

Preferences and utility functions

The probability values computed using the Bayesian belief networks provide information about relative probabilities of various events, possible effects of agent's actions, etc. However, they do not help much in making decisions based on this information.

For example, is an action plan that guarantees the achievement of 90% of the goals with a probability of 0.95 better than another plan that guarantees the achievement of 95% of the goals with a probability of 0.90?

Clearly, this does not depend on probabilities, but on what the agent prefers.

An intelligent agent must have a representation of her **preferences** for making decisions. To represent these preferences, we will use the idea of a **utility function** $U(S)$ which designates how much an agent prefers a certain state. Clearly, such utility is relative, and can only be determined for a specific agent.

The MEU principle

Consider an intelligent agent, having her preferences described with utilities, and using probabilities for evaluating the information she possesses and the possible outcomes of her actions. We will assume, that she acts rationally if she chooses actions with the highest expected utility averaged over all possible outcomes of such actions.

The **expected utility** $EU(s, a)$ of a non-deterministic action a applied by the agent in state s with the set of possible outcomes $\text{RESULT}(s, a)$ and a probability distribution $P(s'|s, a)$ of reaching a state $s' \in \text{RESULT}(s, a)$, is given as:

$$EU(s, a) = \sum_{s' \in \text{RESULT}(s, a)} P(s'|s, a) \times U(s')$$

The principle of **maximum expected utility (MEU)** says that a rational agent should choose the action that maximizes the agent's expected utility:

$$\text{action}(s) = \operatorname{argmax}_a EU(s, a)$$

Note that maximizing expected utility is just an assumption we make. A particular agent might instead want to minimize her worst possible loss, for another example.

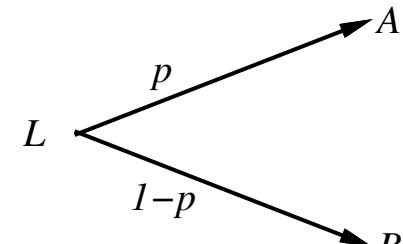
Lotteries and preferences

An intelligent agent must have a clear system of preferences, but it is not always specific and fully defined states that she must decide between. On the contrary, when making decisions under uncertainty the outcomes of any decision are usually uncertain either way. We will call the results for each action a **lottery** L , with a set of possible outcomes S_1, \dots, S_n that occur with probabilities p_1, \dots, p_n . Each outcome S_i of a lottery can be a specific state, or another lottery.

For example, a lottery L with two possible outcomes:

A with probability p and B with probability $1 - p$ we can denote as:

$$L = [p, A; 1 - p, B]$$



As a basis for selecting between lotteries, or states, the agent will use preferences:

$A \succ B$ – A is preferred over B

$A \sim B$ – there is no clear preference between A and B

$A \approx B$ – A is preferred over B or there is no preference

The axioms of the utility theory

We will assume that the agent's preferences have to satisfy the following properties, called the axioms of the utility theory:

orderability

$$(A \succ B) \vee (B \succ A) \vee (A \sim B)$$

transitivity

$$(A \succ B) \wedge (B \succ C) \Rightarrow (A \succ C)$$

continuity

$$A \succ B \succ C \Rightarrow \exists p [p, A; 1 - p, C] \sim B$$

(Continuity means, that if in some state B is in between A and C in preference, then there is some probability p for which the agent will be indifferent between getting B (for sure), and a lottery between A and C with probabilities p and $1 - p$.)

substitutability

$$A \sim B \Rightarrow [p, A; 1 - p, C] \sim [p, B; 1 - p, C]$$

(If an agent is indifferent between two lotteries, A and B , then it is also indifferent between two more complex lotteries, which are the same except that B is substituted for A in one of them. This holds regardless of the probabilities and the other outcomes in the lotteries.)

monotonicity

$$A \succ B \Rightarrow (p \geq q \Leftrightarrow [p, A; 1 - p, B] \succsim [q, A; 1 - q, B])$$

(If an agent prefers A over B , then for two lotteries with possible outcomes A and B it also prefers the one which gives the outcome A with a higher probability B .)

decomposability

$$[p, A; 1 - p, [q, B; 1 - q, C]] \sim [p, A; (1 - p)q, B; (1 - p)(1 - q), C]$$

(Compound lotteries can be reduced to simpler ones using the laws of probability.)

The meaning of the axioms

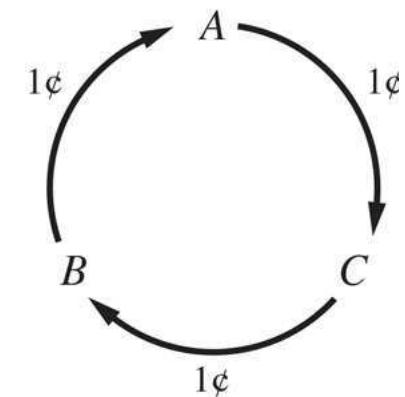
The axioms of the utility theory are constructed so, that violating them must be associated with an overtly irrational behavior.

Consider an agent with a preference system violating the transitivity axiom:
 $A \succ B \succ C$ and $C \succ A$:

If $B \succ C$, then an agent at C should be willing to pay 1 (euro)cent to reach B .

If $A \succ B$, then an agent at B should be willing to pay 1 (euro)cent to reach A .

If $C \succ A$, then an agent at A should be willing to pay 1 (euro)cent to reach C .



Such an agent could be induced into giving away all her money as a consequence of her preferences.

Utility functions

The axioms of the utility theory do not say anything about utilities, but only preferences. We will consider such preferences, satisfying the above axioms, to be the basic property of rational agents, and the starting point to determine utilities.

Fact: for a set of agent's preferences satisfying the preference axioms there exists a real-valued function U defined for the set of lotteries $U : \mathcal{L} \rightarrow \mathbb{R}$, such that:

$$U(A) > U(B) \Leftrightarrow A \succ B$$

$$U(A) = U(B) \Leftrightarrow A \sim B$$

For a lottery with the outcomes S_1, \dots, S_n and their respective probabilities p_1, \dots, p_n , this utility function assumes the value:

$$U([p_1, S_1; \dots; p_n, S_n]) = \sum_i p_i U(S_i)$$

In other words, once the utilities of the component lotteries are specified, and their probabilities are known, the utility of a compound lottery is determined.

The above fact stating that the utility function exists for an agent with a rational set of preferences does not allow to derive it directly, or even mean that it is unique. On the contrary, it is easy to see that the behavior of an agent with a utility function $U(S)$ would not change if her utility function was replaced with another one:

$$U'(S) = aU(S) + b$$

where a and b are constants with $a > 0$. So there are always multiple utility functions which equally well represent a specific agent's preferences.

Utility functions vs. preferences

In summary, the utility functions provide a technique by which we can model an intelligent agent's preferences, to mimic her decision making. The agent herself does not have to define or compute the utility function. The basis for her decision making are always her preferences.

Utility theory with respect to money

Consider the various possible utility functions for the states associated with having some specific sums of money. It is rational to assume that such functions will be monotonic for the specific (certain) values. But what about lotteries?

For an example, let us assume we won a T.V. competition and are offered a prize of one million dollars, or, alternatively, a flip of a coin. If the coin comes up tails we shall receive three million dollars, but if it comes up heads, then we receive nothing. Most people would choose to pocket the million in this case, but if we compute the **expected monetary value (EMV)** of the gamble variant then we obtain

$$\frac{1}{2}(\$0) + \frac{1}{2}(\$3,000,000) = \$1,500,000$$

while the EMV for the sure choice is \$1,000,000.

