Chapter 2

Decision Theory

2.1 Introduction

Everyday one faces problems which call for decisions.

Example 2.1. At a summer evening, a person makes plans for the next day, and chooses between going to the beach or to go shopping. If next day is sunny, the beach is the best choice, but if the weather is less good, it is best to go shopping. How should the person decide what plans to make (if really having to plan ahead, already the evening before)?

The difficulty in the above example is that the weather for tomorrow cannot be known with complete certainty, so that it is a question of a *decision under uncertainty*. Decision theory makes science of the art of making such decisions.

2.2 Decision theory and utility

Let us consider a situation in which we can take different actions, denoting the set of all possible actions available to us by $\mathcal{A}$. In the example above this set is quite small, it consists of only two possibilities: $a_1 =$ "go to the beach" and $a_2 =$ "go shopping". However, observe that generally the number of actions to choose from need not be finite and the set $\mathcal{A}$ can even be continuous. Let us also assume that there is a true state of nature, which we have no knowledge of, but which will determine how much we gain or lose by choosing one action out of those in $\mathcal{A}$. Call this set of all possible "nature states" for Θ. In the example this set again contains only two values: $\theta_1 =$ "sunny tomorrow" and $\theta_2 =$ "rains tomorrow". Similarly to $\mathcal{A}$, in more complex situations this set may be infinite and continuous. Finally, assume that for each a and θ it is possible for us to determine a function $u(a, \theta)$ that describes our gain if we perform a given that θ is true. This function, called utility, is a *subjective* measure of gain that is constructed according to a persons' views and beliefs and can not be "right" or "wrong". It is simply something that we choose to use for this particular situation. For the example above it can look something like in Table 2.1.

		States of nature	
		θ = "sunny"	θ = "rainy"
Actions	a_1 = "beach"	10	-5
	a_2 = "shopping"	2	6

Table 2.1: The value of a utility function $u(a, \theta)$ for Example 1.

Given θ it easy to see which a we should choose. Of course, the whole point is that we must pick an a while not knowing θ, that is we have to make a decision under uncertainty. So how, exactly, do we do that?

The answer is, we have to make a guess about θ. While we do not know its exact value, we may have reasons to believe (for example, from past experience) that one state of nature is more likely than another. If it is sunny today, it is reasonable to conclude that it is more probable that it will be sunny tomorrow than that it will rain. That is, according to our belief, different values of θ can be true with different probabilities and so we can define a probability distribution on the space Θ. This will enable us to talk about the *expected* utility $U_\pi(a)$ that will describe how much we expect to gain by choosing a given a probability distribution $\pi(\theta)$ on all the possible states of nature. Formally, we have

$$U_\pi(a) = \mathbf{E}_\Theta[u(a, \theta)] = \begin{cases} \int_\Theta u(a, \theta)\pi(\theta)d\theta & \text{for a continuous state space} \\ \sum_\Theta u(a, \theta)\pi(\theta) & \text{for a discrete state space} \end{cases}$$

For our simple example, the probability distribution $\pi(\theta)$ may be $\{\pi(\theta_1), \pi(\theta_2\} = \{0.8, 0.2\}$ leading to expected utilities

$$\begin{cases} U_\pi(a_1) & = 10 * 0.8 + (-5) * 0.2 = 7 \\ U_\pi(a_2) & = 2 * 0.8 + 6 * 0.2 = 2.8 \end{cases}$$

Clearly, given these particular $u(a, \theta)$ and $\pi(\theta)$ action a_1 = "go to the beach" is to be preferred.

2.2.1 Utility of money

Let us say that you have an amount of money S at your disposal. Assuming that you don't spend it right away, you can think of two possible actions: a_1 = "store it away" and a_2 = "invest it stocks". In the first case the amount S will stay the same tomorrow as it is today, while in the second case it may increase but it may also decrease. Which of the actions do you prefer? The answer depends on how much you value the possible monetary gain and loss, that is on your utility.

