Decision Making under Uncertainty

Part 2: Subjective probability and utility

Christos Dimitrakakis

April 4, 2014

1 Subjective probability

In order to make decisions, we need to be able to make predictions about the possible outcomes of each decision. Usually, we have *uncertainty* about what those outcomes are. This can be due to *stochasticity*, which is frequently used to model games of chance and inherently unpredictable physical phenomena. It can also be due to *partial information*, a characteristic of many natural problems. For example, it might be hard to guess at any one moment how much change you have in your wallet, whether you will be able to catch the next bus, or to remember where you left your keys.

In either case, this uncertainty can be expressed as a *subjective belief*. This does not have to correspond to reality. For example, some people believe, quite inaccurately, that if a coin comes up tails for a long time, it is quite likely to come up heads very soon. Or, you might quite happily believe your keys are in your pocket, only to realise that you left them at home as soon you arrive at the office.

In this book, we assume the view that subjective beliefs can be modelled as *probabilities*. This allows us to treat uncertainty due to stochasticity and due to partial information in a unified framework. In doing so, we shall treat each part of the problem as specifying a space of possible outcomes. What we wish to do is to find a *consistent way* of defining probabilities in the space of outcomes.

1.1 Relative likelihood

Let us consider the simple example of guessing whether a tossed coin will come up head, or tails. Let S be the sample space, and let $A \subset S$ be the set of tosses where the coin comes up heads, and $B \subset S$ be the set of tosses where it comes up tails. Here $A \cap B = \emptyset$, but there may be some other events such as the coin becoming lost, so it does not necessarily hold that $A \cup B = S$. Nevertheless, we only care about whether A is more likely to occur than B. We can express that via the concept of relative likelihood:

Definition 1.1. (The relative likelihood of two events A and B)

- If one thinks that A is more likely than B, then we write $A \succ B$, or equivalently $B \prec A$.
- If one thinks A is as likely as B, then we write $A \equiv B$.

We also use \succeq and \preceq for at least as likely as and for no more likely than.

Let us now speak more generally about the case where we have defined an appropriate σ -field \mathcal{F} on \mathcal{S} . Then each element $A_i \in \mathcal{F}$ will be a subset of \mathcal{S} . Furthermore, we have defined a relative likelihood relation for all elements $A_i \in \mathcal{F}$.¹

As we would like to use the language of probability to talk about likelihoods, we would like to be able to define a probability measure that agrees with our given relations.

A probability measure $P : \mathcal{F} \to [0, 1]$ is said to *agree* with a relation $A \preceq B$, if it has the property that: $P(A) \leq P(B)$ if and only if $A \preceq B$, for all $A, B \in \mathcal{F}$.

Of course, there are many possible measures that can agree with a given relation. It could even be that a given relational structure is incompatible with any possible probability measure. For that reason, we shall have to make some assumptions about relative likelihoods of events.

1.2 Subjective probability assumptions

Our beliefs must be *consistent*. This can be achieved if they satisfy some assumptions. First of all, it must always be possible to say whether one event is more likely than the other. Consequently, we are not allowed to claim ignorance.

Assumption 1.1 (SP1). For any pair of events $A, B \in \mathcal{F}$, one of the following must hold: Either $A \succ B$, $A \prec B$, or $A \eqsim B$.

If we can partition A, B in such a way that each part of A is less likely than its counterpart in B, then A is less likely than B. For example, suppose that A_1 the event that it rains non-stop between 14:00 and 15:00 tomorrow, A_2 the event that it rains intermittently. Let B_1, B_2 be the corresponding events for cloudy weather, and , but no rain during the next hour. If we think that it is more likely to rain constantly than to be dry and cloudy, and that it is more likely to rain intermittently than to be sunny, then we must conclude that it is more likely to rain than not.

Assumption 1.2 (SP2). Let $A = A_1 \cup A_2$, $B = B_1 \cup B_2$ with $A_1 \cap A_2 = B_1 \cap B_2 = \emptyset$. If $A_i \preceq B_i$ for i = 1, 2 then $A \preceq B$.

We also require the simple technical assumption that any event $A \in \mathcal{F}$ is at least as likely as the empty event \emptyset , which never happens.

Assumption 1.3 (SP3). If S is the certain event, and \emptyset never occurs, then: $\emptyset \preceq A$ and $\emptyset \prec S$.