How can we interpret this result? Let us try to compute the utilities of the possible outcome states. Denoting by S_k the state in which we initially own $\$k$, we have:

$$\begin{aligned} EU(\text{coin_flip}) &= \frac{1}{2}U(S_k) + \frac{1}{2}U(S_{k+3,000,000}) \\ EU(\text{pocket_\$1M}) &= U(S_{k+1,000,000}) \end{aligned}$$

To determine the utility of being in possession of various sums of money we may assume, that the utility gain for the first million is higher then the difference between one and three millions, eg.: $U(S_k) = 5$, $U(S_{k+1,000,000}) = 8$, $U(S_{k+3,000,000}) = 10$. In this case we obtain $EU(\text{coin_flip}) = 7.5$ and we now have the basis to accept the certain million.

On the other hand, if we had many millions of dollars already, then these values could possibly come out different, and perhaps it could turn out that selecting the coin flip would turn out advantageous.

St.Petersburg paradox

Assume we are offered participating in a game (Bernoulli, 1738), in which we repeatedly flip a coin until it comes up heads. At that point the game ends and we win $\$2^n$ if the heads comes first on the n -th throw. Clearly, it is good to play this game. The question is, however, how much would we be willing to pay for the opportunity to play it? Since the probability of getting heads on the n -th throw is $1/2^n$, so:

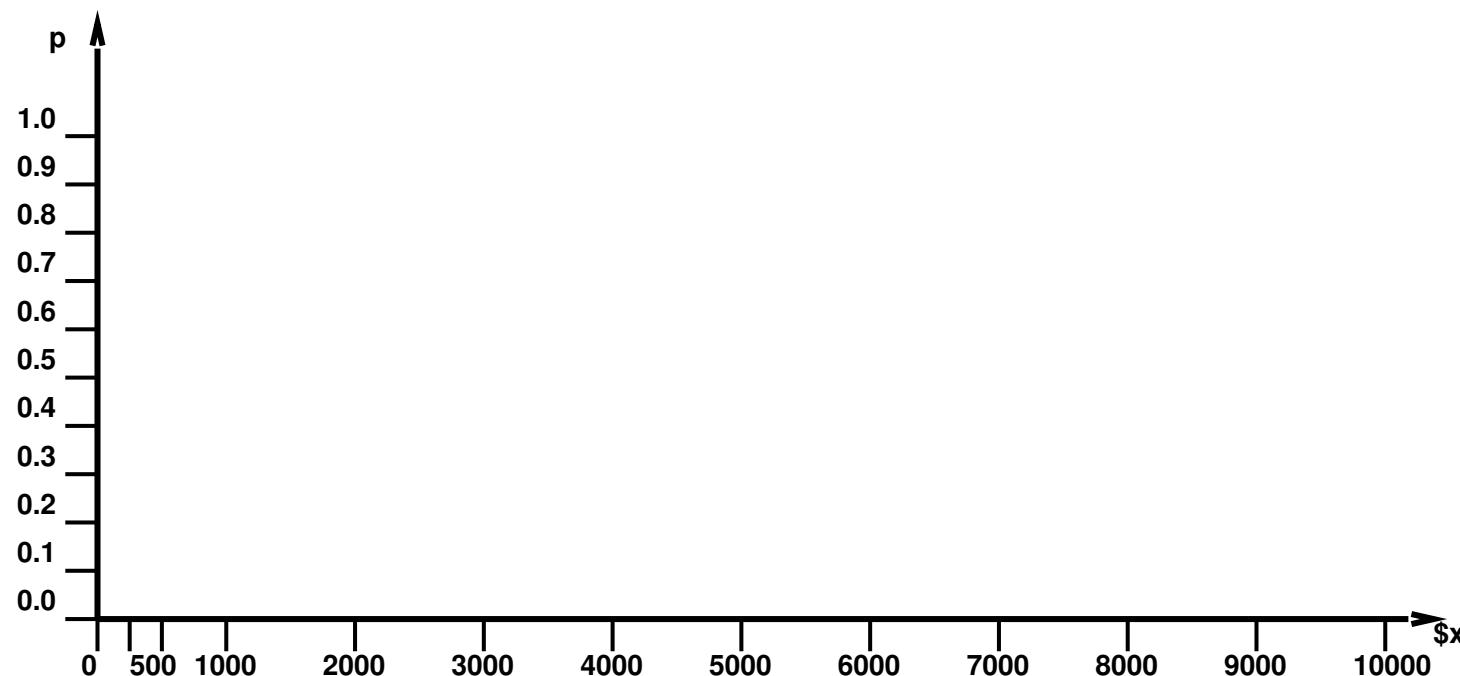
$$\text{EMV}(St.P.) = \sum_i P(\text{Heads}_i) \text{MV}(\text{Heads}_i) = \sum_i \frac{1}{2^i} 2^i = \frac{2}{2} + \frac{4}{4} + \frac{8}{8} + \dots = \infty$$

Does it mean that we should be willing to pay any (definite) amount of money? This does not sound rational, and so thought Bernoulli, who proposed using a logarithmic utility function for money, $U(S_k) = \log_2 k$, by which we obtain:

$$EU(St.P.) = \sum_i P(\text{Heads}_i) U(\text{Heads}_i) = \sum_i \frac{1}{2^i} \log_2 2^i = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots = 1$$

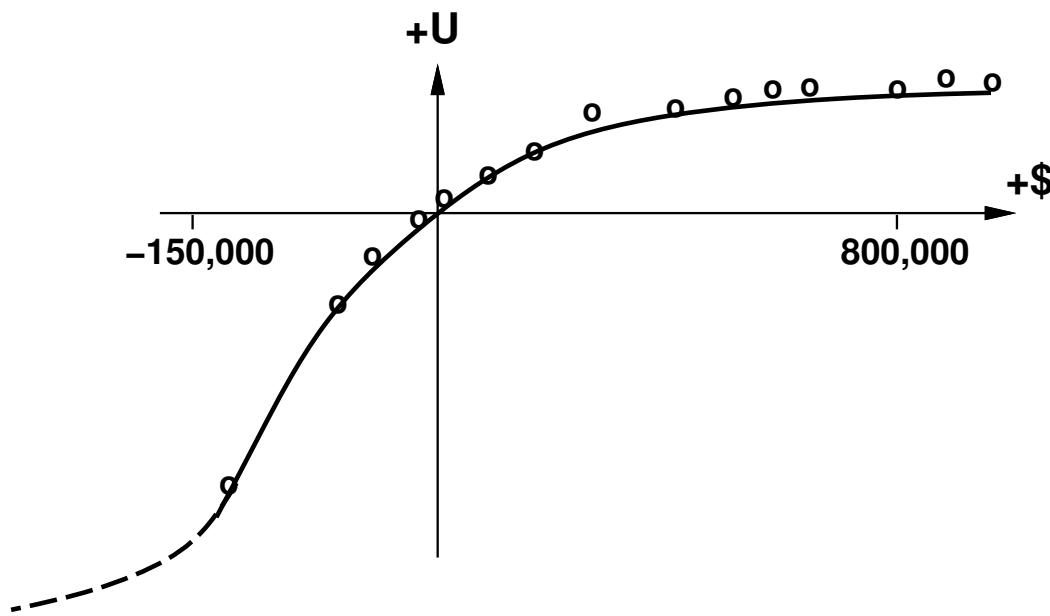
The students' utility of money

For each value x , let us determine, by voting, the probability p , at which half of the class prefers the lottery $[p, \$10,000; (1 - p), \$0]$ over the certain payout of $\$x$:



The curves of the utility of money

It is accepted, that the utility of money is a logarithm-like function, which is concave for the positive monetary values. In some research the specific utility function of money was determined for some person, which can be approximated by the formula $U(S_k) = -263.31 + 22.09 \log(k + 150,000)$ in the range between $-\$150,000$ and $\$800,000$:



For small negative values the function remains concave, since being in debt usually causes a state of panic with most people. However, for very large amounts of money the function will start to saturate, as the negative perception of a debt does not rise linearly with the debt amount.

