Introduction to reliability

Reliability has gained increasing importance in the last few years in manufacturing organisations, the government and civilian communities. With recent concern about government spending, agencies are trying to purchase systems with higher reliability and lower maintenance costs. As consumers, we are mainly concerned with buying products that last longer and are cheaper to maintain, i.e., have higher reliability. The reasons for wanting high product or component or system reliability are obvious:

- Higher customer satisfaction
- Increased sales
- Improved safety
- Decreased warranty costs
- Decreased maintenance costs, etc.

This document covers

- The definition and measurement of reliability and failure rates
- Modelling reliability with the exponential distribution
- Mean time to and before failure
- The addition and multiplication rules of probability
- System reliability and fault tree analysis

Obviously, we only cover a few of the basics here. For more detail see Besterfield (2013), O'Connor (2002), or many other books and websites. There is also a presentation by Professor Ashraf Labib covering asset management, maintenance management and various other practical aspects of the process of managing reliability at http://woodm.myweb.port.ac.uk/q/AshrafHotTopic2012.pdf.

It is important to work through the exercises, and check your answers against the notes on the answers at the end of the document.

What is reliability

The reliability of a product (or system) can be defined as the probability that a product will perform a required function under specified conditions for a certain period of *time*. If we have a large number of items that we can test over time, then the Reliability of the items at time *t* is given by

$$R(t) = \frac{number\ of\ \ survivors\ at\ \ time\ \ t}{number\ of\ \ items\ \ put\ \ on\ \ test\ \ at\ time\ t = 0}$$

At time t = 0, the number of survivors is equal to number of items put on test. Therefore, the reliability at t = 0 is

$$R(0) = 1 = 100\%$$

After this, the reliability, *R(t),* will decline as some components fail (to perform in a satisfactory manner).

The failure rate

The *failure rate* (usually represented by the Greek letter λ) is a very useful quantity. This is defined as the probability of a component failing in one (small) unit of time.

Let N_F = number of failures in a small time interval, say, Δt.

N_S = number of survivors at time *t*.

The failure rate can then be calculated by the equation:

$$\lambda(t) = \frac{N_F}{N_S * \Delta t}$$

For example, if there are 200 surviving components after 400 seconds, and 8 components fail over the next 10 seconds, the failure rate after 400 seconds is given by

λ *(400) = 8 / (200* x *10) = 0.004 = 0.4%*

This simply means that 0.4% of the surviving components fail in each second.

(You may wonder why the above equation defines λ *(400)* and not λ *(410)*. The reason is that Δt is a *small* time interval, so it is reasonable to assume that the failure rate will not change appreciably during the interval. We then define the failure rate using the beginning of the interval for convenience. In the extreme we can make Δt *infinitesimally small*—which is the basis of the differential calculus.)

We can define failure rates for individual components, and also for complex products like cars or washing machines. In the latter case we need to be clear about what is meant by a failure – for a car, for example, these can range from complete breakdown to relatively minor problems. And we should also remember that failures in complex products are generally repairable, whereas this may not be true for individual components.

The Concept of the Bath-tub Curve

The so-called bath-tub curve represents the pattern of failure for many products – especially complex products such as cars and washing machines. The vertical axis in the figure is the failure rate at each point in time. Higher values here indicate higher probabilities of failure.

The bath-tub curve is divided into three regions: infant mortality, useful life and wear-out.