The space of the states of nature Θ in this case consists of all possible differences in the prices of the stock between tomorrow and today. That is, Θ is continuous. However, small changes should be far more likely than large ones. Assume, for the sake of simplicity, that we have reason to believe that the probability of a negative difference is the same as a positive one. Then it is natural to choose our prior distribution $\pi(\theta)$ to be symmetric with light tails, centered around 0. The Normal distribution fulfills the requirements, so we let $\pi(\theta)$, again for simplicity, be $N(0,1)$.

Next, we choose a utility $u(a_2,\theta)$. Consider three choices: $u_1(a_2,\theta) = 1 - e^{-\theta}$ and $u_2(a_2,\theta) = e^{\theta} - 1$ and $u_3(a_2,\theta) = \theta$. The utilities are plotted in Figure 2.1. The third utility, the straight line, represents neutral behaviour. Using it reflects the belief that for you the worth of money is proportional (or, in this case, exactly equal) to its amount. The other two, $u_1(a_2,\theta)$ and $u_2(a_2,\theta)$ describe the so called risk averse and risk seeking behaviours. A risk averse utility is concave and favors gaining a small amount of money with certainty. It will also put a large negative weight on loosing your capital. According to the risk seeking utility, on the other hand, you find only large gains to be useful and do not care too much if your original capital S is lost.

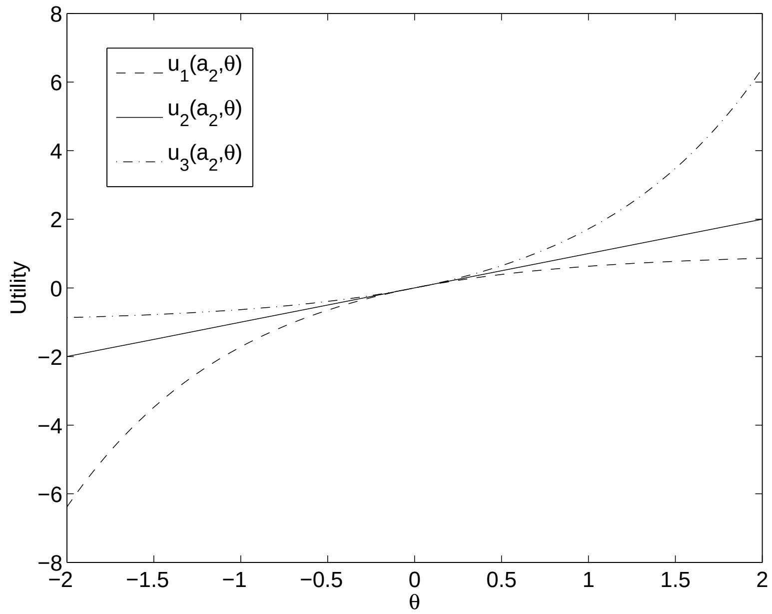

Figure 2.1: The risk seeking (dash-dotted), risk neutral (solid) and risk averse (dashed) utilities.

The expected utilities in those three cases are going to be

$$U_\pi(a_2)_1 = \int_{-\infty}^{\infty} (1 - e^{-\theta})\pi(\theta)d\theta = -0.65$$
$$U_\pi(a_2)_2 = \int_{-\infty}^{\infty} (e^{\theta} - 1)\pi(\theta)d\theta = 0.65$$
$$U_\pi(a_2)_3 = \int_{-\infty}^{\infty} \theta\pi(\theta)d\theta = 0$$

Now we can answer the question of whether we should save our money or invest it. If you are a risk averse person, then your expected utility is going to be smaller than 0 (which is, of course, the expected utility of a_1) and choosing a_1 is preferable. If you are a risk seeker, then you should invest. Finally, if you are of the neutral disposition then a_1 and a_2 are equivalent in the sense that they have equal expected utilities.

