

Meeting 7: Utility

“Solution” to Problem 3 from Meeting 6

Probability patient has disease A, $P(A) = 0.3$

Probability patient has disease B, $P(B) = 0.4$

Probability patient has just temporary symptoms, $P(A \wedge B) = 0 \Rightarrow P(\neg(A \vee B)) = 0.3$

Express the consequences in terms of percentage of symptom reduction and treat them as “gains” (similar to payoff)

<i>Action</i>	<i>Patient's status</i>		
	Disease A <i>prob. 0.3</i>	Disease B <i>prob. 0.4</i>	Temporary symptoms <i>prob. 0.3</i>
Treat for A	100%	0 %	0 %
Treat for B	10 %	40 %	100 %
No treatment	0 %	0 %	100 %

\Rightarrow Expected “gains”:

With Treat for A: $100 \cdot 0.3 + 0 \cdot 0.4 + 0 \cdot 0.3 = 30$

With Treat for B: $10 \cdot 0.3 + 40 \cdot 0.4 + 100 \cdot 0.3 = 49$

Wit No treatment: $0 \cdot 0.3 + 0 \cdot 0.4 + 100 \cdot 0.3 = 30$

\Rightarrow The Bayes’ action is therapy treatment for B

Utility

What do we mean by *utility* ?

Utility is naturally coupled with the consequence of a particular action for a certain state of the world, i.e.

$$c = c(a, \boldsymbol{\theta})$$

Consequences can often be expressed in monetary terms, i.e. using payoff functions we may have $c(a, \boldsymbol{\theta}) = R(a, \boldsymbol{\theta})$.

Cash equivalents may be used to transform non-monetary consequences into monetary ones.

But is utility always about the value of money and is there always a numerical scale that can be used for representing it?

Utility always relates to the relative preferences of consequences with the decision maker.

The following notation is used for comparison of consequences, when the consequences are such that it is not possible to use a simple numerical ordering

$c_i \prec c_j$ means that consequence c_j is preferred to consequence c_i

$c_i \sim c_j$ means that c_i and c_j are equally preferred

$c_i \not\prec c_j$ means that c_i is not preferred to c_j

The Bayes action (maximising the expected payoff) is not always the obvious probabilistic criterion.

A consequence of an action can be preferred to another consequence by one decision-maker, while the opposite can hold for another decision-maker.

Example

Assume that when the temperature is above 25 °C and you have decided to wear long trousers and a long sleeves shirt, you will as a consequence feel unusually hot

$$c_1 = c(a = \text{“longs”}, \theta > 25 \text{ °C})$$

Moreover, assume that when the temperature is below 15 °C and you have decided to wear shorts and a t-shirt you will as a consequence feel unusually cold

$$c_2 = c(a = \text{“shorts”}, \theta < 15 \text{ °C})$$

Your preference order would be one of $c_1 \prec c_2$, $c_2 \prec c_1$ and $c_1 \sim c_2$

$$c_1 = c \text{ (} a = \text{“longs”, } \theta > 25 \text{ }^\circ\text{C) }$$

$$c_2 = c \text{ (} a = \text{“shorts”, } \theta < 15 \text{ }^\circ\text{C) }$$

If you think it is always better to feel warm than cold your preference order will be

$$c_2 \prec c_1$$

Another person feeling the same as you may really dislike feeling too warm and hence has the preference order

$$c_1 \prec c_2$$

A third person also feeling the same as you may be someone who would always complain as soon as weather condition and choice of garments do not “fit” well probably has the preference order

$$c_1 \sim c_2$$

To be able to allow for a relative desirability that deviates from the linear comparability of monetary consequences we introduce a so-called *utility function*:

$$U(c) = U(c(a, \boldsymbol{\theta})) = U(a, \boldsymbol{\theta})$$

If the difference in payoff between two pairs of action and state of world is d_R , *i.e.*

$$d_R = R(a_1, \boldsymbol{\theta}_1) - R(a_2, \boldsymbol{\theta}_2)$$

the following three differences in utility may hold

$$U(a_1, \boldsymbol{\theta}_1) - U(a_2, \boldsymbol{\theta}_2) < k \cdot d_R$$

$$U(a_1, \boldsymbol{\theta}_1) - U(a_2, \boldsymbol{\theta}_2) = k \cdot d_R$$

$$U(a_1, \boldsymbol{\theta}_1) - U(a_2, \boldsymbol{\theta}_2) > k \cdot d_R$$

where k is any constant > 0 that can take care of that utility and payoff may be given on different scales.

