Decision Problems

Christos Dimitrakakis

April 11, 2014

Contents

1 Introduction

In this chapter we describe how to actually formulate statistical decision problems. The simplest such problem arises when we have a choice between a number of different decisions, where each decision gives us different *rewards* with different probabilities. If these probabilities are known, then the framework of expected utility maximisation gives a solution to the problem. An example includes *gambling*, where we must choose between a number of possible lotteries, each one having different payoffs and winning probabilities.

Another classical setting is parameter estimation. Therein, we stipulate the existence of a parameterised *law of nature*, and we wish to choose a best-guess set of parameters for the law through measurements and some prior information, such as for example determining the gravitational attraction constant from observing planetary movements. These measurements are always obtained through experiments, and the automatic design of those experiment is a topic we shall consider in later chapters.

Finally, these decisions will necessarily depend on our prior information, even if have access to some additional measurements. The last section of this chapter will examine how sensitive our decisions are to the prior, and how we can choose a prior distribution so that our decisions are robust.

2 Rewards that depend on the outcome of an experiment

Consider the problem of choosing between a two different types of tickets in raffle. Each type of ticket gives you the chance to win a different prize. The first is a bicycle and the second is a tea set. As most people opt for the bicycle, your chance of actually winning it is much smaller. However, if you prefer winning a bicycle to winning the tea set, it is not clear what choice you should make in the raffle. The above is the quintessential example for problems where the reward that we obtain depends not only on our decisions, but also in the outcome of an *experiment*.

More formally, we must make a decision $d \in \mathcal{D}$ *before* knowing the outcome ω of an experiment with outcomes in Ω. After the experiment is performed, we obtain a *reward* $r \in \mathcal{R}$ which depends on both the outcome of the experiment ω and our decision. As discussed in the previous chapter, our preferences for some rewards over others is determined by a *utility* function $U : \mathcal{R} \to \mathbb{R}$, such that we prefer r to r' if and only if $U(r) \geq U(r')$. Now, however, we cannot choose rewards directly. Another example, which will be used throughout this section, is the following.

Example 2.1 (Taking the umbrella). *We must decide whether to take an umbrella to work. Our reward is a combination of whether we get wet and the amount of objects that we carry. We would rather not get wet and not carry too many things, which can be made more precise by choosing an appropriate utility function. In this example, the only events of interest are whether it rains or not.*

2.1 Formalisation of the problem setting

We are now ready to formulate the problem setting more precisely.

Assumption 2.1 (Outcomes). *There exists a probability measure P on $(\Omega, \mathfrak{F}_\Omega)$ such that the probability of the random outcome ω being in $A \subset \Omega$ is:*

$$\mathbb{P}(\omega \in A) = P(A), \qquad \forall A \in \mathfrak{F}_\Omega. \tag{2.1}$$

Assumption 2.2 (Utilities). *Preferences about rewards in $\mathcal{R}$ are transitive, all rewards are comparable and there exists a utility function U, measurable with respect to $\mathfrak{F}_\mathcal{R}$ such that $U(r') \geq U(r)$ iff $r \succ^* r'$.*

Definition 2.1 (Reward function). *A reward function $\rho : \Omega \times \mathcal{D} \to \mathcal{R}$ defines the reward we obtain if we select $d \in \mathcal{D}$ and the experimental outcome is $\omega \in \Omega$:*

$$r = \rho(\omega, d). \tag{2.2}$$

Definition 2.2 (Statistical decision problem). *A statistical decision problem is a tuple $\langle \mathcal{D}, \Omega, P, U \rangle$, where $\mathcal{D}$ is a measurable decision space, Ω is a measureable space, P is a probability measure on $\mathcal{D}$, and $U : \Omega \times \mathcal{D} \to \mathcal{R}$ is a reward function and $U : \mathcal{R} \to \mathbb{R}$ is a utility function. The problem is to find:*

$$d^* \in \arg\max_{d \in \mathcal{D}} \mathbb{E}_P(U), \tag{2.3}$$

assuming that the supremum is attained and $\mathcal{D}, \Omega, P, U$ are given.