2.3 Autocorrelation

Let $S(t)$ denote the price of a stock at time t and $\{S(t)\}_{t\in T}$ a collection of prices for different time points. This collection can be thought of as a non-stationary stochastic process, which is really hard to model. Instead, one tends to look at the standardized differences between the consecutive time points, that is *returns* and *log-returns*, defined as

$$\begin{aligned} X(t) &= (S(t) - S(t-1))/S(t-1) && \textit{returns} \\ X(t) &= \log(S(t)/S(t-1)) = \log(S(t)) - \log(S(t-1)) && \textit{log-returns} \end{aligned}$$

Based on historical evidence those standardized differences are believed to be independent from one day to another and normally distributed. The independence assumption means that no one should be able to predict the stock market tomorrow and be able to make money without risk.

The independence assumption, although generally believed to be valid, needs not always be correct. In order to see whether it is plausible, we can try to model this possible dependence with an *autocorrelation function* defined as

$$r_X(s,t) = \frac{\mathbf{Cov}\{X(s), X(t)\}}{\sqrt{\mathbf{Var}\{X(s)\}\,\mathbf{Var}\{X(t)\}}}.$$

where $X(t)$ and $X(s)$ correspond to the values of a stochastic process at time points t and s. Note that $r_{X_1}(t,t) = \mathbf{Var}\{X(t)\}$. Under the assumption of *stationarity* (the distribution of $\{X(t)\}_{t=m}^{m+h}$ is independent of t for a fixed h), the ACF only depends on the time lag, h. This means that

$$r_X(h) = r_X(t, t+h) = \frac{\mathbf{Cov}\{X(t), X(t+h)\}}{\mathbf{Var}\{X(t)\}}.$$

We can estimate the ACF by

$$\hat{r}_X(h) = \frac{(T+1)\sum_{i=0}^{T-h}(X(i) - \overline{X})(X(i+h) - \overline{X})}{(T-h+1)\sum_{i=0}^{T}(X(i) - \overline{X})^2},$$

where

$$\overline{X} = \frac{1}{T+1}\sum_{i=0}^{T} X(i).$$

The function $\hat{r}_X(h)$ is known as the sample auto correlation function.

To investigate linear independence the question is how small the sample acf must be to be interpreted as linear independence. It turns out that, under some assumptions, for independent data,

$$\lim_{n\to\infty} \sqrt{n}\,\hat{r}_X(h) \sim N(0,1).$$

This can be used to construct a 95% confidence interval for independent data. Hence, for n sufficient large we assume linear independence if

$$|\hat{r}_X(h)| \leq 1.96/\sqrt{n}.$$

To investigate independence it is common to continue the analysis with estimation of $f(X(0)), \ldots, f(X(t))$, for instance with $f(x) = |x|$. If one still has no linear dependence for the transformed data one assumes independence.

Example 2.2. A process $X(t) = \epsilon(t) - a\epsilon(t-1)$ where $\{\epsilon(t)\}_{t\in\mathbb{Z}}$ are zero-mean iid random variables is called a *moving average* process of degree 1 (MA(1)). Clearly, we have a dependence between $X(t)$ and $X(t+h)$ for $h = 1$, while we have independence for $h > 1$. In Figure 2.2 we have simulated the process with $\epsilon(t) \sim N(0,1)$ and $a = 0.5$ and estimated the sample acf. We have also plotted the noise and its sample acf. Although the processes looks similar they have very different dependence structure.

Example 2.3. A process $X(t) = \epsilon(t) + aX(t-1)$ where $\{\epsilon(t)\}_{t\in\mathbb{Z}}$ are zero-mean iid random variables is called an *auto regressive* process of degree 1 (AR(1)). Clearly, we will have decreasing dependence between $X(t)$ and $X(t+h)$ for $h > 1$. In Figure 2.3 we have simulated the process with $\epsilon(t) \sim N(0,1)$, $a = 0.5$ and $X(0) = 0$ and estimated the sample acf. We have also plotted the noise and its sample acf. Although the processes looks simular they have very different dependence structure.

2.4 Computer assignment: Portfolio optimization

2.4.1 Setting

In this lab our objective is to create a portfolio (i.e. divide our capital S between n different stocks) in such a way that the difference between the value of the portfolio today and tomorrow is maximized. That is, we would like to optimize it, with optimization performed over all possible divisions of S. It is those divisions that the action space $\mathcal{A}$ consists of, and we can describe each action a as a vector of weights $w = [w_1, ..., w_n]$ such that $\sum_{i=1}^{n} w_i = 1$, where each weight corresponds to how much of our capital we invest in a particular stock.