Back to the positive monetary values, we may conclude that the agents having a concave utility function will generally prefer obtaining the expected value of a lottery (for sure) over participating in the lottery:

$$U(S_L) < U(S_{\text{EMV}(L)})$$

Such behavior may be termed **risk-averse**. A convex (range of) utility function may be termed **risk-seeking**. In any small range the utility function is typically almost linear, and for such function the behavior is termed **risk-neutral**.

Irrationality

Accepting a concave, logarithm-like function of the utility of money does not explain the whole psychology of economic decision-making by humans. It turns out they constantly violate the axioms of the utility theory. For example, given the choice between lotteries A or B , and C or D :

A : 80% of winning \$4000

B : 100% of winning \$3000

C : 20% of winning \$4000

D : 25% of winning \$3000

the majority of people choose B over A , but C over D . However, assuming $U(0) = 0$, the first choice implies $0.8 \times U(4000) < U(3000)$, while the second choice means exactly the opposite.

One possible explanation of this is simply inadequacy of the utility axioms to describe human behavior. But another possibility is that it does not take into account the **regret**. People know, that in the case of lottery A they will feel stupid if they choose, and then loose this lottery knowing, that they could choose the safe and profitable (though not as much) lottery B . In the case of C versus D this regret does not apply, so the behavior is rational.

Normalized utility functions

We have previously noted, that the utility axioms do not unambiguously determine the utility function from the preferences alone. An agent using the utility function: $U'(S) = k_1 + k_2 U(S)$ where k_1 and k_2 are constants ($k_2 > 0$), will behave and act identically to an agent with the function $U(S)$, if both agents use the same set of beliefs about their environment.

The utility function may then be scaled and shift up or down by a constant amount, and the behavior of an agent will not change. Agents may therefore use a **normalized utility function**.

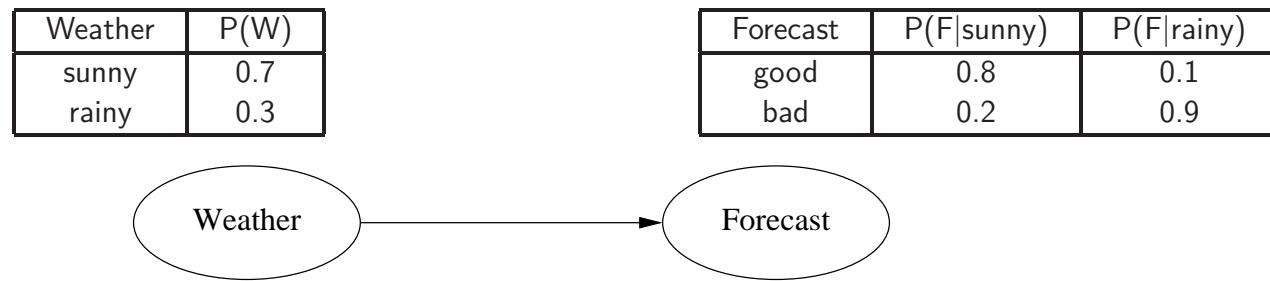
Let us denote by u_\perp the utility of the state of the „worst disaster” $u_\perp = U(S_\perp)$ for the original utility function $U(S)$, and by u_\top the utility of the state of the „highest reward” $u_\top = U(S_\top)$. Then for the normalized utility function U' we will take $U'(S_\perp) = 0$ and $U'(S_\top) = 1$, and the utilities of the intermediate states $U'(S)$ we will determine asking the agent to give the probability p , for which the agent is indifferent between the state S and a **standard lottery** $[p, S_\top; (1 - p), S_\perp]$

$$U'(S) = p, \text{ when } S \sim [p, S_\top; (1 - p), S_\perp]$$

Making decisions

The Bayesian belief networks provide answers about the probability distributions of any random variables, assuming prior knowledge about any combination of other variables. Having the distributions of utilities make it possible to use this knowledge, on the grounds of the MEU principle.

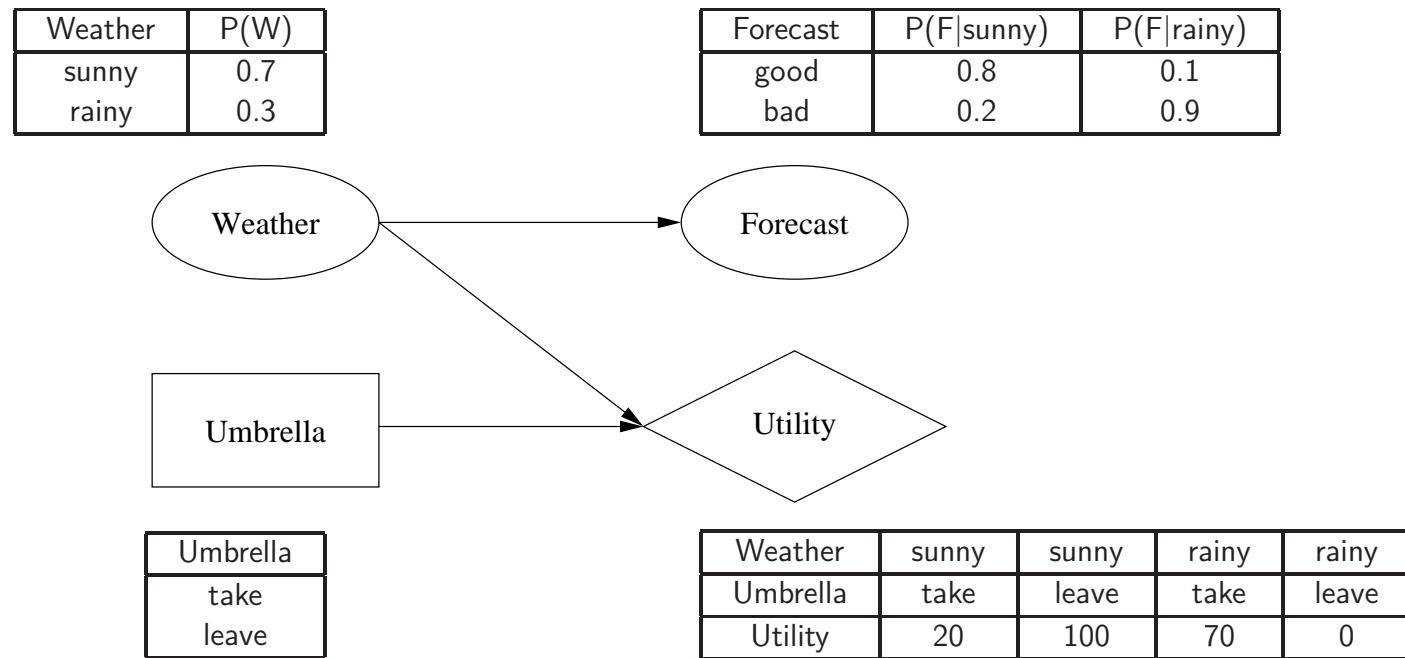
Consider an important decision: when leaving home, should I take an umbrella? It will be useful only in the rain. Otherwise it is clumsy to carry and can get lost. But how do I know if it will rain? A weather forecast can be helpful.



On a side note, this network provides an interesting example of a probabilistic dependency which goes backward to the chronological sequence. The weather influences the forecast, even though the forecast comes earlier. How is it possible? Ask a meteorologist.

Influence diagrams

Both the action to consider and the utilities of the situations can be added to the belief network graph as special action and utility nodes. There should be links from appropriate chance nodes and action nodes to the utility nodes.