(a) The combined decision problem

(b) The separated decision problem

Figure 1: Decision diagrams for the combined and separated formulation of the decision problem. Squares denote decision variables, diamonds denote utilities. All other variables are denoted by circles. Arrows denote the flow of dependency.

The decision space might be arbitrarily more complex than the one we have seen so far. For example, our decisions may be distributions over simple decisions, or functions whose value depends on future events. We shall examine those problems later in the chapter.

When we discussed the problem of choosing between distributions, in section **??**, we had directly defined probability distributions on rewards. We can now formulate our problem in that setting. First, we define a set of distributions $\{P_i \mid i \in \mathcal{D}\}$ on $(R, \mathfrak{F}_R)$, such that the i-th decision amounts to choosing a particular distribution P_i on the rewards.

The probability measure induced by decisions
For every $d \in \mathcal{D}$, the function $\rho : \Omega \times D \to \mathcal{R}$ induces a probability distribution P_d on $\mathcal{R}$. In fact, for any $B \in \mathfrak{F}_{\mathcal{R}}$:

$$P_d(B) \triangleq \mathbb{P}(\rho(\omega, d) \in B) = P(\{\omega \mid \rho(\omega, d) \in B\}). \tag{2.4}$$

Assumption 2.3. *The sets* $\{\omega \mid \rho(\omega, d) \in B\}$ *must belong to* $\mathfrak{F}_\Omega$*. That is,* ρ *must be* $\mathfrak{F}_\Omega$*-measurable for any* d*.*

In other words, while the outcome of the experiment is independent of the decision, the distribution of rewards is effectively chosen by our decision, as before. However, this structure allows us to clearly distinguish the controllable from the random part of the rewards.

In either case, we employ the expected utility hypothesis (Assumption **??**). Thus, we should choose the decision that results in the highest expected utility.

Expected utility
The expected utility of any decision $d \in \mathcal{D}$ under P is: the expected utility is:

$$\mathbb{E}_{P_d}(U) = \int_R U(r)\,\mathrm{d}P_d(r) = \int_\Omega U[\rho(\omega, d)]\,\mathrm{d}P(\omega) \tag{2.5}$$

$\rho(\omega, d)$	d_1	d_2
ω_1	dry, carrying umbrella	wet
ω_2	dry, carrying umbrella	dry
$U[\rho(\omega, d)]$	d_1	d_2
ω_1	0	-10
ω_2	0	1
$\mathbb{E}_P(U \mid d)$	0	-1.2

Table 1: Rewards, utilities, expected utility for 20% probability of rain.

From now on, we shall use the simple notation

$$U(P, d) \triangleq \mathbb{E}_{P_d} U \tag{2.6}$$

to denote the expected utility of d under distribution P.

Instead of viewing the decision as effectively choosing a distribution over rewards (Fig. 1(a)) we can separate the random part of the process from the deterministic part (Fig. 1(b)) by considering a measure P on some space of outcomes Ω, such that the reward depends on both d and the outcome $\omega \in \Omega$ through a function $\rho(\omega, d)$. The optimal decision is of course always the $d \in \mathcal{D}$ maximising $\mathbb{E}(U \mid P_d)$.

The dependency structure of this problem in either formulation can be visualised in the *decision diagram* shown in Figure 1(a).

Example 2.2. *You are going to work, and it might rain. The forecast said that the probability of rain* (ω_1) *was* 20%. *What do you do?*

- d_1*: Take the umbrella.*
- d_2*: Risk it!*

The reward of a given outcome and decision combination, as well as the expected utility is given in table 1.

2.2 Decision diagrams

Decision diagrams are also known as *decision networks* or *influence diagrams.* Like the examples shown in Figure 1, they are used to show dependencies between different variables. In general, these include the following types of nodes:

- Choice nodes, denoted by squares. These are nodes whose values the decision maker can directly choose. Sometimes there is more than one decision maker involved.
- Value nodes, denoted by diamonds. These are the nodes that the decision maker is interested in influencing. The utility of the decision maker is a direct function of the value nodes.
- Circle nodes are used to denote all other types of variables. These include deterministic, stochastic, known or unknown variables.