As it turns out, these assumptions are sufficient for proving the following theorems [1]. The first theorem tells us that our belief must be consistent with respect to transitivity.

Theorem 1.1 (Transitivity). For all events A, B, D, if $A \preceq B$ and $B \preceq D$, then $A \preceq D$.

¹More formally, we can define three classes: $C_{\succ}, C_{\prec}, C_{\eqsim} \subset \mathcal{F}^2$ such that a pair $(A_i, A_j) \in C_R$ if an only if it satisfies the relation $A_i R A_j$, where $R \in \{\succ, \prec, \eqsim\}$. It is easy to see that the three classes form a partition of \mathcal{F}^2 .

The second theorem says that if two events have a certain relation, then their negations have the converse relation.

Theorem 1.2 (Complement). For any $A, B: A \preceq B$ iff $A^{\complement} \succ B^{\complement}$.

Finally, note that if $A \subset B$, then it must be the case that whenever A happens, B must happen and hence B must be at least as likely as A. This is demonstrated in the following theorem.

Theorem 1.3 (Fundamental property of relative likelihoods). If $A \subset B$ then $A \preceq B$. Furthermore, $\emptyset \preceq A \preceq S$ for any event A.

Since we are dealing with σ -fields, we need to introduce properties for infinite sequences of events. While these are not necessary if the field \mathcal{F} is finite, it is good to include them for generality.

Assumption 1.4 (SP4). If $A_1 \supset A_2 \supset \cdots$ is a decreasing sequence of events in \mathcal{F} and $B \in \mathcal{F}$ is such that $A_i \succeq B$ for all i, then $\bigcap_{i=1}^{\infty} A_i \succeq B$.

As a consequence, we obtain the following dual theorem:

Theorem 1.4. If $A_1 \subset A_2 \subset \cdots$ is an increasing sequence of events in \mathcal{F} and $B \in \mathcal{F}$ is such that $A_i \preceq B$ for all i, then $\bigcup_{i=1}^{\infty} A_i \preceq B$.

We are now able to state a theorem for the unions of infinite sequences of disjoint events.

Theorem 1.5. If $(A_i)_{i=1}^{\infty}$ and $(B_i)_{i=1}^{\infty}$ are infinite sequences of disjoint events in \mathcal{F} such that $A_i \preceq B_i$ for all i, then $\bigcup_{i=1}^{\infty} A_i \preceq \bigcup_{i=1}^{\infty} B_i$.

Exercise 1. Here we prove that a probability measure P always satisfies the stipulated assumptions.

- (i) For any events P(A) > P(B), P(A) < P(B) or P(A) = P(B).
- (ii) If A_i , B_i are partitions of $A, B, \forall i P(A_i) \leq P(B_i) \Rightarrow P(A) \leq P(B)$.
- (iii) For any A, $P(\emptyset) \leq P(A)$ and $P(\emptyset) < P(S)$

Solution. Part (i) is trivial, as $P : \mathcal{F} \to [0,1]$. Part (ii) follows from $P(A) = P(\bigcup_i A_i) = \sum_i P(A_i) \leq \sum_i P(B_i) = P(B)$. Part (iiI) $P(\emptyset) = 0, P(A) \geq 0$. Also, $P(\mathcal{S}) = 1$.

1.3 Assigning unique probabilities

In many cases, and particularly when \mathcal{F} is a finite field, there is a large number of probability distributions agreeing with our relative likelihoods.

How can we assign probabilities to events in an unambiguous manner?

Example 1.1. Consider $\mathcal{F} = \{\emptyset, A, A^{\complement}, \mathcal{S}\}$ and say $A \succ A^{\complement}$. Consequently, P(A) > 1/2. But this is insufficient for assigning a specific value to P(A).

Let A be an interval on the real line, with length $\lambda(A)$.

Definition 1.2 (Uniform distribution). $x : S \to [0, 1]$ has a uniform distribution on [0, 1] if, for any subintervals A, B of [0, 1],

 $(x \in A) \precsim (x \in B)$ iff $\lambda(A) \le \lambda(B)$

This means that *any* larger interval is more likely than *any* smaller interval. Now we shall connect the uniform distribution to the original sample space S by assuming that there is some function with uniform distribution.