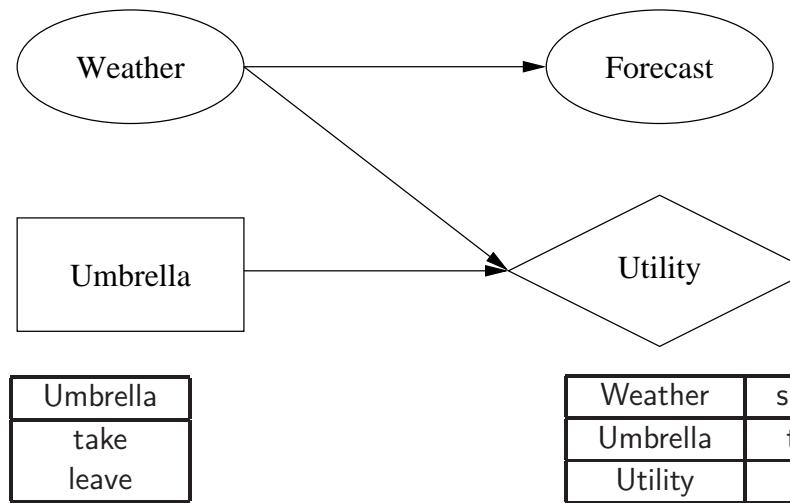
Such extended belief networks are called **influence diagrams**, or elsewhere **decision networks**. Some computer tools for building and processing belief networks also support influence diagrams like this.

Computing a decision

First consider the case, when there is no additional information.

| Weather | $P(W)$ |
|---------|--------|
| sunny | 0.7 |
| rainy | 0.3 |

| Forecast | $P(F sunny)$ | $P(F rainy)$ |
|----------|--------------|--------------|
| good | 0.8 | 0.1 |
| bad | 0.2 | 0.9 |



$$EU(\text{leave}) = P(\text{sunny}) * U(\text{leave, sunny}) + P(\text{rainy}) * U(\text{leave, rainy})$$

$$= 0.7 * 100 + 0.3 * 0$$

$$= 70$$

$$EU(\text{take}) = P(\text{sunny}) * U(\text{take, sunny}) + P(\text{rainy}) * U(\text{take, rainy})$$

$$= 0.7 * 20 + 0.3 * 70$$

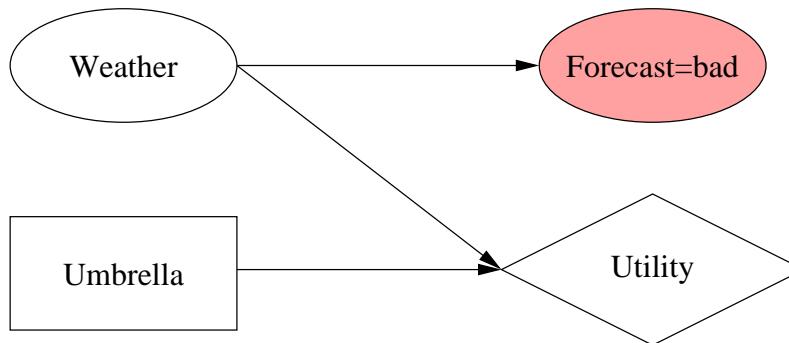
$$= 35$$

The utility of leaving the umbrella home is greater in this case.

Suppose now the forecast is for bad weather. Querying the network for the probability distribution of the weather gives: $P(\text{sunny}, \text{rainy}|\text{bad}) \approx (0.34, 0.66)$.

| Weather | $P(W)$ |
|---------|--------|
| sunny | 0.7 |
| rainy | 0.3 |

| Forecast | $P(F \text{sunny})$ | $P(F \text{rainy})$ |
|----------|---------------------|---------------------|
| good | 0.8 | 0.1 |
| bad | 0.2 | 0.9 |



| Umbrella |
|----------|
| take |
| leave |

| Weather | sunny | sunny | rainy | rainy |
|----------|-------|-------|-------|-------|
| Umbrella | take | leave | take | leave |
| Utility | 20 | 100 | 70 | 0 |

$$\begin{aligned}
 EU(\text{leave}|\text{bad}) &= P(\text{sunny}|\text{bad}) * U(\text{leave, sunny}) + P(\text{rainy}|\text{bad}) * U(\text{leave, rainy}) \\
 &= 0.34 * 100 + 0.66 * 0 \\
 &= 34
 \end{aligned}$$

$$\begin{aligned}
 EU(\text{take}|\text{bad}) &= P(\text{sunny}|\text{bad}) * U(\text{take, sunny}) + P(\text{rainy}|\text{bad}) * U(\text{take, rainy}) \\
 &= 0.34 * 20 + 0.66 * 70 \\
 &= 53
 \end{aligned}$$

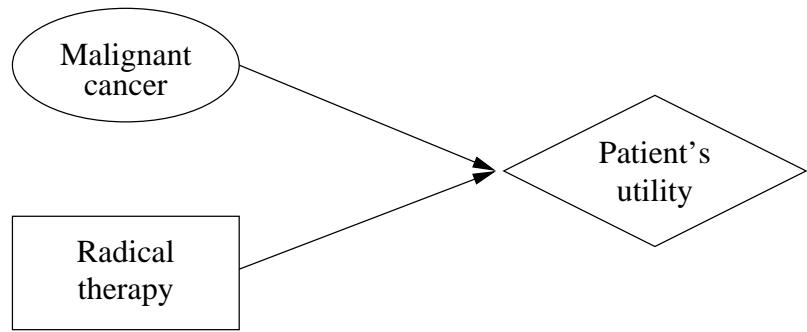
This time the decision should be to take the umbrella along.

Short review

1. Consider a patient with some probability of a malignant cancer and considered for a radical therapy. The situation is presented in the following influence diagram with a partially filled-in patient utility values.

(a) Fill in the missing utilities in a rational way, and justify why you chose such values.

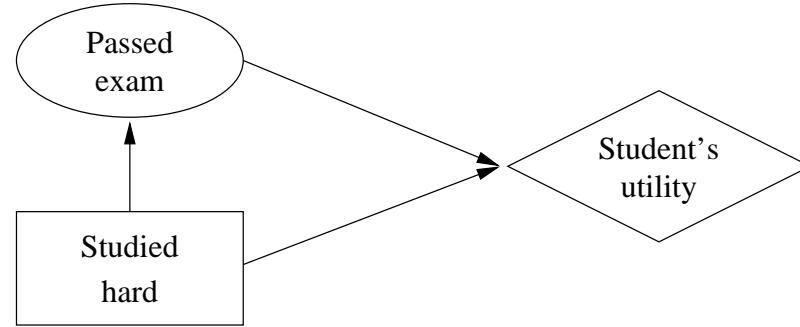
(b) Compute the cancer probability that switches the therapy decision.



| Malignant cancer | Radical therapy | Patient utility |
|------------------|-----------------|-----------------|
| N | N | 100 |
| Y | N | 0 |
| N | Y | |
| Y | Y | |

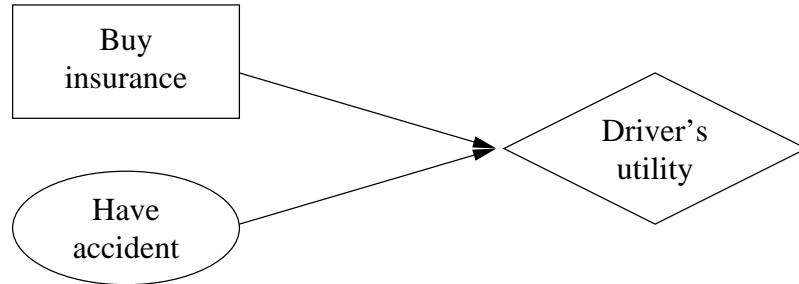
2. Consider a student approaching an examination, and can either study hard, or take it easy, counting on her good luck.

- Consider two different alternative attitudes toward life and hard work, and assign their corresponding utilities in the range from 0 to 100.
- For the selected course, determine the prior probability of passing the exam, and with the two described attitudes compute the study decision according to the MEU principle.
- Next, treating the study decision as an independent random variable, determine the conditional probability distribution for passing the exam selected before. Set the prior probability of studying and compute the expected values of the utility for the two described attitudes.