Bath tub curve

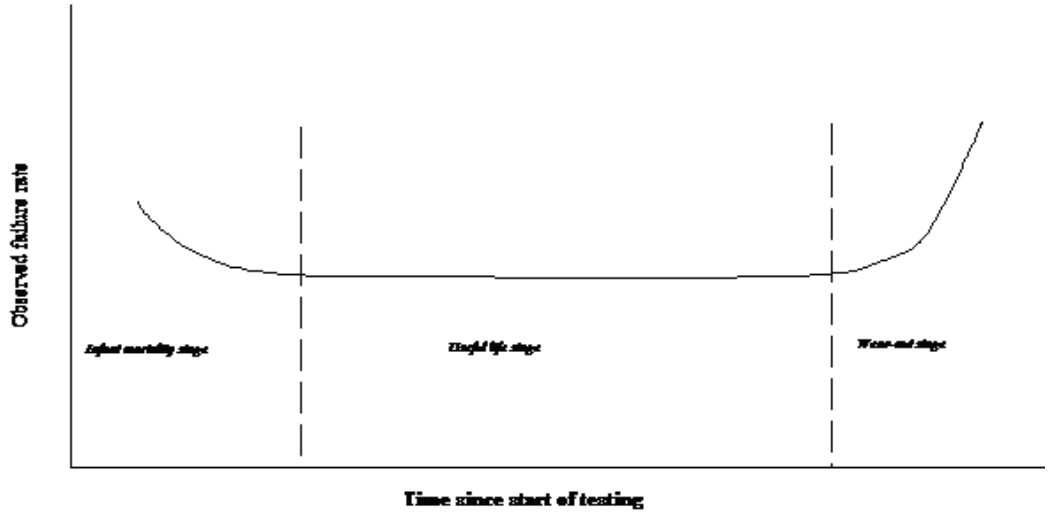

***Infant Mortality**:* This stage is also called early failure or debugging stage. The failure rate is high but decreases gradually with time. During this period, failures occur because engineering did not test products or systems or devices sufficiently, or manufacturing made some defective products. Therefore the failure rate at the beginning of infant mortality stage is high and then it decreases with time after early failures are removed by burn-in or other stress screening methods. Some of the typical early failures are:

- poor welds
- poor connections
- contamination on surface in materials
- incorrect positioning of parts, etc.

Useful life: This is the middle stage of the bath-tub curve. This stage is characterised by a constant failure rate. This period is usually given the most consideration during design stage and is the most significant period for reliability prediction and evaluation activities. Product or component reliability with a constant failure rate can be predicted by the exponential distribution (which we come to later).

Wear-out stage: This is the final stage where the failure rate increases as the products begin to wear out because of age or lack of maintenance. When the failure rate becomes high, repair, replacement of parts etc., should be done.

Exercises

1 Would you expect the bath tub curve to apply to a car? What about a human being?

2 One thousand transistors are placed on life test, and the number of failures in each time

interval are recorded. Find the reliability and the failure rate at 0, 100, 200, etc hours. (You may find it helpful to set this up on a spreadsheet.)

Time interval	Number of failures
0-100	160
100-200	86
200-300	78
300-400	70
400-500	64
500-600	58
600-700	52
700-800	43
800-900	42
900-1000	36

Draw a graph to show the change in the failure rate as the transistors get older.

Do you think this component shows the bath tub pattern of failure?

Draw a graph to show how the reliability changes over time.

Measuring reliability

To see the level and pattern of the reliability of a product or component in practice, it is necessary to make some measurements. The simplest way to do this is to test a large number of products or components until they fail, and then analyse the resulting data. Exercise 2 above shows how this works. This enables us to estimate the failure rate and reliability after different lengths of time - and decide, empirically, if the bath tub curve applies, or if the failure rate shows some other pattern.

There are a number of obvious difficulties which may arise. If the useful life is large it may be not be practical to wait until products or components fail. It may be too expensive to test large samples, so small samples may have to suffice. And for some products (e.g. space capsules) it may be difficult to simulate operating conditions at all closely. There are a number of approaches to these difficulties - most of these are beyond the scope of this unit, but they are discussed in the reading suggested for this session. (An exception is the prediction of the reliability of a product from information about the reliability of its components - this avoids the necessity to test the whole product - which is discussed in a later section of the notes for this session).

It is also possible that the failure rate will depend on environmental conditions in a predictable way. For example, one of the key factors affecting the reliability of electronic components and systems is temperature – basically the higher the temperature of the device the higher the failure rate. Most computer equipment therefore has some form of cooling, ranging from a simple fan to forced chilled air cooling.

Reliability Distributions

There are many statistical distributions used for reliability analysis—for example, the exponential

distribution, the Weibull distribution, the normal distribution, the lognormal distribution, and the gamma distribution. Here we look at the exponential distribution only, as this is the simplest and the most widely applicable.