Assumption 1.5 (SP5). There exists a random variable $x : S \to [0,1]$ with a uniform distribution in [0,1].

Constructing the probability distribution

We can now use the uniform distribution to create a unique probability measure that agrees with our likelihood relation. First, we have to map each event in S to an equivalent event in [0, 1].

Theorem 1.6 (Equivalent event). For any event $A \in \mathcal{F}$, there exists some $\alpha \in [0,1]$ such that $A = (X \in [0,\alpha])$.

This means that we can now define the probability of an event A by matching it to a specific equivalent event on [0, 1].

Definition 1.3 (The probability of A). Given any event A, define P(A) to be the α with $A = (X \in [0, \alpha])$.

Hence

$$A \eqsim (X \in [0, P(A)]).$$

The above is sufficient to show the following theorem.

Theorem 1.7 (Relative likelihood and probability). If assumptions SP1-SP5 are satisfied, then the probability measure P defined above is unique. Furthermore, for any two events $A, B, A \preceq B$ iff $P(A) \leq P(B)$.

1.4 Conditional likelihoods

Conditional likelihood

So far we have only considered the problem of forming opinions about which events are more likely a priori. However, we also need to have a way to incorporate evidence which may adjust our opinions. For example, while we ordinarily may think that $A \preceq B$, we may have additional information D, given which we think the opposite is true. We can formalise this through the notion of conditional likelihoods.

Example 1.2. Say that A is the event that it rains in Gothenburg, Sweden tomorrow. We know that Gothenburg is quite rainy due to its oceanic climate, so we set $A \succeq A^{\complement}$. Now, let us try and incorporate some additional information. Let D denote the fact that good weather is forecast. I personally believe that $(A \mid D) \preceq (A^{\complement} \mid D)$, i.e. that good weather is more probable than rain, given the evidence of the weather forecast.

Conditional likelihoods

Define $(A \mid D) \preceq (B \mid D)$ to mean that B is at least as likely as A when it is known that D has occurred.

Assumption 1.6 (CP). For any events A, B, D,

 $(A \mid D) \precsim (B \mid D) \quad iff \quad A \cap D \precsim B \cap D.$

Theorem 1.8. If a relation \preceq satisfies assumptions SP1 to SP5 and CP, then P is the unique probability distribution such that:

For any A, B, D such that P(D) > 0,

$$(A \mid D) \precsim (B \mid D)$$
 iff $P(A \mid D) \le P(B \mid D)$

Definition 1.4 (Conditional probability).

$$P(A \mid D) = \frac{P(A \cap D)}{P(D)}$$
(1.1)

1.5 Probability elicitation

Probability elicitation is the problem of quantifying the subjective probabilities that a particular individual uses. One of the simplest, and most direct, methods, is to simply ask. However, because we cannot simply ask somebody to completely specify a probability distribution, we can ask for this distribution iteratively.

Example 1.3 (Temperature prediction). Let τ be the temperature tomorrow at noon in Gothenburg. What are your estimates?

Eliciting the prior / forming the subjective probability measure P

- Select temperature x_0 s.t. $(\tau \le x_0) = (\tau > x_0)$.
- Select temperature x_1 s.t. $(\tau \le x_1 \mid \tau \le x_0) = (\tau > x_1 \mid \tau \le x_0)$.

Note that, necessarily, $P(\tau \leq x_0) = P(\tau > x_0) = p_0$. Since $P(\tau \leq x_0) + P(\tau > x_0) = P(\tau \leq x_0 \cup \tau > x_0) = P(\tau \in \mathbb{R}) = 1$, it follows that $p_0 = 1/2$. Similarly, $P(\tau \leq x_1 \mid \tau \leq x_0) = P(\tau > x_1 \mid \tau \leq x_0) = 1/4$.