Two axioms of utility:

1. If $c_1 \prec c_2$ then $U(c_1) < U(c_2)$ and if $c_1 \sim c_2$ then $U(c_1) = U(c_2)$

2. If

- $O_1 =$ Obtaining consequence c_1 for certain
- $O_2 =$ Obtaining consequence c_2 with probability p and obtaining consequence c_3 with probability $1-p$
- $O_1 \sim O_2$

then $U(c_1) = p \cdot U(c_2) + (1-p) \cdot U(c_3)$

“ p -mixture”

Hence, it is not necessary to work with preferences and their notations (\prec , \sim , \preceq).

All preferences can be expressed in terms of the utility function:

$$c_1 \prec c_2 \quad \Leftrightarrow \quad U(c_1) < U(c_2)$$

$$c_1 \sim c_2 \quad \Leftrightarrow \quad U(c_1) = U(c_2)$$

$$c_1 \preceq c_2 \quad \Leftrightarrow \quad U(c_1) \leq U(c_2)$$

Now, assume $U(a, \theta)$ is a utility function and let $W(a, \theta) = b + d \cdot U(a, \theta)$ where b and d are constants with $d > 0$.

$$\begin{aligned}
 &\text{If } U(a_i, \theta_k) < U(a_j, \theta_l) \text{ [where } i \neq j \text{ or } k \neq l \text{ or both; } c(a_i, \theta_k) < c(a_j, \theta_l)] \\
 &\Rightarrow \\
 &W(a_i, \theta_k) = b + d \cdot U(a_i, \theta_k) \\
 &W(a_j, \theta_l) = b + d \cdot U(a_j, \theta_l) \\
 &\Rightarrow W(a_j, \theta_l) - W(a_i, \theta_k) = b + d \cdot U(a_j, \theta_l) - (b + d \cdot U(a_i, \theta_k)) \\
 &= \underbrace{d}_{>0} \cdot \underbrace{(U(a_j, \theta_l) - U(a_i, \theta_k))}_{>0} > 0
 \end{aligned}$$

$$\begin{aligned}
 &\text{If } U(a_i, \theta_k) = U(a_j, \theta_l) \text{ [where } i \neq j \text{ or } k \neq l \text{ or both}] \\
 &\Rightarrow W(a_j, \theta_l) - W(a_i, \theta_k) = b + d \cdot U(a_j, \theta_l) - (b + d \cdot U(a_i, \theta_k)) \\
 &= \underbrace{d}_{>0} \cdot \underbrace{(U(a_j, \theta_l) - U(a_i, \theta_k))}_{=0} = 0
 \end{aligned}$$

$$\begin{aligned}
 &\text{If } U(a_i, \theta_k) \\
 &= p \cdot U(a_{j_1}, \theta_{l_1}) + (1 - p) \cdot U(a_{j_2}, \theta_{l_2}) \text{ [utilities for 3 different consequences]} \\
 &\Rightarrow W(a_i, \theta_k) = b + d \cdot U(a_i, \theta_k) = b + d \cdot (p \cdot U(a_{j_1}, \theta_{l_1}) + (1 - p) \cdot U(a_{j_2}, \theta_{l_2})) = \\
 &= b \cdot p + b \cdot (1 - p) + d \cdot p \cdot U(a_{j_1}, \theta_{l_1}) + d \cdot (1 - p) \cdot U(a_{j_2}, \theta_{l_2}) = \\
 &= p \cdot (b + d \cdot U(a_{j_1}, \theta_{l_1})) + (1 - p) \cdot (b + d \cdot U(a_{j_2}, \theta_{l_2})) = \\
 &= p \cdot W(a_{j_1}, \theta_{l_1}) + (1 - p) \cdot W(a_{j_2}, \theta_{l_2})
 \end{aligned}$$

Now, assume $U(a, \theta)$ is a utility function and let $W(a, \theta) = b + d \cdot U(a, \theta)$ where b and d are constants with $d > 0$.