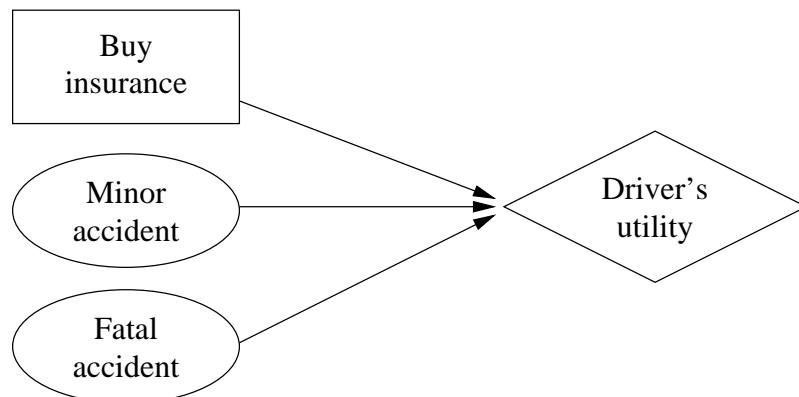
| Passed exam | Studied hard | Student utility |
|-------------|--------------|-----------------|
| N | N | |
| Y | N | |
| N | Y | |
| Y | Y | |

3. Consider purchasing car insurance. It is generally better to have insurance, although it can be questioned in case we did not have an accident. Let us assume the following model of the situation:



| Buy insurance | Have accident | Driver's utility |
|---------------|---------------|------------------|
| N | N | 100 |
| Y | N | 100 |
| N | Y | 5 |
| Y | Y | 25 |

Now, “accident” is a generalization of different kinds of accidents, with varying consequences. We want to create a more detailed model, which distinguishes two distinct types of accidents: a minor collision, and a fatal accident in which people are seriously injured or killed. Make some reasonable assumptions and fill in the missing utilities in a way that is consistent with the previous model.



| Buy insurance | Minor accident | Fatal accident | Driver's utility |
|---------------|----------------|----------------|------------------|
| N | N | N | 100 |
| Y | N | N | 100 |
| N | Y | N | |
| Y | Y | N | |
| N | N | Y | |
| Y | N | Y | |
| N | Y | Y | |
| Y | Y | Y | |

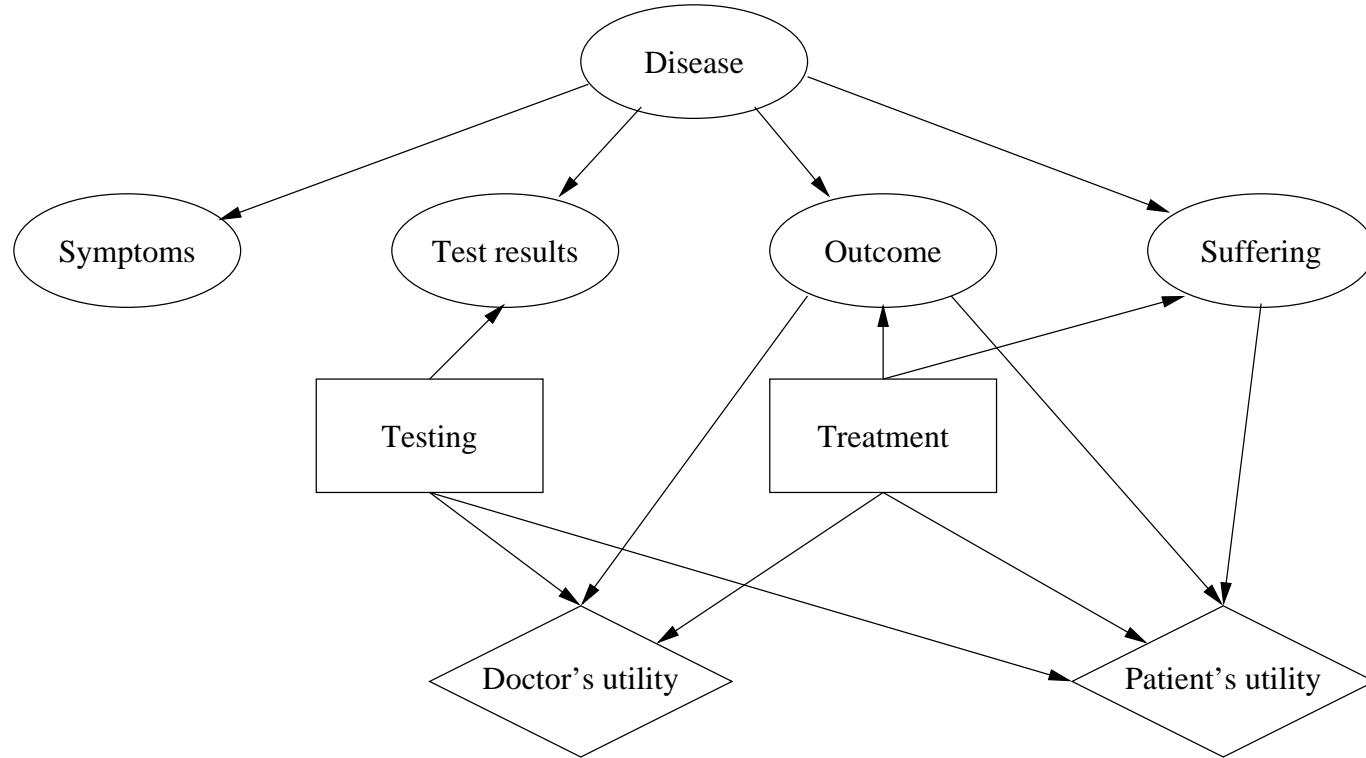
Multiple actions and multiple utilities

There are typically many chance nodes in a belief network, since this is essentially its primary purpose — to simplify the calculations of conditional probabilities in complex cases.

However, it is perfectly normal to have exactly one decision and one utility node in a network. This is because a belief network typically represents a situation to make a single decision. Also, to select the decision unambiguously, it is best to have a single criterion.

But multiple action nodes are perfectly possible in an influence diagram. They represent a situation, when the agent must make just one of these decisions, or must make in one step a joint decision, based on the information in the network. They do not permit to compute a sequence of decisions, when the consequences of one action affect the choice of another.

On the other hand, multiple utilities, if present, must be aggregated using one of the multiattribute utility models described below.



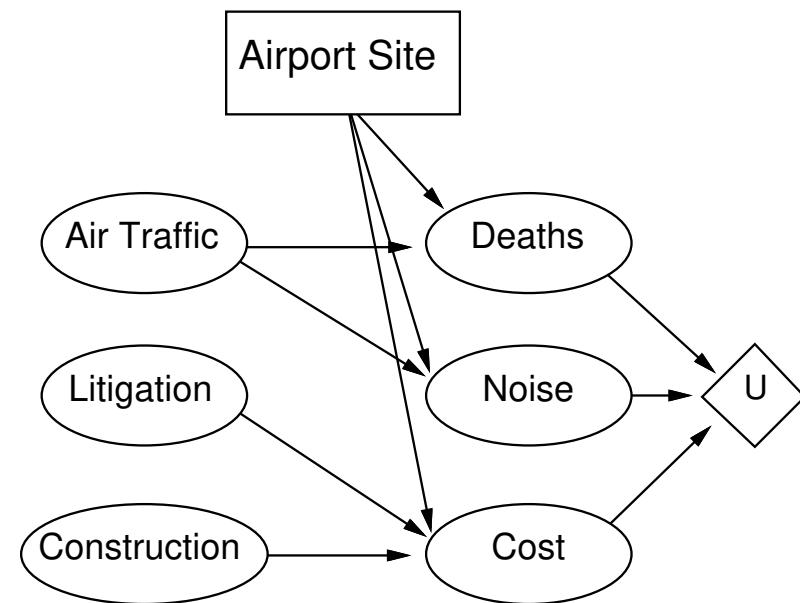
From the above network a doctor can prescribe additional tests, or make a treatment decision, regardless of whether test results are available or not. Prescribing tests and deciding a treatment based on the test results cannot be computed in one step.