Updating beliefs

Although we always start with a particular belief, this belief must be adjusted when we receive new evidence. In probabilistic inference, the updated beliefs are simply the probability of future events conditioned on observed events. This idea is captured neatly by Bayes' theorem, which links together the prior probability of events $P(A_i)$ with their posterior probability $P(A_i \mid B)$ given some event B and the probability $P(B \mid A_i)$ of observing the evidence B given that hypothesis A_i is true. **Theorem 1.9** (Bayes' theorem). Let A_1, A_2, \ldots be a (possibly infinite) sequence of disjoint events such that $\bigcup_{i=1}^{n} A_i = S$ and $P(A_i) > 0$ for all *i*. Let *B* be another event with P(B) > 0. Then

$$P(A_i \mid B) = \frac{P(B \mid A_i)P(A_i)}{\sum_{j=1}^{n} P(B \mid A_j)P(A_j)}$$
(1.2)

Proof. By definition, $P(A_i | B) = P(A_i \cap B)/P(B)$, and $P(A_i \cap B) = P(B | A_i)P(A_i)$, so:

$$P(A_i \mid B) = \frac{P(B \mid A_i)P(A_i)}{P(B)},$$
(1.3)

As $\bigcup_{i=1}^{n} A_i = S$, we have $B = \bigcup_{j=1}^{n} (B \cap A_j)$. Since A_i are disjoint, so are $B \cap A_i$. As P is a probability, the union property and an application of 1.3 gives

$$P(B) = P\left(\bigcup_{j=1}^{n} (B \cap A_j)\right) = \sum_{j=1}^{n} P(B \cap A_j) = \sum_{j=1}^{n} P(B \mid A_j) P(A_j).$$

A simple exercise in updating beliefs

The area of Germany Form a subjective probability for the area a of Germany in km². $A_1 : a < 10^5 \text{ km}^2$ $A_2 : a \in [10^5, 2.5 \cdot 10^5) \text{ km}^2$ $A_3 : a \in [2.5 \cdot 10^5, 5 \cdot 10^5) \text{ km}^2$ $A_4 : a \in [5 \cdot 10^5, 10^6) \text{ km}^2$ $A_5 : a \ge 10^6 \text{ km}^2$. Choose $P(A_i)$ for all i.

Additional information

- The EU's largest country is France $(6.7 \cdot 10^5 \text{km}^2)$ and the smallest is Malta with 316km^2 .
- Germany is the 4th largest of the 27 EU states
- UK $(2.4 \cdot 10^5 \text{km}^2)$ is the 8th largest EU state

The correct answer is A_3 , since $a = 3.57 \cdot 10^5$

2 Utility theory

While probability can be used to describe how likely an event is, utility can be used to describe how desirable it is. More concretely, our subjective probabilities are numerical representations of our beliefs and information. They can be taken to represent our "internal model" of the world. By analogy, our utilities are numerical representations of our tastes and preferences. Even if they are not directly known to us, we assume that our actions are such that we act so as to obtain maximum utility, in some sense.

2.1 Rewards and preferences

Rewards

Consider that we have to choose a *reward* r from a set R of possible rewards. While the elements of R may be abritrary, we shall in general find that we prefer some rewards to others. In fact, some elements of R may not even be desirable. As an example, R might be a set of tickets to different musical events, or a set of financial commodities.

Preferences

Example 2.1 (Musical event tickets). We have a set of tickets R, and we must choose the ticket $r \in R$ we prefer best.

- Case 1: R are tickets to different music events at the same time, at equally good halls with equally good seats and the same price. Here preferences simply coincide with the preferences for a certain type of music or an artist.
- Case 2: R are tickets to different events at different times, at different quality halls with different quality seats and different prices. Here, preferences may depend on all the factors.

Example 2.2 (Route selection). We have a set of alternate routes and must pick one.

- R contains two routes, one short and one long, of the same quality.
- *R* contains two routes, one short and one long, but the long route is more scenic.

Preferences among rewards

We will treat preferences in a similar manner as we have treated probabilities. That is, we will define a linear ordering among possible rewards.

Let $a, b \in R$ be two rewards. When we prefer a to b, we write $a \succ^* b$. Conversely, when we like a less than b we write $a \prec^* b$. If we like a as much as b, we write $a \equiv^* b$. We also use \succeq^* and \preceq^* for I like at least as much as and for I don't like any more than, respectively.

Properties of the preference relations.

- (i) For any $a, b \in R$, one of the following holds: $a \succ^* b, a \prec^* b, a \equiv^* b$.
- (ii) If $a, b, c \in R$ are such that $a \preceq^* b$ and $b \preceq^* c$, then $a \preceq^* c$.