If $U(a_i, \theta_k) < U(a_j, \theta_l)$ [where $i \neq j$ or $k \neq l$ or both ; $c(a_i, \theta_k) < c(a_j, \theta_l)$]

\Rightarrow

$$W(a_i, \theta_k) = b + d \cdot U(a_i, \theta_k)$$

$$W(a_j, \theta_l) = b + d \cdot U(a_j, \theta_l)$$

$$\Rightarrow W(a_j, \theta_l) - W(a_i, \theta_k) = b + d \cdot U(a_j, \theta_l) - (b + d \cdot U(a_i, \theta_k))$$

$$= \underbrace{d}_{>0} \cdot \left(\underbrace{U(a_j, \theta_l) - U(a_i, \theta_k)}_{>0} \right)$$

A utility function is only unique up to a linear transformation

If $U(a_i, \theta_k) = U(a_j, \theta_l)$ [where $i \neq j$ or $k \neq l$ or both] \Rightarrow

$$W(a_j, \theta_l) - W(a_i, \theta_k) = b + d \cdot U(a_j, \theta_l) - (b + d \cdot U(a_i, \theta_k))$$

$$= \underbrace{d}_{>0} \cdot \left(\underbrace{U(a_j, \theta_l) - U(a_i, \theta_k)}_{=0} \right) = 0$$

If $U(a_i, \theta_k)$

$$= p \cdot U(a_{j_1}, \theta_{l_1}) + (1 - p) \cdot U(a_{j_2}, \theta_{l_2}) \text{ [utilities for 3 different consequences]}$$

$$\Rightarrow W(a_i, \theta_k) = b + d \cdot U(a_i, \theta_k) = b + d \cdot (p \cdot U(a_{j_1}, \theta_{l_1}) + (1 - p) \cdot U(a_{j_2}, \theta_{l_2})) =$$

$$= b \cdot p + b \cdot (1 - p) + d \cdot p \cdot U(a_{j_1}, \theta_{l_1}) + d \cdot (1 - p) \cdot U(a_{j_2}, \theta_{l_2}) =$$

$$= p \cdot (b + d \cdot U(a_{j_1}, \theta_{l_1})) + (1 - p) \cdot (b + d \cdot U(a_{j_2}, \theta_{l_2})) =$$

$$= p \cdot W(a_{j_1}, \theta_{l_1}) + (1 - p) \cdot W(a_{j_2}, \theta_{l_2})$$

The *expected utility* of an action a with respect to a probability distribution of the states of the world is obtained – analogously to how expected payoff and expected loss are obtained – by integrating the utility function with the probability distribution of θ using its probability density (or mass) function $g(\theta)$:

$$EU = E_g \left(U(a, \tilde{\theta}) \right) = \int_{\theta} U(a, \theta) \cdot g(\theta) d\theta$$

When data is not taken into account $g(\theta)$ is the prior pdf/pmf $f'(\theta) = p(\theta)$:

$$EU = \int_{\theta} U(a, \theta) \cdot f'(\theta) d\theta = \int_{\theta} U(a, \theta) \cdot p(\theta) d\theta$$

When data (\mathbf{x}) is taken into account, $g(\theta)$ is the posterior pdf/pmf $f''(\theta|\mathbf{x}) = q(\theta|\mathbf{x})$:

$$EU = \int_{\theta} U(a, \theta) \cdot f''(\theta|\mathbf{x}) d\theta = \int_{\theta} U(a, \theta) \cdot q(\theta|\mathbf{x}) d\theta$$

Assessing/finding a utility function

Let

c_1 be the worst consequence and c_2 be the best consequence

Normalise – without loss of generality – the utility function $U(c)$ such that $U(c_1) = 0$ and $U(c_2) = 1$

For a particular action with consequence c it must always hold that $0 \leq U(c) \leq 1$

Now, assume a gamble in which you should choose between

1. Obtaining consequence c for certain
2. Obtain consequence c_1 with probability $1-p$ and consequence c_2 with probability p

With the first choice the expected utility is $U(c)$.