The nodes are connected via directed edges. These denote the dependencies between nodes. For example, in Figure 1(b), the reward is a function of both ω and d, i.e. $r = \boldsymbol{r}(\omega, d)$, while ω depends only on the probability distribution P. Typically, there must be a path from a choice node to a value node, otherwise nothing the decision maker can do will influence its utility. Nodes belonging to or observed by different players will usually be denoted by different lines or colors. In Figure 1(b), ω, which is not observed, is shown in a lighter color.

2.3 Statistical estimation

This is especially the case in statistical problem of *parameter estimation,* such as estimating the covariance matrix of a Gaussian random variable. A simple example is estimating the distribution of votes in an election from a small sample.

Example 2.3 (Voting)**.** *Let us say for example that you wish to estimate the number of votes for different candidates in an election. The* unknown parameters *of the problem mainly include: the percentage of likely voters in the population, the probability that a likely voter is going to vote for each candidate. One simple way to estimate this is by polling.*

Consider a nation with k political parties. Let $\omega = (\omega_1, \ldots, \omega_k) \in [0,1]^k$ be the voting percentages for each party. We wish to make a guess $d \in [0,1]^k$. How should we guess, given a distribution $P(\omega)$? How should we select U and ρ? This depends on what our goals is, when we make the guess.

If we wish to give a reasonable estimate about all the k parties votes, we can use the squared error: First, set $\rho(\omega, d) = (\omega_1 - d_1, \ldots, \omega_k - d_k)$, our error vector $r \in [0,1]^k$. Then we set $U(r) \triangleq -\|r\|^2$, where $\|r\|^2 = \sum_i |x_i|^2$.

If on the other hand, we just want to predict the winner of the election, then the actual percentages of all individual parties are not important. In that case, we can set $\rho(\omega, d) = 1$ if $\arg\max_i \omega_i = \arg\max_i d_i$ and 0 otherwise, and $U(r) = r$.

- The unknown outcome of the experiment ω is called a *parameter.*
- The set of outcomes Ω is called the *parameter space.*

Losses and risks

In such problems, it is common to specify a loss instead of a utility. This is usually the negative utility:

Definition 2.3 (Loss)**.**

$$\ell(\omega, d) = -U[\rho(\omega, d)]. \tag{2.7}$$

Given the above, instead of the expected utility, we consider the expected loss, or risk.

Definition 2.4 (Risk)**.**

$$\sigma(P, d) = \int_{\Omega} \ell(\omega, d) \, \mathrm{d}P(\omega). \tag{2.8}$$

Of course, the optimal decision is d minimising σ.

3 Bayes decisions

Definition 3.1 (Bayes-optimal utility). *Consider parameter space Ω, decision space $\mathcal{D}$, and a utility function $U : \Omega \times \mathcal{D} \to \mathbb{R}$. For any probability distribution P on Ω, the Bayes-optimal utility $U^*(P)$ is defined as the smallest upper bound on $U(P, d)$ for all decisions $d \in \mathcal{D}$. That is,*

$$U^*(P) = \sup_{d \in D} U(P, d). \tag{3.1}$$

Remark 3.1. *For any function $f : X \to Y$, where Y is equipped with a complete binary relation $<$, we define, for any $A \subset X$*

$$M = \inf_{x \in A} f(x)$$

such that $M \leq f(x)$ for any $x \in A$. In other words, M is a lower bound on $f(x)$. Furthermore, for any $M' > M$, there exists some $x' \in A$ s.t. $M' > f(x')$. In other words, there exists no greater lower bound than M.

As can be seen from Figure **??**, for absolute loss, the optimal decision is to choose the d that is closest to the most likely ω. However, for quadratic loss, Figure **??** appears to indicate that the optimal choice should be equal to the expected value of ω. This is actually true in general for quadratic loss, and for $d, \omega \in \mathbb{R}$, as shall be seen from the following example.