Let's note the separate utilities computed from the patient's and the doctor's point of view. While the final outcome is important for both, they consider additional factors, but different, and in different ways (eg. the patient's suffering and the doctor's prestige). The decisions based on them could be different.

Multiattribute utilities

In many practical problems one needs to consider many criteria for the correctness of the decisions being made. In other words, multiple utility functions may exist taking into account different attributes of the states under consideration.

For an example, consider a problem of siting a new airport for some city. Many factors need to be taken into account: the cost of the land, increased road traffic, local weather conditions, and other threats. For each possible site one can determine the critical attributes: the total cost, the increase in accidents (deaths), the noise, etc.



Multiattribute utility functions

We will try to describe a model for making rational decisions for the multiattribute problems. We shall denote the attributes by X_1, X_2, \dots and their values by x_1, x_2, \dots . For simplicity we assume that the higher attribute values correspond to higher utility values, with respect to the given attribute.

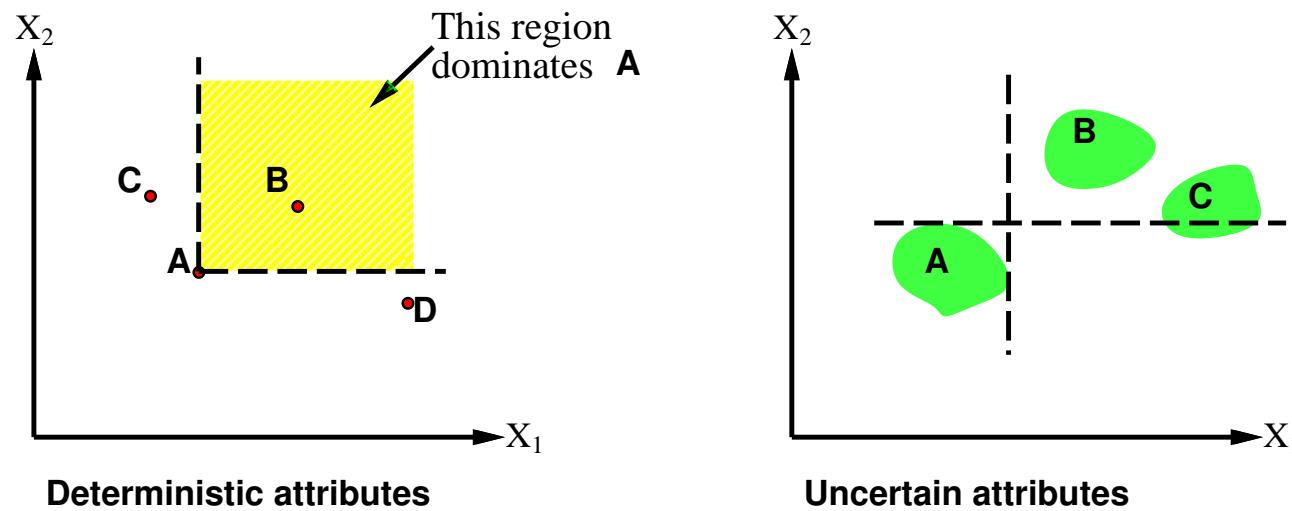
Assuming there exist utility functions for all the attributes: $f_1(x_1), f_2(x_2), \dots$ we may try to express the global utility function with the formula:

$$U(x_1, x_2, \dots x_n) = f(f_1(x_1), f_2(x_2), \dots, f_n(x_n))$$

However, in many cases the $f()$ function is hard to express. Therefore we shall first examine some special cases.

Strict dominance

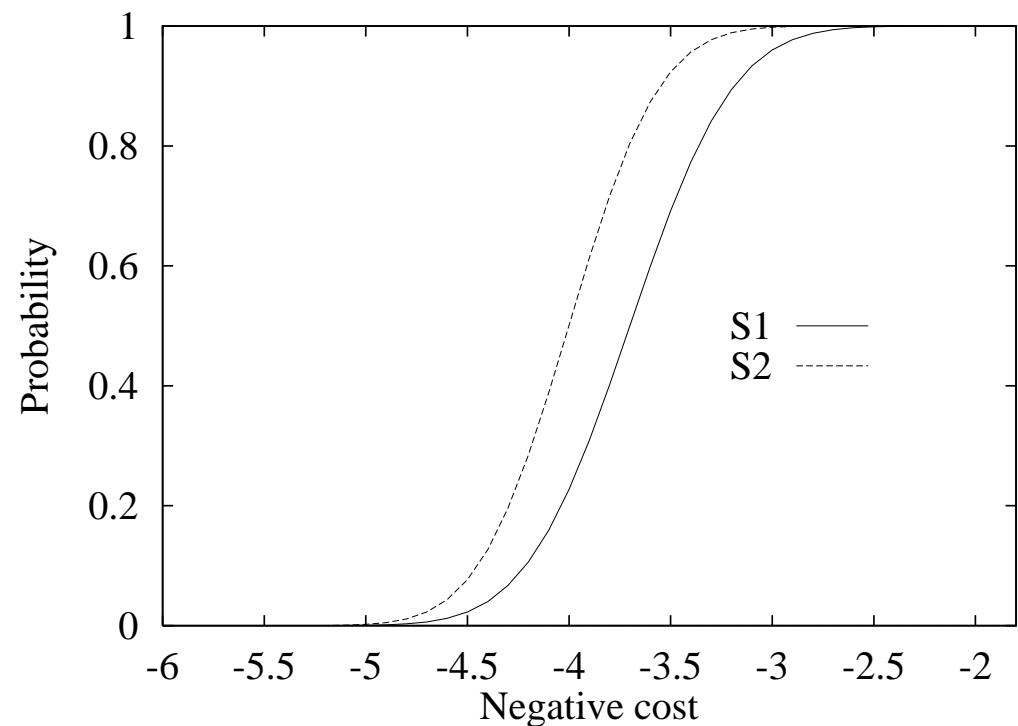
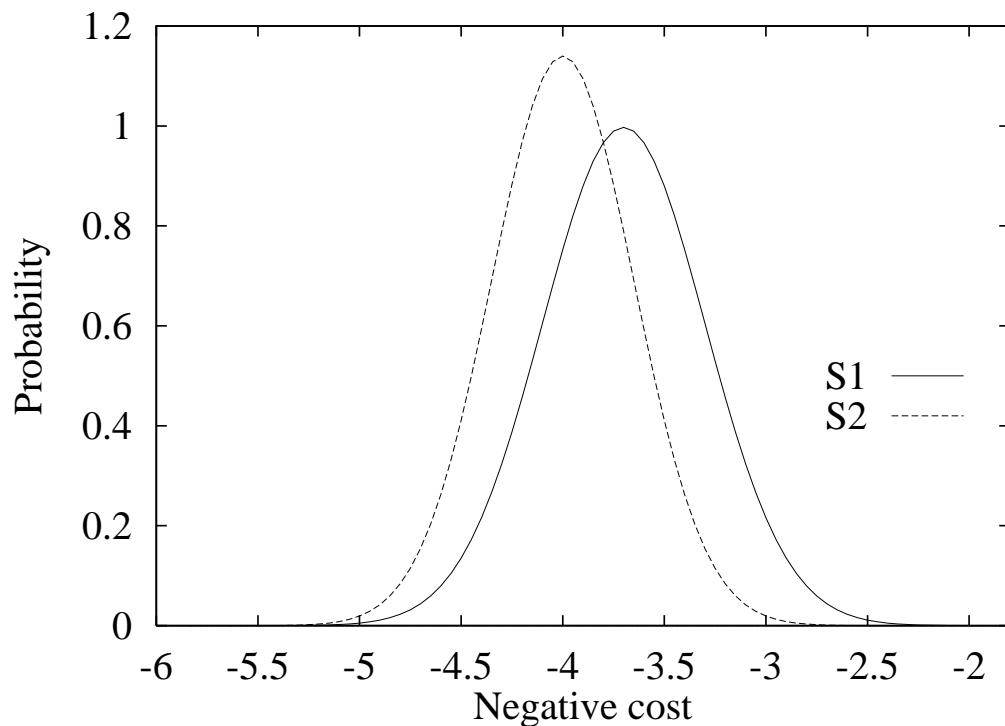
A **strict dominance** case we will call a situation where one choice B has all the attributes better than another choice A . For example, one siting B for the airport may be cheaper, safer, and cause less negative impact for the environment and people than A . In such case, we can strike the solution A right off from further considerations. However, it is not possible to make such decisions between possibilities A and C or A and D (diagram on the left):



In general, definite values of the attributes may not be known. But a strict dominance may also occur in the probabilistic case, when all choices result in certain probability distributions of the attribute values (diagram on the right).