2.2 Preferences among distributions

When we cannot select rewards directly

In most problems, we cannot choose the rewards directly. Rather, we must make some decision, and then obtain a reward depending on this decision. Since we may be uncertain about the outcome of a decision, we can specify our uncertainty regarding the rewards obtained by a decision in terms of a probability distribution.

- **Example 2.3** (Route selection). Each reward $r \in R$ is the time it takes to travel from A to B.
 - We prefer shorter times.
 - There are two routes, P_1 , P_2 .
 - Route P₁ takes 10 minutes when the road is clear, but 30 minutes when the traffic is heavy. The probability of heavy traffic on P₁ is q₁.
 - Route P₂ takes 15 minutes when the road is clear, but 25 minutes when the traffic is heavy. The probability of heavy traffic on P₂ is q₂.

Preferences among probability distributions

Consequently, we have to define preferences between probability distributions, rather than rewards. We use the same notation as before. Let P_1, P_2 be two distributions on (R, \mathcal{F}_R) . If we prefer P_1 to P_2 , we write $P_1 \succ^* P_2$. If we like P_1 less than P_2 , write $P_1 \prec^* P_2$. If we like P_1 as much as P_2 , we write $P_1 \eqsim^* P_2$. Finally, we also use \succeq^* and \preceq^* do denote strict preference relations.

2.3 Utility

The concept of utility allows us to create a unifying framework, such that given a particular set of rewards and probability distributions on them, we can define preferences among distributions automatically. The first step is to define utility as a way to define a preference relation among rewards.

Definition 2.1 (Utility). The utility is a function $U : R \to \mathbb{R}$, such that for all $a, b \in R$

$$a \succeq^* b \quad iff \quad U(a) \ge U(b),$$

$$(2.1)$$

The above definition is very similar to how we defined probability in terms of relative likelihood. For a given utility function, its expectation for a distribution over rewards as: **Definition 2.2** (Expected utility). The expected utility of a distribution P on R is:

$$\mathbb{E}_P(U) = \int_R U(r) \,\mathrm{d}P(r) \tag{2.2}$$

Finally, we make the assumption that the utility function is such that the expected utility remains consistent with the preference relations between all probability distributions we are choosing between.

Assumption 2.1. The expected utility hypothesis The utility of P is equal to the expected utility of the reward under P. Consequently,

$$P \succeq^* Q \quad iff \quad \mathbb{E}_P(U) \ge \mathbb{E}_Q(U).$$
 (2.3)

Example 2.4. Consider the following decision problem. You have the option of entering a lottery, for 1 CU, that gives you a prize of 10 CU. The probability of winning is 0.01. This can be formalised by making it a choice between P, where you do not enter the lottery and Q, which represents entering the lottery. Now we can calculate the expected utility for each choice. This is simply $\mathbb{E}(U \mid P)$

r	U(r)	P	Q
did not enter	0	1	0
paid 1 CU and lost	-1	0	0.99
paid 1 CU and won 10	9	0	0.01

Table 1: A simple gambling problem

$$P) = \sum_{r} U(r)P(r)$$
 and $\mathbb{E}(U \mid Q) = \sum_{r} U(r)Q(r)$ respectively.

$$\begin{array}{c|c} & \mathbf{P} & \mathbf{Q} \\ \hline \mathbb{E}(U \mid \cdot) & \mathbf{0} & -\mathbf{0.9} \end{array}$$

Table 2: Expected utility for the gambling problem

Monetary rewards

_

Example 2.5. Choose between the following two gambles:

- 1. The reward is 500,000 with certainty.
- 2. The reward is 2,500,000 with probability 0.10. It is 500,000 with probability 0.89, and 0 with probability 0.01.

Example 2.6. Choose between the following two gambles:

- 1. The reward is 500,000 with probability 0.11, or 0 with probability 0.89.
- 2. The reward is: 2,500,000 with probability 0.1, or 0 with probability 0.9.

Exercise 2. Show that if gamble 1 is preferred in the first example, gamble 1 must also be preferred in the second example.

The St. Petersburg Paradox

A simple game [Bernoulli, 1713]

- A fair coin is tossed until a head is obtained.
- If the first head is obtained on the *n*-th toss, our reward will be 2^n currency units.

How much are you willing to pay, to play this game once?

- The probability to stop at round n is 2^{-n} .
- Thus, the expected monetary gain of the game is

$$\sum_{n=1}^{\infty} 2^n 2^{-n} = \infty.$$

• If your utility function were linear you'd be willing to pay any amount to play.