Example 3.1 (Quadratic loss). *Now consider $\Omega = \mathbb{R}$ with measure P and $\mathcal{D} = \mathbb{R}$. For any point $\omega \in \mathbb{R}$, we define the utility as:*

$$U(\omega, d) = -|\omega - d|^2. \tag{3.2}$$

The optimal decision maximises

$$U(P, d) = -\int_{\mathbb{R}} |\omega - d|^2 \, \mathrm{d}P(\omega).$$

Then, as long as $\partial/\partial d |\omega - d|^2$ is measurable with respect to $\mathfrak{F}_{\mathbb{R}}$

$$\begin{aligned}
\frac{\partial}{\partial d} \int_{\mathbb{R}} |\omega - d|^2 \, \mathrm{d}P(\omega) &= \int_{\mathbb{R}} \frac{\partial}{\partial d} |\omega - d|^2 \, \mathrm{d}P(\omega) && (3.3)\\
&= 2 \int_{\mathbb{R}} (d - \omega) \, \mathrm{d}P(\omega) && (3.4)\\
&= 2 \int_{\mathbb{R}} d \, \mathrm{d}P(\omega) - 2 \int_{\mathbb{R}} \omega \, \mathrm{d}P(\omega) && (3.5)\\
&= 2d - 2\,\mathbb{E}(\omega), && (3.6)
\end{aligned}$$

so the expected utility is maximised for $d = \mathbb{E}(\omega)$.

3.1 Convexity of the Bayes-optimal utility

We shall show now the expected utility is linear. Consequently, the Bayes-utility is convex with respect to the distribution P. This firstly implies that there is a unique “worst” distribution P, against which we cannot do very well. Secondly, we can approximate the Bayes-utility very well for all possible distributions by generalising from a small number of distributions. In order to define linearity and convexity, we must first introduce the concept of a mixture of distributions.

A mixture of distributions

Consider two probability measures P, Q on $(\Omega, \mathfrak{F}_\Omega)$.

These define two alternative distributions for ω. For any P, Q and $\alpha \in [0,1]$, we define

$$Z_\alpha \triangleq \alpha P + (1-\alpha)Q \tag{3.7}$$

to mean the probability measure such that $Z_\alpha(A) = \alpha P(A) + (1-\alpha)Q(A)$ for any $A \in \mathfrak{F}_\Omega$.

Remark 3.2 (Linearity of the expected utility)**.** *If Z_α is as defined in* (3.7)*, then, for any $d \in \mathcal{D}$:*

$$U(Z_\alpha, d) = \alpha U(P, d) + (1-\alpha)U(Q, d). \tag{3.8}$$

Proof. This follows from the linearity of expectation, i.e.

$$\begin{aligned} U(Z_\alpha, d) &= \int_\Omega U(\omega, d)\,\mathrm{d}Z_\alpha(\omega) && (3.9)\\ &= \alpha \int_\Omega U(\omega, d)\,\mathrm{d}P(\omega) + (1-\alpha)\int_\Omega U(\omega, d)\,\mathrm{d}Q(\omega) && (3.10)\\ &= \alpha U(P, d) + (1-\alpha)U(Q, d). && (3.11) \end{aligned}$$

□

Theorem 3.1. *For probability measures P, Q on Ω and any $\alpha \in [0,1]$,*

$$U^*[Z_\alpha] \le \alpha U^*(P) + (1-\alpha)U^*(Q), \tag{3.12}$$

where $Z_\alpha = \alpha P + (1-\alpha)Q$.

Proof. From the definition of the expected utility (3.8), for any decision $d \in \mathcal{D}$,

$$U(Z_\alpha, d) = \alpha U(P, d) + (1-\alpha)U(Q, d).$$

Hence, by definition (3.1) of the Bayes-utility:

$$\begin{aligned} U^*(Z_\alpha) &= \sup_{d \in D} U(Z_\alpha, d)\\ &= \sup_{d \in D}[\alpha U(P, d) + (1-\alpha)U(Q, d)]. \end{aligned}$$

Use $\sup_x[f(x) + g(x)] \le \sup_x f(x) + \sup_x g(x)$ to bound r.h.s:

$$\begin{aligned} U^*[Z_\alpha] &\le \alpha \sup_{d \in D} U(P, d) + (1-\alpha)\sup_{d \in D} U(Q, d)\\ &= \alpha U^*(P) + (1-\alpha)U^*(Q). \end{aligned}$$