Stochastic dominance

Strict dominance may not occur frequently in practical problems. In some cases one can take advantage of a more general case of **stochastic dominance**.



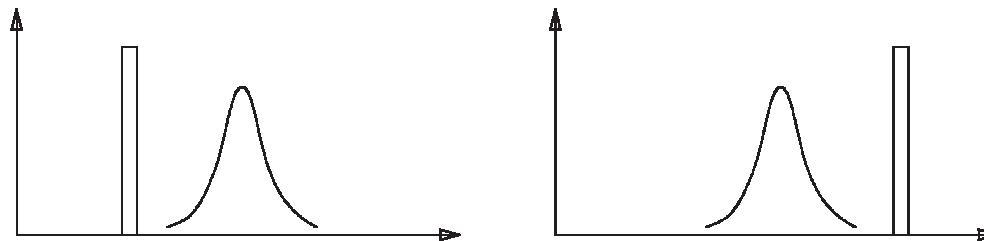
For example, if the siting cost of an airport at S_1 was estimated with the normal distribution with the expected value of \$3,700M and the variance of \$400M, and at site S_2 with the normal distribution with the expected value of \$4,000M and the variance of \$350M, then S_1 dominates stochastically S_2 , which can be seen in the cumulative distributions.

If the agent considers two possible actions A_1 and A_2 , which lead to probability distributions $p_1(x)$ and $p_2(x)$ for the attribute X , then we can say that A_1 dominates stochastically A_2 if:

$$\forall x \int_{-\infty}^x p_1(x')dx' \leq \int_{-\infty}^x p_2(x')dx'$$

In the extreme case, for one of the cases (outcome of the action A_2) the value x_2 of attribute X may even be certain (ie. with 100% probability), which does not mean that A_2 should be preferred, nor that preferred should be action A_1 with the outcome as a probability distribution for the attribute X .

Depending on the specific distribution the uncertain action may stochastically dominate the certain action, or the other way around.



Multiattribute utilities — the deterministic case

In a general case no dominance may occur, and computing multiattribute preferences is harder. However, it can often happen that the state attributes X_1 and X_2 are **preferentially independent** of X_3 , if the preferences between $\langle x_1, x_2, x_3 \rangle$ and $\langle x_1', x_2', x_3 \rangle$ do not depend on the specific value of x_3 .

If each pair of the attributes X_i and X_j is preferentially independent of each of the remaining attributes X_k , then this set of attributes has the property of **mutual preferential independence** (MPI).

It turns out, that in such case there exist an **additive value function** describing an agent's preferences:

$$V(S) = \sum_i V_i(X_i(S))$$

In many practical cases this function can properly model the preferences and make correct decisions.

Multiattribute utilities — the probabilistic case

The MPI property can be extended to lotteries: a set of attributes **X** is **utility independent** from the set of attributes **Y**, if the preferences between the lotteries on the attributes from **X** are independent of the specific values of the attributes from **Y**. A set of attributes is **mutually utility-independent** (MUI), if each subset of its attributes is utility independent from the other attributes.

For attributes which are MUI the agent's behavior can be described using a **multiplicative utility function**, which for a three attribute case can be written as:

$$\begin{aligned} U(S) &= k_1 U_1(X_1(S)) + k_2 U_2(X_2(S)) + k_3 U_3(X_3(S)) \\ &+ k_1 k_2 U_1(X_1(S)) U_2(X_2(S)) + k_2 k_3 U_2(X_2(S)) U_3(X_3(S)) + k_3 k_1 U_3(X_3(S)) U_1(X_1(S)) \\ &+ k_1 k_2 k_3 U_1(X_1(S)) U_2(X_2(S)) U_3(X_3(S)) \end{aligned}$$

In some specific cases there also exists a purely additive utility function.

The value of information

So far we assumed that our hypothetical intelligent agent has unlimited access to all the available knowledge about the world. This assumption is not very realistic. In the practical cases of making decisions one of the important and very difficult questions to answer is which questions pertaining the problem under considerations should the agent try to find answers to.

For example, consider a doctor, who does not have all the information about the patient when he becomes familiar with the case. Prescribing some tests to gain more information about a patient may lead to a better diagnosis statement, but on the other hand it is costly, and also delays the commencement of the treatment.

The importance of information depends on two factors: (1) whether various possible outcomes will significantly affect the decision, and (2) the probabilities of different outcomes.

The value of information theory allows to make decisions about which information the agent should collect.

Value of information — an example

Assume a drilling company considers the possibility of buying drilling rights in one of n identical ocean blocks. Further assume, that it is known that exactly one of the blocks contains oil worth C dollars, and that the price of one block is C/n . Note that the expected value of the profit EP from this transaction is 0:

$$EP = \frac{1}{n} \left(C - \frac{C}{n} \right) + \frac{n-1}{n} \left(-\frac{C}{n} \right) = 0$$

Suppose now, that a geologist comes along, claiming he knows for sure whether one specific block has oil or not. What could be the value of such information?

Let us consider all the cases. With probability $\frac{1}{n}$ the selected block has oil, and in such case the company will surely buy it and make C less C/n for the drilling rights. With probability $\frac{n-1}{n}$ the block does not have oil, and, knowing that, the company will buy another block which may have oil with probability $\frac{1}{n-1}$, to make an expected $C/(n-1)$ dollars, again less C/n :

$$EP' = \frac{1}{n} \left(C - \frac{C}{n} \right) + \frac{n-1}{n} \left(\frac{C}{n-1} - \frac{C}{n} \right) = \frac{C}{n}$$

Conclusion: information has value, which in this case is equal to the worth of the drilling rights to the block.

Value of perfect information

Assume the current knowledge is E and the agent's goal is to select the best action α from all possible actions A . We determine the expected value of the utility of this action (averaged over the action's possible outcomes $Result_i(A)$):