Reliability Prediction Using the Exponential Distribution

The exponential distribution applies when the failure rate is constant - the graph is a straight horizontal line, instead of a "bath tub". (It can be used to analyse the middle phase of a bath tub - e.g. the period from 100 to 1000 hours in Exercise 2 above.) It is one of the most commonly used distributions in reliability, and is used to predict the probability of survival to a particular time. If λ is the failure rate and t is the time, then the reliability, R, can be determined by the following equation:

$$\mathbf{R(t) = e^{-\lambda t}}$$

There is a brief note on the mathematical background to this equation in the Appendix.

To see that it gives sensible results, imagine that there are initially 1000 components and that λ is 10% (0.1) per hour. After one hour about 10% of the original 1000 components will have failed - leaving about 900 survivors. After two hours, about 10% of the 900 survivors will have failed leaving about 810. Similarly there will be about 729 survivors after the third hour, which means that the reliability after 3 hours is 0.729.

Using the exponential distribution the reliability after 3 hours, with λ=0.1, is given by

$$\mathbf{R(t) = e^{-3\lambda} = e^{-0.3} = 0.741}$$

(You can work this out using a calculator or a spreadsheet—see the mathematical appendix for more details.)

This is close to the earlier answer as we should expect. The reason it is not identical is that the method of subtracting 10% every hour to obtain 900, 810, etc ignores the fact that the number of survivors is changing all the time, not just every hour. The exponential formula uses calculus to take this into account.

If the failure rate is small in relation to the time involved a much simpler method will give reasonable results. Let's suppose that λ=0.01 (1%) in the example above. The formula now gives the reliability after three hours as

$$\mathbf{R(t) = e^{-3\lambda} = e^{-0.03} = 0.9704}$$

A simpler way of working this out would be just to say that if the failure rate is 0.01 per hour the total proportion of failures in 3 hours will be 0.03 (3%) so the reliability after three hours is simply

R(t) = 1 – 0.03 = 0.97 (or 100% – 3% = 97%)

This is not exactly right because in each hour the expected number of failures will decline as the surviving pool of working components gets smaller. But when the failure are is 1% and we are interested in what happens after three hours, the error is negligible. On the other hand, if we want to know what happens after 300 hours the simple method gives a silly answer (the reliability will be negative!) and we need to use the exponential method. You should be able to check this answer with your calculator or a computer (the answer should be 0.0498 or about 5%).

Mean time to Failure (MTTF) and Mean time between Failures (MTBF)

MTTF applies to non-repairable items or devices and is defined as "the average time an item may be expected to function before failure". This can be estimated from a suitable sample of items which have been tested to the point of failure: the MTTF is simply the average of all the times to failure. For example, if four items have lasted 3,000 hours, 4000, hours, 4000 hours and 5,000 hours, the MTTF is 16,000/4 or 4,000 hours.

The MTBF applies to repairable items. The definition of this refers to “between” failures for obvious reasons. It should be obvious that

MTBF = Total device hours / number of failures

For example, consider an item which has failed, say, 4 times over a period of 16,000 hours. Then MTBF is 16,000/4 = 4,000 hours. (This is, of course, just the same method as for MTTF.)

For the particular case of an *exponential distribution*,

$\lambda = 1/\text{MTBF}$ (or $1/\text{MTTF}$)

where λ is the failure rate.

For example, the item above fails, on average, once every 4000 hours, so the probability of failure for each hour is obviously 1/4000. This depends on the failure rate being constant - which is the condition for the exponential distribution.

This equation can also be written the other way round:

$\text{MTBF (or MTTF)} = 1/\lambda$

For example, if the failure rate is 0.00025, then

MTBF (or MTTF) = 1/0.00025 = 4,000 hours.

Exercises

3 An industrial machine compresses natural gas into an interstate gas pipeline. The compressor is on line 24 hours a day. (If the machine is down, a gas field has to be shut down until the natural gas can be compressed, so down time is very expensive.) The vendor

knows that the compressor has a constant failure rate of 0.000001 failures/hr. What is the operational reliability after 2500 hours of continuous service?

4 What is the highest failure rate for a product if it is to have a reliability (or probability of survival) of 98 percent at 5000 hours? Assume that the time to failure follows an exponential distribution.

5 Suppose that a component we wish to model has a constant failure rate with a mean time between failures of 25 hours? Find:-

(a) The reliability function.

(b) The reliability of the item at 30 hours.