The states of nature Θ are, similarly to the Section 2.2.1, all possible differences of the stock prices (i.e. returns or log-returns). However, now we do not have one

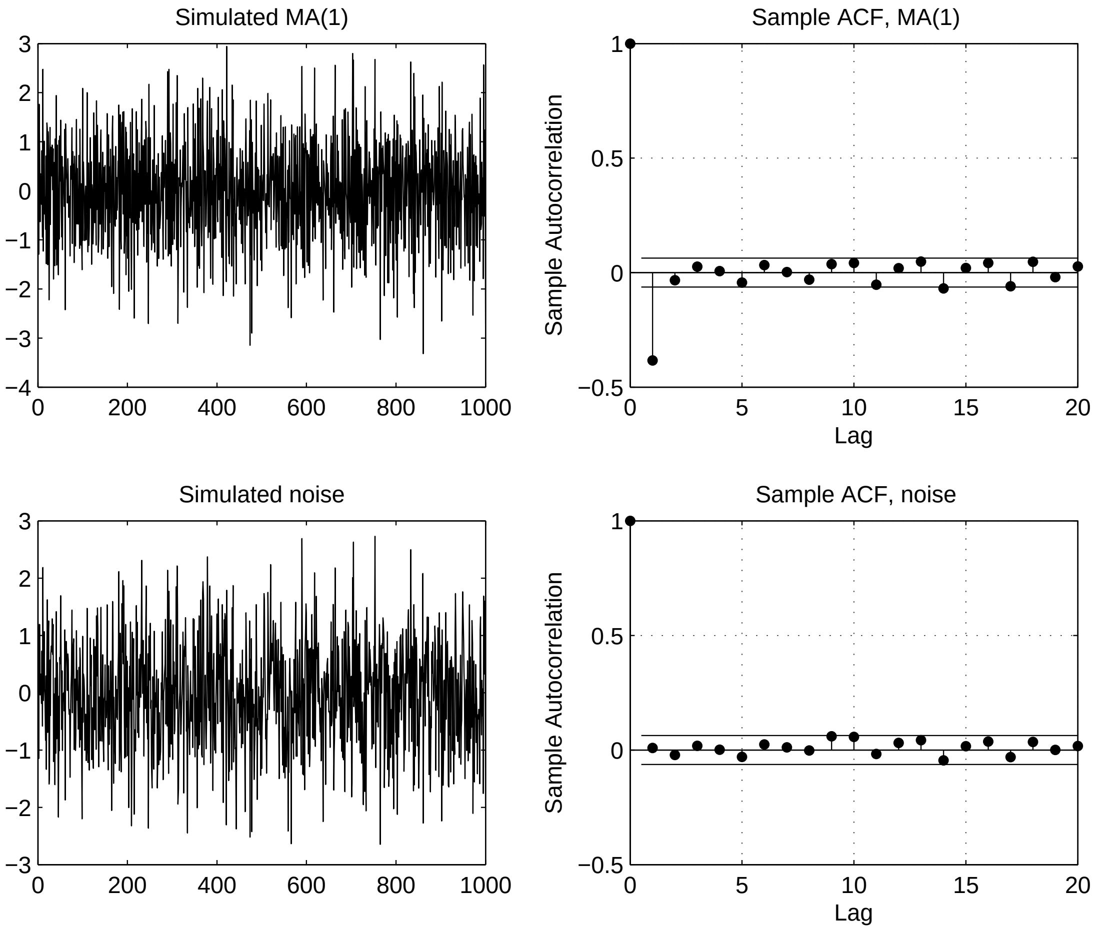

Figure 2.2: Top left plot: Simulated MA(1) with $a = 0.5$. Top right plot: Sample ACF for simulated MA(1). Straight line is the 95% confidence interval for independence. Bottom left plot: Simulated white noise. Bottom right plot: Sample ACF for simulated noise. Straight line is the 95% confidence interval for independence.