□

Convexity of the Bayes utility

As we have proven, the expected utility is linear with respect to P. Thus, for any fixed decision d we obtain one of the lines in Fig. 2. Due to the theorem just proved, the Bayes risk is concave. Furthermore, the minimising decision for any P is tangent to the risk at the point $(P, U^*(P))$. If we take a decision

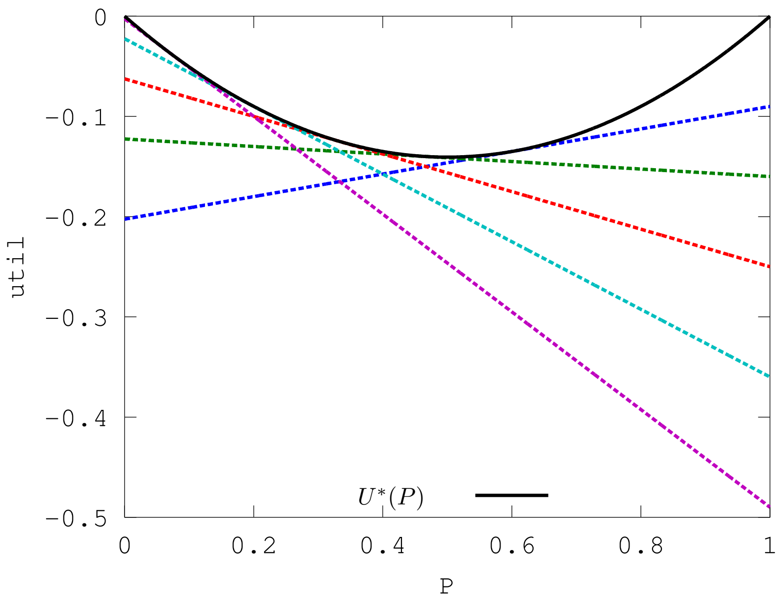

Figure 2: A strictly convex Bayes utility.

that is optimal with respect to some P, but the distribution is in fact $Q \neq P$, then we are not far from the optimal when P and Q are close and U^* is smooth. Consequently, we can trivially lower bound the Bayes utility by examining a finite set of decisions $\hat{D}$:

$$U^*(P) \geq \max_{d \in \hat{D}} U(P, d) \forall P$$

In addition, we can upper-bound the Bayes utility as follows. Take any two distributions P_1, P_2 in the set of allowed distributions. Then, the following upper bound holds

$$U^*(\alpha P_1 + (1 - \alpha) P_2) \leq \alpha U^*(P_1) + (1 - \alpha) U^*(P_2)$$

due to convexity. The two bounds suggest an algorithm for successive approximation of the Bayes risk, by looking for the largest gap between the lower and the upper bounds.

4 Decision problems with observations

So far we have only examined problems where the parameters were drawn from some distribution. This distribution constituted our subjective belief about what the unknown parameter is. Now, we examine the case where we can obtain some observations that depend on the unknown ω before we make our decision. These observations should give us more information about the parameter, before making a decision. Intuitively, we should be able to make decisions by simply considering the posterior distribution. The following section will investigate whether this is true.

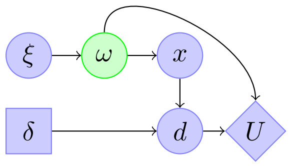

Figure 3: Statistical decision problem with observations

Obtaining information

In this setting, we once more need to take some decision $d \in \mathcal{D}$ so as to maximise expected utility. As before, we have a prior distribution ξ on some parameter $\omega \in \Omega$, representing what we know about ω. Consequently, the expected utility of any fixed decision d is going to be $\mathbb{E}_\xi(U \mid d)$.

