), (32), (36), and (37), the risk index under the decision invest in additional flood defenses is

$$R(I, D_2) = \frac{q \log \frac{m-x-I}{m-I} \left[2\phi + k \sqrt{\phi(1-\phi)}\right]}{3} + q \log \frac{m - I}{m}. \tag{38}$$

The decision theoretical equality

$$R(D_1) = R(I, D_2) \tag{39}$$

represents the equilibrium situation, where we will be undecided between decisions D_1, keep the status quo, and D_2, invest in additional flood defenses. Now, if we substitute (37) and (38) into (39), then we obtain the closed expression for that investment value where we will be undecided between both decisions:

$$\log \frac{m - I}{m} = \frac{1}{3}\left[\left(2\theta + \sqrt{\theta(1-\theta)}\right) \log \frac{m - x}{m} - \left(2\phi + \sqrt{\phi(1-\phi)}\right) \log \frac{m - x - I}{m - I}\right]. \tag{40}$$

Any investment smaller than the numerical solution of I in (40) will turn (39) into an inequality, where D_2 becomes more attractive than D_1. It follows that the equilibrium investment (40) is also the maximal investment we will be willing to make to improve our flood defenses.

Note that the "Weber-constant," q, has fallen away in both the decision theoretical equalities (31) and (40). This will hold in general, as both the expectation values and standard deviations of the utility probability distributions (26) and (27) are linear in the unknown constant q. It follows that we may always set, without any loss of generality, q to one.

SOME NUMERICAL RESULTS

In our simple toy problem we have a decision maker who must decide on how much he is willing to invest in a further improvement of his flood defenses.

After the great Dutch flooding the 'Oosterschelde Waterkering' was built. This was a movable dike that allowed for an improved safety from $\theta = 1/100$ to $\phi = 1/4000$, while keeping the Oosterschelde connected to the North Sea. This open connection to the North Sea was decided upon in order to keep the salt-sea ecological system of the Oosterschelde lake intact.

The total costs of the Oosterschelde Waterkering were about 2.5 billion euros. The bulk of these costs were due to the movable character of this dike. Had the Dutch government decided to build an immovable dike, then the costs would only have been about 175 million euros.

The total value of the assets at risk were about 1/20th of the GDP at the time, so that in (1),

$$x = 3.75 \times 10^9 \text{ euros.} \tag{41}$$

The wealth of the decision maker, that is, the Dutch government, was about 40% of the Dutch GDP at the time. Aggregated over a period of five years to account for the building time of the movable Oosterschelde dike, the relevant wealth was

$$m = 1.5 \times 10^{11} \text{ euros.} \tag{42}$$

Right after the great flood the probability in (3) of a (catastrophic) flood had been estimated to be

$$\theta = \frac{1}{100}, \tag{43}$$

whereas the probability in (4) of a (catastrophic) flood under the improved flood defenses had been estimated as

$$\phi = \frac{1}{4000}. \tag{44}$$

Substituting the values (41) through (44) into (10), (23), (31), and (40), we obtain the solutions for the maximal investments I:

- Expected outcome theory:
  - Any sigma level: $I = 36.6 \times 10^6$ euros

- BDT without utility transformation:
  - 1-sigma level: $I = 129.0 \times 10^6$ euros
  - 2-sigma level: $I = 233.6 \times 10^6$ euros
  - 3-sigma level: $I = 338.2 \times 10^6$ euros

- Expected utility theory:
  - Any sigma level: $I = 37.0 \times 10^6$ euros

- BDT with utility transformation:
  - 1-sigma level: $I = 129.8 \times 10^6$ euros
  - 2-sigma level: $I = 234.9 \times 10^6$ euros
  - 3-sigma level: $I = 340.1 \times 10^6$ euros.

We note here that after the great Dutch flood the discussion was not whether to build additional flood defenses, but, rather, whether to choose for the expensive solution which would keep the Oosterschelde salt-sea ecosystem intact over the 'cheap' solution which would not. Under the expected utility theory solution the cheap solution of an immovable dike would have been too expensive by a factor of three, whereas under the Bayesian decision theory solution with utility transformation the cheap solution was well within the 2-sigma bounds.

We also note that the actual project was justified under neither one of the solutions. This is because the (in)tangible costs of losing the Oosterschelde salt-sea ecosystem and the (in)tangible benefits of human safety were not factored explicitly into this particular decision analysis. But the very fact that the Dutch chose to invest 2.5 billion euros in a movable Oosterschelde Waterkering, rather than opt for the cheap immovable dike solution of 175 million euros, is an important data point which shows that these additional (in)tangibles must have played an important role in the actual decision making process.

DISCUSSION

We give here a comparison of the expected outcome theory, the expected utility theory, and the Bayesian decision theory, by way of a simple toy problem in which we look at the investment willingness to avert a high impact low probability event. We have demonstrated here that the adjusted criterion of choice, in which the mean of the sum of the undershoot corrected lower confidence bound, expectation value, and overshoot corrected upper confidence bound of either outcome or utility probability distributions are maximized, though mathematically trivial [3, 5], has non-trivial practical implications for the modeled investment willingness for high impact low probability events. For it is found that under the alternative criterion of choice of the Bayesian decision theory the investment willingness for such events may increase easily with a factor three or more.

ACKNOWLEDGMENTS

This work has received funding from the European Unions Seventh Programme for research, technological development and demonstration under grant agreements No's 265138, 603960 and 608166.

REFERENCES

[1] Jaynes E.T.: *Probability Theory; the Logic of Science*. Cambridge University Press, (2003).
[2] Lindgren B.W.: *Statistical Theory*, Chapman & Hall, Inc., New York, (1993).
[3] Erp van H.R.N., Linger R.O., and Gelder van P.H.A.J.M.: *Fact Sheet on the Bayesian Decision Theory*, arXiv, (2015).
[4] Bernoulli D.: *Exposition of a New Theory on the Measurement of Risk*. Translated from Latin into English by Dr Louise Sommer from 'Specimen Theoriae Novae de Mensura Sortis', Commentarii Academiae Scientiarum Imperialis Petropolitanas, Tomus V, 175-192, (1738).
[5] Erp van H.R.N., Linger R.O., and Gelder van P.H.A.J.M.: *An Outline of the Bayesian Decision Theory*, Bayesian Inference and Maximum Entropy Methods in Science and Engineering-35th International Workshop, Potsdam, New York, (2016).