$$EU(\alpha|E) = \max_A \sum_i U(Result_i(A))P(Result_i(A)|Do(A), E)$$

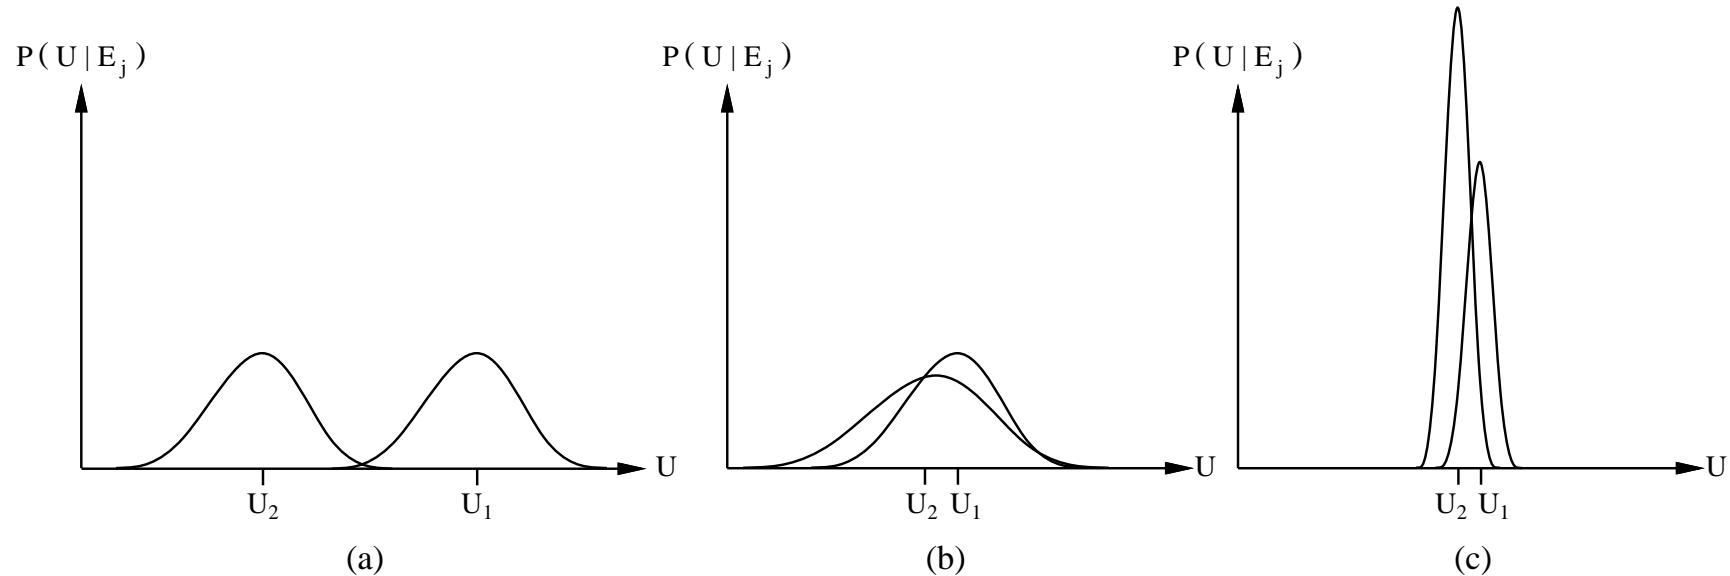
If the agent gained additional knowledge about the value of some random variable E_j then the expected utility value of such action would be: α_{E_j} :

$$EU(\alpha_{E_j}|E, E_j) = \max_A \sum_i U(Result_i(A))P(Result_i(A)|Do(A), E, E_j)$$

However, since the E_j is a random variable with an unknown value, we must base our decision of whether to acquire the knowledge of its value, by taking into account all the possible values, and what we know about them already. The **value of perfect information (VPI)** about a variable E_j can be computed as:

$$VPI_E(E_j) = \left(\sum_k P(E_j = e_{jk}|E)EU(\alpha_{e_{jk}}|E, E_j = e_{jk}) \right) - EU(\alpha|E)$$

For example, consider an agent having the choice of two actions A_1 and A_2 , and their utilities with probability distributions with expected values U_1 and U_2 . Gaining some additional information E_j will cause the expected utilities of these actions to U'_1 i U'_2 . Knowing the values of: U_1 , U_2 , and U'_1 , U'_2 , we may make the decision of whether it is worthwhile to acquire the information.



In the (a) case the difference in expected utilities of these actions is large, so having additional information may not make the agent change her choice of action; thus the value of information is nil. In (b) the difference in the expected values is small, but their variances are large, so the additional information may, by removing some uncertainty, help choose a significantly better action. In (c) the variances are small, so the uncertainty is small, as is the difference between the two actions, and acquiring new information may again not be worthwhile.

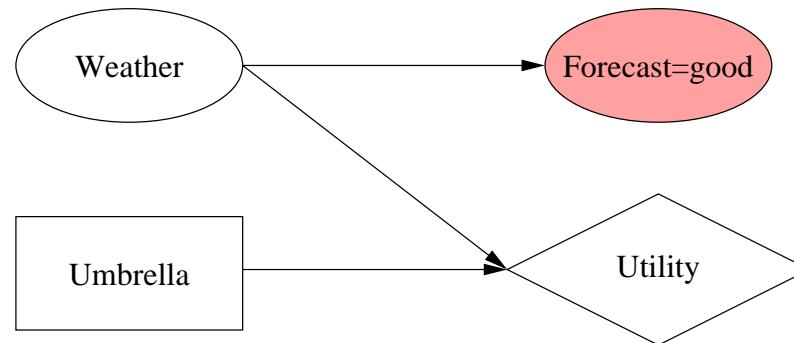
Example: the value of a weather forecast

Recall the weather and umbrella example considered before. We had computed:

$$MEU(\text{Umbrella}) = \max_a EU(a) = 70 \quad (\text{see slide 19})$$

$$MEU(\text{Umbrella}|\text{bad}) = \max_a EU(a|\text{bad}) = 53 \quad (\text{see slide 20})$$

We can also compute the utility of the best action for the good weather case (which should obviously be: „leave” since this decision prevailed even for zero information case):



$$MEU(\text{Umbrella}|\text{good}) = \max_a EU(a|\text{good}) = 95$$

To compute the value of a weather forecast, we need to know the probability distribution of the *Forecast* variable. This can be obtained by querying the belief network: $P(\text{good}, \text{bad}) \approx (0.59, 0.41)$.

$$VPI_E(E_j) = \left(\sum_k P(E_j = e_{jk} | E) EU(\alpha_{e_{jk}} | E, E_j = e_{jk}) \right) - MEU(\alpha | E)$$

$$\begin{aligned} VPI(\text{Forecast}) &= P(\text{good})EU(\alpha_{\text{good}} | \text{good}) + P(\text{bad})EU(\alpha_{\text{bad}} | \text{bad}) - MEU(\alpha) \\ &= P(\text{good})MEU(\text{good}) + P(\text{bad})MEU(\text{bad}) - MEU(\alpha) \\ &= 0.59 * 95 + 0.41 * 53 - 70 \\ &= 7.78 \end{aligned}$$

For the utility distribution defined for this problem the value of the weather forecast is 7.78, expressed in the utility units. If we could purchase a forecast with the reliability as considered in the computation, for a price not exceeding this value, then it would be profitable to do so, to make better umbrella decisions.

Properties of the value of information

Fact: the value of information is non-negative.

It can reach zero if, for example, the knowledge of other facts renders some information useless. So the value of information is not additive.

In turn, the value of knowledge of the values of two random variables does not depend on the order of acquiring these data. If the agent knows the values of E_i and E_j then the conclusions it may make from this combined knowledge does not depend on when and in which order it was acquired.

$$VPI_E(E_j, E_k) = VPI_E(E_j) + VPI_{E,E_j}(E_k) = VPI_E(E_k) + VPI_{E,E_k}(E_j)$$

The value of two different pieces of information may be different, and the agent might try to evaluate which information would bring the higher profit (less their cost).

Shortsighted information-gathering agent

An intelligent agent should ask questions of the user in a reasonable order, avoid asking irrelevant questions, take into account the value of information relative to its cost, and stop asking questions when it no longer makes sense. These capabilities can be achieved using the value of information as a guide.

A reasonable algorithm for an agent: select an information, which has the highest net value (its value of perfect information less the cost to acquire it), and if this value is positive then opt for acquiring this information. If it is negative, then proceed to the proper activities (non information gathering).

This agent algorithm is shortsighted, or **myopic**, since it is based on the consideration of acquiring only one variable, when gaining several data in turn could give the agent more advantage. This is somewhat analogous to greedy search algorithms, and likewise may be successful in some practical cases.

In a general case, an intelligent agent should consider various subsets of random variables, and corresponding information requests with their possible outcomes.

Short review

1. For the cancer patient case from the question on slide 21, and the utilities proposed therein, compute the value of perfect information about the cancer variable.