However, it might be possible to obtain more infomation about ω before making a decision. In particular, each ω corresponds to a *model* of the world ψ_ω. This is expressed as a probability distribution over the observation space $\mathcal{S}$, such that $\psi_\omega(X)$ is the probability that the observation is in $X \subset \mathcal{S}$. The set of parameters Ω thus defines a family of models:

$$\Psi \triangleq \{\psi_\omega \mid \omega \in \Omega\}. \tag{4.1}$$

Now, consider the case where we take an observation x from the true model ψ_{ω^*} before having to make a decision. We can represent the dependency of our decisions on the observations by making our decision a function of x:

Definition 4.1 (Decision function). *A decision function $\delta : \mathcal{S} \to \mathcal{D}$ maps from the set of possible observations $\mathcal{S}$ to the set of possible decisions.*

The expected utility of a decision function δ is:

$$U(\xi, \delta) \triangleq \mathbb{E}_\xi \{U[\omega, \delta(x)]\} = \int_\Omega \left(\int_\mathcal{S} U[\omega, \delta(x)] \, \mathrm{d}\psi_\omega(x) \right) \mathrm{d}\xi(\omega). \tag{4.2}$$

When the set of decision functions includes all fixed decisions, then there is a decision function δ^* at least as good as the best fixed decision d^*. More formally:

Remark 4.1. *Let $\mathscr{D}$ denote the set of decision functions $\delta : \mathcal{S} \to \mathcal{D}$. If, $\forall d \in \mathcal{D}$ $\exists \delta \in \mathscr{D}$ such that $\delta(x) = d$ $\forall x \in \mathcal{S}$, then $\sup_{\delta \in \mathscr{D}} \mathbb{E}_\xi(U \mid \delta) \geq \sup_{d \in \mathcal{D}} \mathbb{E}_\xi(U \mid d)$.*

Proof. The proof follows by setting $\mathscr{D}_0$ to be the set of fixed decision functions. The result follows since $\mathscr{D}_0 \subset \mathscr{D}$. □

This is the standard Bayesian framework for decision making. It may be slightly more intuitive in some case to use the notation $\psi(x \mid \omega)$, in order to emphasize that this is a conditional distribution. However, there is no technical difference between the two notations.

Example 4.1. *Consider the problem of deciding whether or not to go to a particular restaurant. Let $\Omega = [0, 1]$ with $\omega = 0$ meaning the food is in general horrible and $\omega = 1$ meaning the restaurant is great. Let $x_1, \ldots, x_n$ be n expert opinions in $\mathcal{S} = \{0, 1\}$ about the restaurant. Under our model, the probability of* *observing* $x_i = 1$ *when the quality of the restaurant is* ω *is given by* $\psi_\omega(1) = \omega$ *and conversely* $\psi_\omega(0) = 1-\omega$. *The probability of observing a particular*[1] *sequence* x *of length* n *is*

$$\psi_\omega(x) = \omega^s (1-\omega)^{n-s}$$

with $s = \sum_{i=1}^n x_i$.

Maximising utility when making observations

Statistical procedures based on the notion that a distribution can be assigned to any parameter in a statistical decision problem, as we are assuming here, are called *Bayesian statistical methods.* The scope of these methods has been the subject of much discussion in the statistical literature. See e.g. [2].

In the following, we shall look at different expressions for the expected utility. We shall overload the utility operator U for various cases: when the parameter is fixed, when the parameter is random, when the decision is fixed, and when the decision depends on the observation x and thus is random as well.

Expected utility of a fixed decision d with $\omega \sim \xi$

We first consider the expected utility of taking a fixed decision $d \in \mathcal{D}$, when $\mathbb{P}(\omega \in A) = \xi(A)$. This is the case we have dealt with so far.

$$U(\xi, d) \triangleq \mathbb{E}_\xi(U \mid d) = \int_\Omega U(\omega, d) \,\mathrm{d}\xi(\omega). \tag{4.3}$$

Expected utility of a decision function δ with fixed $\omega \in \Omega$

Now assume that ω is fixed, but instead of selecting a decision directly, we select a decision that depends on the random observation x, which is distributed according to ψ_ω on $\mathcal{S}$. We do this by defining a function $\delta : \mathcal{S} \to \mathcal{D}$.

$$U(\omega, \delta) = \int_{\mathcal{S}} U(\omega, \delta(x)) \,\mathrm{d}\psi_\omega(x). \tag{4.4}$$

Expected utility of a decision function δ with $\omega \sim \xi$

Now we generalise to the case where ω is distributed with measure ξ. Note that the expectation of the previous expression (4.4) is by definition written as:

$$U(\xi, \delta) = \int_\Omega U(\omega, \delta) \,\mathrm{d}\xi(\omega), \qquad U^*(\xi) \triangleq \sup_\delta U(\xi, \delta) = U(\xi, \delta^*). \tag{4.5}$$

[1] We obtain a slightly different probability under the binomial model, but the end result is the same.