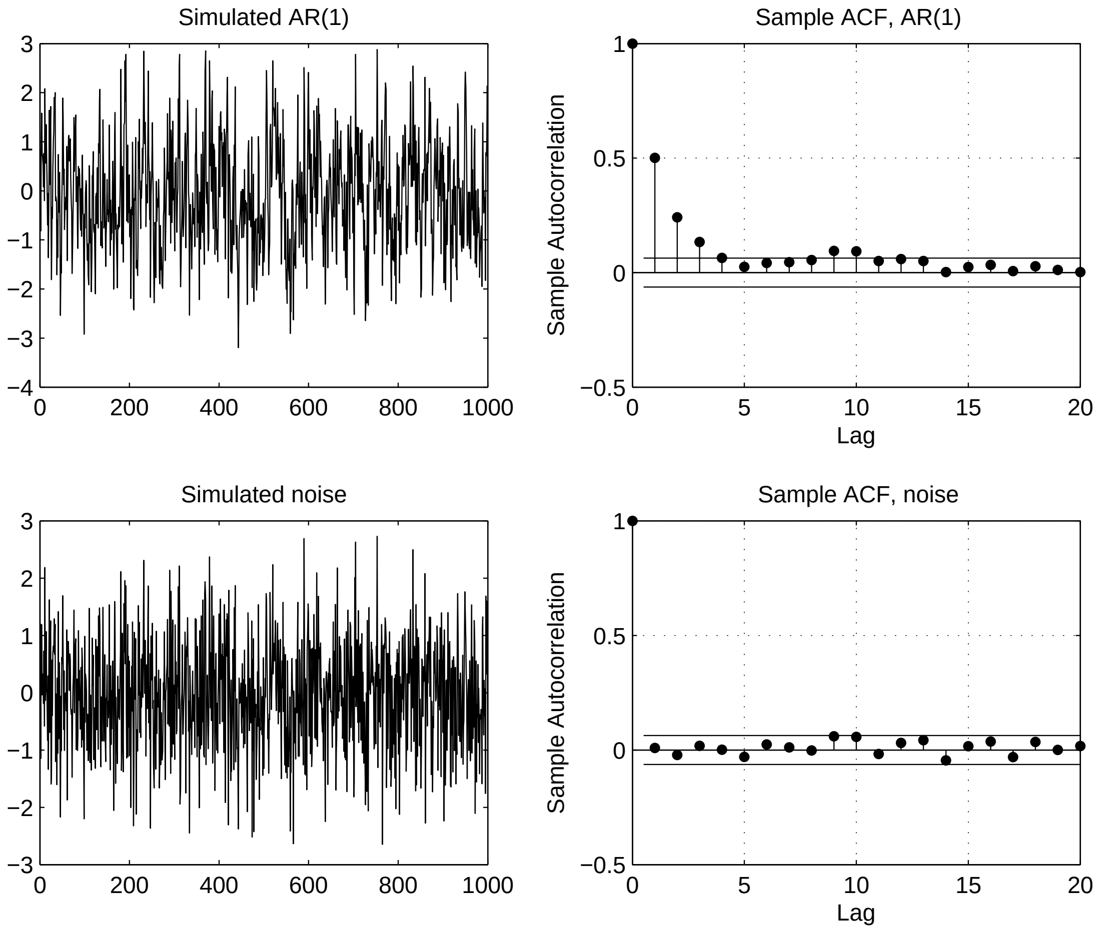

Figure 2.3: Top left plot: Simulated AR(1) with $a = 0.5$. Top right plot: Sample ACF for simulated AR(1). Straight line is the 95% confidence interval for independence. Bottom left plot: Simulated white noise. Bottom right plot: Sample ACF for simulated noise. Straight line is the 95% confidence interval for independence.

single stock but, rather, n. As stated earlier, the (log)returns are considered to be time-independent and normally distributed. If this assumption is correct, we can model all the differences simultaneously with a joint multivariate Normal distribution. The mean vector μ and the covariance matrix Σ can be estimated using historical data. Let $X_i(t)$ be the returns of stock i at time t. Observe also that the returns are assumed to be i.i.d. Normal *in time* (that is, for instance, $X_1(t)$ is independent of $X_1(t-1)$), but it may or may not be true that $X_1(t)$ is independent of $X_2(t)$.