2.4 Measuring utility

Experimental measurement of utility

Example 2.7. We shall try and measure the utility of all monetary rewards in some interval [a, b].

Let $\langle a, b \rangle$ denote a lottery ticket that yields a or b CU with equal probability. Consider the following sequence:

- 1. Find x_1 such that receiving x_1 CU with certainty is equivalent to receiving $\langle a, b \rangle$.
- 2. Find x_2 such that receiving x_2 CU with certainty is equivalent to receiving $\langle a, x_1 \rangle$.
- 3. Find x_3 such that receiving x_3 CU with certainty is equivalent to receiving $\langle x_1, b \rangle$.
- 4. Find x_4 such that receiving x_4 CU with certainty is equivalent to receiving $\langle x_2, x_3 \rangle$.

If $x_1 \neq x_4$, then your preferences do not meet the requirements of a utility function.

Exercise 3. 1. Specify an amount a, then observe random value Y.

- 2. If $Y \ge a$, receive Y.
- 3. If Y < a, receive random reward X with known distribution (independent of Y).
- 4. Show that you should choose a s.t. $U(a) = \mathbb{E}[U(X)]$.

2.5 Convex and concave utility functions

Convex functions

Definition 2.3. A function g is convex on A if, for any points $x, y \in A$, and any $\alpha \in [0, 1]$:

$$\alpha g(x) + (1 - \alpha)g(y) \ge g[\alpha x + (1 - \alpha)y]$$

Theorem 2.1 (Jensen's inequality). If g is convex on S and $x \in S$ with measure P(A) = 1 and $\mathbb{E}(x)$ and $\mathbb{E}[g(x)]$ exist, then:

$$\mathbb{E}[g(x)] \ge g[\mathbb{E}(x)]. \tag{2.4}$$

Example 2.8. If the utility function is convex, then we choose a gamble giving a random gain x rather than one giving a fixed gain $\mathbb{E}(x)$. Thus, a convex utility function implies risk-taking.

Concave functions

Definition 2.4. A function g is concave on S if, for any points $x, y \in S$, and any $\alpha \in [0, 1]$:

$$\alpha g(x) + (1 - \alpha)g(y) \le g[\alpha x + (1 - \alpha)y]$$

For concave functions, an analogue of Jensen's inequality holds (in the other direction). If the utility function is concave, then we choose a gamble giving a fixed gain $\mathbb{E}[X]$ rather than one giving a random gain X. Consequently, a concave utility function implies risk aversion.

Example 2.9 (Insurance). The act of buying insurance can be related to concavity of our utility function.

Let x be the insurance cost, h our insurance cover and ϵ the probability of needing the cover. Then we are going to buy insurance if the utility of losing x with certainty is greater than the utility of losing -h with probability ϵ .

$$U(-x) > \epsilon U(-h) + (1 - \epsilon)U(0).$$
(2.5)

The company has a linear utility, and fixes the premium x high enough for

$$x > \epsilon h. \tag{2.6}$$

Consequently, we see that from (2.6) that $U(-\epsilon h) \ge U(-x)$, as U is an increasing function. Thus, from (2.5) we obtain $U(-\epsilon h) > \epsilon U(-h) + (1-\epsilon)U(0)$. Now the $-\epsilon h$ term is the utility of our expected monetary loss, while the right hand side is our expected utility. Consequently if the inequality holds, our utility function is (at least locally) concave.

3 Summary

Summary

- We can subjectively indicate which events we think are more likely.
- Using relative likelihoods, we can define a *subjective probability* P for all events.

- Similarly, we can subjectively indicate preferences for rewards.
- We can determine a *utility function* for all rewards.
- Hypothesis: we prefer the probability distribution (over rewards) with the highest *expected utility*.
- Concave utility functions imply *risk aversion* (and convex, risk-taking).

References

- [1] Morris H. DeGroot. Optimal Statistical Decisions. John Wiley & Sons, 1970.
- [2] Milton Friedman and Leonard J. Savage. The expected-utility hypothesis and the measurability of utility. *The Journal of Political Economy*, 60(6):463, 1952.
- [3] Leonard J. Savage. The Foundations of Statistics. Dover Publications, 1972.