Bayes decision rules

We wish to construct the Bayes decision rule, that is, the decision function with maximal ξ-expected utility. However, doing so by examining all possible decision functions is hard, because (usually) there are many more decision functions than decisions. It is however, easy to find the Bayes decision for each possible observation.

Theorem 4.1. *If U is non-negative or bounded, then we can reverse the integration order of*

$$U(\xi, \delta) = \mathbb{E}\left\{U[\omega, \delta(x)]\right\} = \int_\Omega \int_\mathcal{S} U[\omega, \delta(x)] \,\mathrm{d}\psi_\omega(x) \,\mathrm{d}\xi(\omega),$$

which is the normal form, to obtain the risk in extensive form.

$$U(\xi, \delta) = \int_\mathcal{S} \int_\Omega U[\omega, \delta(x)] \,\mathrm{d}\xi(\omega \mid x) \,\mathrm{d}f(x), \tag{4.6}$$

where $f(x) = \int_\Omega \psi_\omega(x) \,\mathrm{d}\xi(\omega)$.

Proof. To prove this when U is non-negative, we shall use Tonelli's theorem. First we need to construct an appropriate product measure. Note that the original is written Let $p(x \mid \omega) \triangleq \frac{\mathrm{d}\psi_\omega(x)}{\mathrm{d}\nu(x)}$ be the Radon-Nikodym derivative of ψ_ω with respect to some dominating measure ν on $\mathcal{S}$. Similarly, let $p(\omega) \triangleq \frac{\mathrm{d}\xi(\omega)}{\mathrm{d}\mu(x)}$ be the corresponding derivative for ξ. Now, the utility can be written as:

$$\begin{aligned} U(\xi, \delta) &= \int_\Omega \int_\mathcal{S} U[\omega, \delta(x)] p(x \mid \omega) p(\omega) \,\mathrm{d}\nu(x) \,\mathrm{d}\mu(\omega) && (4.7)\\ &= \int_\Omega \int_\mathcal{S} h(\omega, x) \,\mathrm{d}\nu(x) \,\mathrm{d}\mu(\omega). && (4.8) \end{aligned}$$

Clearly, if U is non-negative, then h is non-negative. Then, Tonelli's theorem applies and

$$\begin{aligned} U(\xi, \delta) &= \int_\mathcal{S} \int_\Omega h(\omega, x) \,\mathrm{d}\mu(\omega) \,\mathrm{d}\mu(x) && (4.9)\\ &= \int_\mathcal{S} \int_\Omega p(x \mid \omega) p(\omega) \,\mathrm{d}\mu(\omega) \,\mathrm{d}\nu(x) && (4.10)\\ &= \int_\mathcal{S} \int_\Omega p(\omega \mid x) \,\mathrm{d}\mu(\omega) p(x) \,\mathrm{d}\nu(x) && (4.11)\\ &= \int_\mathcal{S} \left[\int_\Omega p(\omega \mid x) \,\mathrm{d}\mu(\omega)\right] p(x) \,\mathrm{d}\nu(x) \quad = \int_\mathcal{S} \left[\int_\Omega \mathrm{d}\xi(\omega \mid x)\right] \mathrm{d}f(x), && (4.12) \end{aligned}$$

where $p(x) = \mathrm{d}f(x) / \mathrm{d}\nu(x)$. □

We can construct an optimal decision function δ^* as follows. For any specific observed $x \in S$, we set $\delta^*(x)$ to:

$$\delta^*(x) \triangleq \operatorname*{arg\,max}_{d \in D} \mathbb{E}_\xi(U \mid x, d) = \operatorname*{arg\,max}_{d \in D} \int_\Omega U(\omega, d) \,\mathrm{d}\xi(\omega \mid x).$$

So now we can plug δ^* in the extensive form to obtain:

$$\int_{\mathcal{S}} \int_{\Omega} U[\omega, \delta^*(x)] \,\mathrm{d}\xi(\omega \mid x) \,\mathrm{d}f(x) = \int_{\mathcal{S}} \left\{ \min_d \int_{\Omega} U[\omega, d] \,\mathrm{d}\xi(\omega \mid x) \right\} \mathrm{d}f(x).$$

Consequently, there is no need to completely specify the decision function before we have seen x. In particular, this would create problems when $\mathcal{S}$ is large.