With the second choice the expected utility is

$$U(c_1) \cdot [1 - p] + U(c_2) \cdot p = 0 \cdot (1 - p) + 1 \cdot p = p$$

1. Obtaining consequence c for certain
2. Obtain consequence c_1 with probability $1-p$ and consequence c_2 with probability p

For a certain value of p , p_0 say, you will be indifferent between 1 and 2

$$\text{Hence } U(c) = (1 - p_0) \cdot \underbrace{U(c_1)}_{=0} + p_0 \cdot \underbrace{U(c_2)}_{=1} = p_0$$

This means that $U(c)$ can be seen as proportional to the probability of obtaining the best consequence.

$$\Pr(\text{Best consequence}|a, \boldsymbol{\theta}) \propto U(a, \boldsymbol{\theta})$$

\Rightarrow

$$\Pr(\text{Best consequence}|a) \propto \int_{\boldsymbol{\theta}} U(a, \boldsymbol{\theta}) \cdot g(\boldsymbol{\theta}) d\boldsymbol{\theta} = \overline{U(a, g)} (= E_{g(\boldsymbol{\theta})}(U(a, \boldsymbol{\theta})))$$

Hence, the optimal action is the action that *maximises the expected utility* under the probability distribution that rules the state of nature

$$a_g^{(\text{optimal})} = \arg \max_{a \in \mathcal{A}} \overline{U(a, g)} \quad \mathcal{A} \text{ is the set of possible actions}$$

\Rightarrow The *Bayes action (decision)* is

$$a^{(B)} = \begin{cases} \arg \max_{a \in \mathcal{A}} \left(\overline{U(a, p)} \right) & \text{when no data are used} \\ \arg \max_{a \in \mathcal{A}} \left(\overline{U(a, q, \boldsymbol{x})} \right) & \text{when data, } \boldsymbol{x}, \text{ is used} \end{cases}$$

Loss function

Utilities can – like payoffs – be both positive and negative. A negative utility – which may be called *disutility* – comes with a non-desirable consequence.

When all consequences are non-desirable it is common to describe the decision problem in terms of losses instead of utilities. The loss function in Bayesian decision theory is defined as

$$L(a, \boldsymbol{\theta}) = \max_{a' \in \mathcal{A}} (U(a', \boldsymbol{\theta})) - U(a, \boldsymbol{\theta})$$

Note the analogy with the relation between loss and payoff function.

Then, the Bayes action with the use of data can be written

$$\begin{aligned} a^{(B)} &= \arg \max_{a \in \mathcal{A}} \left(\int_{\boldsymbol{\theta}} U(a, \boldsymbol{\theta}) q(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta} \right) = \\ &= \arg \max_{a \in \mathcal{A}} \left(\int_{\boldsymbol{\theta}} \left[\max_{a' \in \mathcal{A}} (U(a', \boldsymbol{\theta})) - L(a, \boldsymbol{\theta}) \right] q(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta} \right) = \\ &= \arg \min_{a \in \mathcal{A}} \left(\int_{\boldsymbol{\theta}} L(a, \boldsymbol{\theta}) q(\boldsymbol{\theta} | \mathbf{x}) d\boldsymbol{\theta} \right) = \arg \min_{a \in \mathcal{A}} \overline{L(a, q, \mathbf{x})} \end{aligned}$$

i.e. the action that minimises the expected posterior loss.

Example

Assume you are choosing between fixing the interest rate of your mortgage loan for one year or keeping the floating interest rate for this period.

Let us say that the floating rate for the moment is 4 % and the fixed rate is 5 %.

The floating rate may however increase during the period and we may approximately assume that with probability $g_1 = 0.10$ the average floating rate will be 7 %, with probability $g_2 = 0.20$ the average floating rate will be 6 % and with probability $g_3 = 0.70$ the floating rate will stay at 4 %.