In this assignment you are given the prices for 7 different stocks available to play with, namely AstraZenica, Electrolux, Ericsson, Gambio, Nokia, Swedish Match and Svenska Handelsbanken. These are recorded in form of time series (that is $S(t)$ and not $X(t)$) from 2004-06-03 to 2006-06-01 and can be found in the file *stockdata.tsv*. The first column in the file is a date vector, with the other columns containing the 7 stock price developments.

Our portfolio is going to consist of the weighted sum of the log-returns, that is $\sum_{i=1}^{7} w_i X_i(t)$. Under the above multivariate Normal assumption, the portfolio becomes normally distributed with expected value $z = \mu^T w$ and variance $\sigma^2 = w^T \Sigma w$.

We are going to examine a family of utility functions, where each utility has the form $u(a, \theta) = 1 - e^{-kw^T\theta}$, with k a positive. By denoting total (transformed) gain that the portfolio yields through $x = w^T\theta$, we can define the utility in terms of x rather than a and θ as $u(x) = 1 - e^{-kx}$. Then, the expected utility of the whole portfolio becomes

$$U_\pi(w) = \int_{-\infty}^{\infty} (1 - \mathrm{e}^{-kx}) \frac{1}{\sqrt{2\pi}\,\sigma} \mathrm{e}^{-((x-z)^2/(2\sigma^2))}\, dx \tag{2.1}$$

and we face the problem of maximizing $U_\pi(w)$ under the constraints

$$w_1, \ldots, w_n \geq 0 \quad \text{and} \quad \sum_{i=1}^{n} w_i = 1.$$

It can be shown that this problem is equivalent with the quadratic problem to maximize

$$\mu^T w - \frac{k}{2} w^t \Sigma w \tag{2.2}$$

under the same constraints.

2.4.2 Software

The optimization will be performed in the Matlab environment. For this assignment, you may find the following functions particularly useful:

textread, hist, fmincon

2.4.3 Tasks

1. **Data exploration**. This part of the assignment aims at getting a feel of the data set by ways of plots, histograms etc.

   · Load the data and calculate log-returns. Display those in histograms. Does normal assumption seem plausible?

   · Estimate the autocorrelation function of the log-returns. Plot it. Does the assumption of independence between time points seem plausible?

   · Estimate the mean vector and the covariance/correlation structure between the log-returns corresponding to different stocks. Do they seem to be independent?

   · Explore the utility function for different values of k. Present in a plot. Interpret in terms of risk aversion/seeking.

   · Consider each stock at a time. Approximate the distribution of the corresponding log-returns. Calculate the expected utility for a few different k. Which of the stocks is the most advantageous (may be different for different parameters)?

2. **Optimization with two stocks**. In this part we are going to optimize the portfolio using only two stocks, namely Ericsson and Gambio.

   · Take out the Ericsson and Gambio stocks from the data set. Estimate the mean vector and covariance matrix.

   · Using those estimates and some k, calculate $U_\pi(w)$ for $w_1 \in [0, 1]$, $w_2 = 1 - w_1$. Approximately, for what w_1 do we have maximum expected utility?

   · Find the optimum point more accurately with **fmincon** procedure. Repeat for several different k. Comment on the change of weights.

3. **Optimization with seven stocks**. (2 p)

   · Repeat the previous task using all seven stocks simultaneously.

   · Calculate $U_\pi(w)$ for the "naive" ways of dividing money equally between the stocks or putting it all on a single stock (see Task 1). Comment.

4. **Theoretical extras**. (2 p)

   · Show that the problem of maximizing Equation 2.1 is, in fact, equivalent to the problem of maximizing Equation 2.2.

In order to pass the lab tasks 1 and 2 have to be completed.