Definition 4.2 (Prior distribution)**.** *The distribution ξ is called the prior distribution of ω.*

Definition 4.3 (Marginal distribution)**.** *The distribution f is called the (prior) marginal distribution of x.*

Definition 4.4 (Posterior distribution)**.** *The conditional distribution $\xi(\cdot \mid x)$ is called the posterior distribution of ω.*

Bayes decision rule.

The *optimal decision* given x, is the optimal decision with respect to the *posterior* $\xi(\omega \mid x)$. Thus, we do not need to pre-compute the complete Bayes-optimal decision rule.

4.1 Calculating posteriors

Posterior distributions for multiple observations

We now consider how we can re-write the posterior distribution over Ω incrementally. Assume that we observe $x^n \triangleq x_1, \ldots, x_n$. We have a prior ξ on Ω. For the observations, we write:

Observation probability given history x^{n-1} and parameter ω

$$\psi_\omega(x_n \mid x^{n-1}) = \frac{\psi_\omega(x^n)}{\psi_\omega(x^{n-1})}$$

Now we can write the posterior as follows:

Posterior recursion

$$\xi(\omega \mid x^n) = \frac{\psi_\omega(x^n)\xi(\omega)}{f(x^n)} = \frac{\psi_\omega(x_n \mid x^{n-1})\xi(\omega \mid x^{n-1})}{f(x_n \mid x^{n-1})}. \tag{4.13}$$

Here $f(\cdot \mid \cdot) = \int_\Omega \psi_\omega(\cdot \mid \cdot) \,\mathrm{d}\xi(\omega)$ is a marginal distribution.

Posterior distributions for multiple independent observations

Now we consider the case where, given the parameter, the next observation does not depend on the history: If $\psi(x_n \mid \omega, x^{n-1}) = \psi_\omega(x_n)$ then $\psi_\omega(x^n) = \prod_{k=1}^n \psi_\omega(x_k)$. Then:

Posterior recursion with conditional independence

$$\xi_n(\omega) \triangleq \xi_0(\omega \mid x^n) = \frac{\psi_\omega(x^n)\xi_0(\omega)}{f_0(x_n)} \tag{4.14}$$

$$= \xi_{n-1}(\omega \mid x_n) = \frac{\psi_\omega(x_n)\xi_{n-1}(\omega)}{f_{n-1}(x_n)}, \tag{4.15}$$

where we define ξ_t to be the belief at time t. Here $f_n(\cdot \mid \cdot) = \int_\Omega \psi(\cdot \mid \cdot, \omega) \, \mathrm{d}\xi_n(\omega)$ is the marginal distribution with respect to the n-th posterior.

Conditional independence allows us to write the posterior update as an identical recursion at each time t. We shall take advantage of that when we look at *conjugate prior* distributions. For such models, the recursion involves a particularly simple parameter update.

Quick summary

- We want to make a decision against an unknown parameter ω.
- The risk is the negative expected utility.
- The Bayes risk is the minimum risk, and it is concave with respect to the distribution of ω.
- Our decisions can depend on observations, via a decision function.
- We can construct a complete decision function by computing $U(\xi, \delta)$ for all *decision functions* (normal form).
- We can instead wait until we observe x and compute $U[\xi(\cdot \mid x), d]$ for all *decisions* (extensive form).
- The posterior given multiple observations can be computed recursively using independence.

References

[1] Morris H. DeGroot. *Optimal Statistical Decisions*. John Wiley & Sons, 1970.

[2] Leonard J. Savage. *The Foundations of Statistics*. Dover Publications, 1972.