Let a_1 = Fix the interest rate and a_2 = Keep the floating interest rate

Let θ = average floating rate for the coming period

a_1 = Fix the interest rate
 a_2 = Keep the floating interest rate
 θ = average floating rate (in %) for the coming period
 $P_g(\theta = 4) = 0.7; P_g(\theta = 6) = 0.2; P_g(\theta = 7) = 0.1$

A utility function can be defined from the potential changes in average floating rate:

$$U(a_1, \theta) = \begin{cases} 4 - 5 = -1 & \theta = 4 \\ 4 - 5 = -1 & \theta = 6 \\ 4 - 5 = -1 & \theta = 7 \end{cases} \quad U(a_2, \theta) = \begin{cases} 4 - 4 = 0 & \theta = 4 \\ 4 - 6 = -2 & \theta = 6 \\ 4 - 7 = -3 & \theta = 7 \end{cases}$$

Note that since a utility function is unique up to a linear transformation we do not need to include the mortgage loan amount into this function.

$$\Rightarrow$$

$$\overline{U(a_1, g)} = (-1) \cdot 0.7 + (-1) \cdot 0.2 + (-1) \cdot 0.1 = -1$$

$$\overline{U(a_2, g)} = 0 \cdot 0.7 + (-2) \cdot 0.2 + (-3) \cdot 0.1 = -0.7$$

$$\Rightarrow a^{(B)} = a_2$$

a_1 = Fix the interest rate
 a_2 = Keep the floating interest rate
 θ = average floating rate (in %) for the coming period
 $P_g(\theta = 4) = 0.7; P_g(\theta = 6) = 0.2; P_g(\theta = 7) = 0.1$

$$U(a_1, \theta) = \begin{cases} -1 & \theta = 4 \\ -1 & \theta = 6 \\ -1 & \theta = 7 \end{cases} \quad U(a_2, \theta) = \begin{cases} 0 & \theta = 4 \\ -2 & \theta = 6 \\ -3 & \theta = 7 \end{cases}$$

The loss function can be calculated as

$$L(a_1, \theta) = \begin{cases} 0 - (-1) = 1 & \theta = 4 \\ -1 - (-1) = 0 & \theta = 6 \\ -1 - (-1) = 0 & \theta = 7 \end{cases}$$

$$U(a_2, \theta) = \begin{cases} 0 - 0 = 0 & \theta = 4 \\ -1 - (-2) = 1 & \theta = 6 \\ -1 - (-3) = 2 & \theta = 7 \end{cases}$$

$$L(a, \theta) = \max_{a' \in \mathcal{A}} (U(a', \theta)) - U(a, \theta)$$

\Rightarrow

$$\overline{L(a_1, g)} = 1 \cdot 0.7 + 0 \cdot 0.2 + 0 \cdot 0.1 = 0.7$$

$$\overline{L(a_2, g)} = 0 \cdot 0.7 + 1 \cdot 0.2 + 2 \cdot 0.1 = 0.4$$

$$\Rightarrow a^{(B)} = a_2$$

The utility function used so far is linear in payoff.

a_1 = Fix the interest rate
 a_2 = Keep the floating interest rate
 θ = average floating rate (in %) for the coming period
 $P_g(\theta = 4) = 0.7; P_g(\theta = 6) = 0.2; P_g(\theta = 7) = 0.1$

Another utility function may be obtained by taking into account the mortgage loan amount and also include inconveniences that may appear in the process of changing the type of rate.

Least preferable payoff is -3% (occurs when keeping the floating rate and its average becomes 7%)

Most preferable payoff is 0% (occurs when keeping the floating interest rate and its average stays at 4% .)

Hence, let $U(a_2, \theta = 7) = 0$ and $U(a_2, \theta = 4) = 1$.

For any other consequence (payoff), c , the utility will be equal to the probability p_0 with which you are indifferent between c and $0 \cdot (1 - p_0) + 1 \cdot p_0 = p_0$

a_1 = Fix the interest rate
 a_2 = Keep the floating interest rate
 θ = average floating rate (in %) for the coming period
 $P_g(\theta = 4) = 0.7; P_g(\theta = 6) = 0.2; P_g(\theta = 7) = 0.1$

$$U(a_2, \theta = 7) = 0 ; U(a_2, \theta = 4) = 1$$

With the utility function used before

$$U(a_1, \theta) = \begin{cases} -1 & \theta = 4 \\ -1 & \theta = 6 \\ -1 & \theta = 7 \end{cases} \quad U(a_2, \theta) = \begin{cases} 0 & \theta = 4 \\ -2 & \theta = 6 \\ -3 & \theta = 7 \end{cases}$$

...we note that $U(a_2, 6) = -2 = (-3) \cdot \left(\frac{-2}{-3}\right) = (-3) \cdot \left(1 - \frac{1}{3}\right)$

Hence using the (0,1)-scale this utility is equal to 1/3. Correspondingly, the utilities $U(a_1, 4)$, $U(a_1, 6)$, and $U(a_1, 7)$ are all equal to 2/3 on this scale (as a consequence of the utility function being linear in payoff).

...but this also means that you would be indifferent between a certain increase by 1 %-unit of the rate on one hand and an increase by 3 %-units with probability 1/3 and no increase with probability 2/3 on the other.

a_1 = Fix the interest rate
 a_2 = Keep the floating interest rate
 θ = average floating rate (in %) for the coming period
 $P_g(\theta = 4) = 0.7; P_g(\theta = 6) = 0.2; P_g(\theta = 7) = 0.1$

$$U(a_2, \theta = 7) = 0 ; U(a_2, \theta = 4) = 1$$

However, would you instead be indifferent between an increase in interest rate of 1 %-unit on one hand and no increase with probability 4/5 and an increase of 3 %-units with probability 1/5 on the other hand – reflecting potential inconveniences with changing interest rate – your utilities $U(a_1, 4)$, $U(a_1, 6)$, and $U(a_1, 7)$ would all be equal to 3/4.

It might still be the case that your utility $U(a_2, 6)$ is 1/3.

With this utility function your expected utilities for the two actions are

$$\overline{U(a_1, g)} = (4/5) \cdot 0.7 + (4/5) \cdot 0.2 + (4/5) \cdot 0.1 = 0.8$$

$$\overline{U(a_2, g)} = 1 \cdot 0.7 + (1/3) \cdot 0.2 + 0 \cdot 0.1 \approx 0.77$$

$$\Rightarrow a^{(B)} = a_1$$

Example – medical treatment from a GP:s perspective

A general practitioner (GP) is supposed to state the diagnosis of a patient, who has declared some symptoms.

The GP sees three possibilities for the symptoms declared:

1. The patient has disease D1
2. The patient has disease D2
3. The patient has no disease

These are the three possible states of nature and can be denoted θ_1 , θ_2 and θ_3 respectively.

Now, the GP can choose to either...

...treat the patient for disease D1 (action a_1) ...or...

...treat the patient for disease D2 (action a_2) ...or...

...give no treatment (action a_3)

Note! This description of the decision problem is not realistic, but just to illustrate. A GP would do more than just select a treatment directly from the declared symptoms if there is more than one explanation to them.

Naturally, different combinations of action and state of the nature would lead to less or more preferable consequences.

Assume that

- the least preferable (worst) consequence is obtained when the patient has **disease D2** and no treatment is given
- the most (and equally) preferable consequences are obtained when the patient has **disease D1** and is **treated for D1**, and when the patient has **disease D2** and is **treated for D2** respectively.

We could also obtain this most preferable consequence when the patient has no disease and is not treated.

Hence,

$$\begin{aligned} U(a_3, \theta_2) &= 0 \\ U(a_1, \theta_1) &= U(a_2, \theta_2) = 1 \end{aligned}$$

a_1 = “Treat for disease D1”

a_2 = “Treat for disease D2”

a_3 = “Do not treat”

θ_1 = “The patient has disease D1”

θ_2 = “The patient has disease D2”

θ_3 = “The patient has no disease”

Now, to find the utilities for all other combinations of action and state of the nature (treatment and presence/absence of disease) the GP should do the following:

a_1 = “Treat for disease D1”

a_2 = “Treat for disease D2”

a_3 = “Do not treat”

θ_1 = “The patient has disease D1”

θ_2 = “The patient has disease D2”

θ_3 = “The patient has no disease”

Find $U(a_i, \theta_j)$ such that he is indifferent between

- I. Obtaining for certain the consequence of taking action a_i when the state of nature is θ_j , i.e. $c(a_i, \theta_j)$
- II. Obtain consequence $c(a_3, \theta_2)$ with probability $1 - U(a_i, \theta_j)$ and consequence $c(a_1, \theta_1)$ ($= c(a_2, \theta_2)$) with probability $U(a_i, \theta_j)$

A possible table of utilities may then be

	θ_1	θ_2	θ_3
a_1	1	0.3	0.6
a_2	0.4	1	0.6
a_3	0.1	0	0.9

	θ_1	θ_2	θ_3
a_1	1	0.2	0.6
a_2	0.4	1	0.6
a_3	0.3	0	0.9

a_1 = “Treat for disease D1”

a_2 = “Treat for disease D2”

a_3 = “Do not treat”

θ_1 = “The patient has disease D1”

θ_2 = “The patient has disease D2”

θ_3 = “The patient has no disease”

Some rational explanations for the numbers in the table may be:

- Treating when no disease is present may lead to inconveniences for the patient (undertaking some time-consuming activities and/or suffering from side-effects of prescribed drugs).
- Treating for one disease while the other is present would have similar consequences like in a) with the addition that the other disease would possibly not be cured

Now, the GP must assign probabilities to the three states of nature.

From epidemiological studies it might be known that

- the incidence rate of disease D1 is 5 in 100 persons per week
- the incidence rate of disease D2 is 1 in 100 persons per week.

Both these incidence rates are per week at the time of the year (think flus).

To assign prior probabilities from these epidemiological statistics we must assume that the GP is confident with restricting the set of possible states of nature to the three used here.

Calculating conditional incidence rates may then have given the following conditional prior probabilities (assuming a proportion of people with no disease among those visiting the GP):

$$P(\tilde{\theta} = \theta_1 \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) \approx 0.20, P(\tilde{\theta} = \theta_2 \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) \approx 0.01, \\ P(\tilde{\theta} = \theta_3 \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) \approx 0.79$$

The much higher probability of θ_1 can be due to experience that people with disease D1 tend to visit the GP much more often than people with disease D2.

a_1 = “Treat for disease D1”

a_2 = “Treat for disease D2”

a_3 = “Do not treat”

θ_1 = “The patient has disease D1”

θ_2 = “The patient has disease D2”

θ_3 = “The patient has no disease”

	θ_1	θ_2	θ_3
a_1	1	0.2	0.6
a_2	0.4	1	0.6
a_3	0.3	0	0.9

a_1 = “Treat for disease D1”

a_2 = “Treat for disease D2”

a_3 = “Do not treat”

θ_1 = “The patient has disease D1”

θ_2 = “The patient has disease D2”

θ_3 = “The patient has no disease”

$$P(\tilde{\theta} = \theta_1 \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) \approx 0.20, P(\tilde{\theta} = \theta_2 \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) \approx 0.01,$$

$$P(\tilde{\theta} = \theta_3 \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) \approx 0.79$$

Hence, the expected utilities of each action become

$$E(U(a_1, \tilde{\theta}) \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) = EU_1 = 1 \cdot 0.20 + 0.2 \cdot 0.01 + 0.6 \cdot 0.79 \approx 0.68$$

$$E(U(a_2, \tilde{\theta}) \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) = EU_2 = 0.4 \cdot 0.20 + 1 \cdot 0.01 + 0.6 \cdot 0.79 \approx 0.56$$

$$E(U(a_3, \tilde{\theta}) \mid \tilde{\theta} = \theta_1, \theta_2 \text{ or } \theta_3) = EU_3 = 0.3 \cdot 0.20 + 0 \cdot 0.01 + 0.9 \cdot 0.79 \approx 0.77$$

...and the optimal action using the *EU*-criterion is a_3 , i.e. “Do not treat”