6 The equipment in a packaging plant has a MTBF of 1000 hours. What is the probability that the equipment will operate for a period of 500 hours without failure?

7 TALCO manufactures microwave ovens. In order to develop warranty guidelines, TALCO randomly tested 10 microwave ovens continuously to failure. The failure information of the 10 ovens is shown in the table.

Microwave	Hours
1	2300
2	2150
3	2800
4	1890
5	2790
6	1890
7	2450
8	2630
9	2100
10	2120

What is the mean time to failure of the microwave ovens?

Does the evidence suggest that the reliability of the ovens follows the exponential distribution (with a constant failure rate)?

The addition and multiplication rules of probability

The next aspects of reliability theory we will consider depend on some probability theory—the addition and multiplication rules. I will explain these in terms of dice and cards.

Suppose that you throw a single dice. The probability of getting a 6 is

$$P(6) = 1/6$$

And the probability of getting a 5 is also

$$P(5) = 1/6$$

Equally obvious is that the probability getting a 5 or a 6 is

$$P(5 \text{ or } 6) = 2/6 = 1/3$$

And that

$$P(5 \text{ or } 6) = P(5) + P(6)$$

Similarly if you choose a single card from a pack of (52) cards, the probability of getting an ace or a picture card (jack, queen or king) is

$$P(\text{ace or picture}) = 4/52 + 12/52 = 16/52$$

because 4 of the 52 cards are aces, and another 12 are picture cards.

These examples suggest that if we want the probability of something *or* something else happening, we can *add* the probabilities.

But … what about the probability of an ace or a red card (hearts or diamonds). Can we say that

$$P(\text{ace or red}) = P(\text{ace}) + P(\text{red}) = 4/52 + 26/52 = 30/52 \text{ ??}$$

This is obviously wrong because two of the aces are also red, so we are in effect double counting these aces if we add the probabilities. Before adding the probabilities you need to check that the two events cannot both occur—i.e. they do not overlap or are *mutually exclusive* (i.e. each excludes the other).

The complete addition rule is

$$P(A \text{ or } B) = P(A) + P(B)$$ *if A and B are mutually exclusive (i.e. they don't overlap).*

It can easily be extended to three or more events:

$$P(A \text{ or } B \text{ or } C \text{ or } \ldots) = P(A) + P(B) + P(C) + \ldots$$ *if A, B, C … are mutually exclusive.*

To explain the *multiplication* rule we need to do more than one thing, so let's throw the dice *and* draw a card from the pack. How can we work out the probability of getting a 6 *and* a spade?

This is a little more complicated than the *addition* rule. It helps to imagine doing the experiment lots of times—say 1000 times.

Obviously you will get a 6 on about 1/6 of these thousand times—i.e. about 167 times. Now think about how many times you will get a spade as well. This will happen on about ¼ of these 167 times, or about 42 times out of the 1000 times we imagined doing the experiment. So the probability is about 42/1000.

Now do the same thought experiment, but working in terms of the probabilities this time. We will get a 6 on about one sixth of the times, and a spade as well on about one quarter of this one sixth. One quarter of one sixth means the same as one quarter times one sixth, so we simply multiply the probabilities:

P(6 and spade) = P(6)*P(spade) = 1/6 * 1/4 = 1/24

which is, of course, about 42/1000 as before.

Just like the addition rule there is an important assumption here. We assumed that the two events are *statistically independent*. This means that the two probabilities are independent of each other: knowing whether one has happened is of no help in assessing the probability of the other happening. In the example above, knowing that we've got a 6 is of no relevance to what will happen with the cards—so these events are statistically independent.

But in some situations you need to be very careful about this assumption. Suppose you know that in a particular place the probability of rain falling on a given day is 1/3. Can you say that

P(rain today *and* rain tomorrow) = 1/3 * 1/3 = 1/9 ??

In practice this is likely to be wrong because the two events are not likely to be independent. In England certainly, if it rains today the probability of rain tomorrow is likely to be higher than if it did not rain today because the weather tends to set in to wet or dry spells. This means that the second probability should probably be rather more than 1/3, so the result is likely to be substantially more than 1/9.

The multiplication rule, then, also has an important condition:

P(A and B) = P(A)*P(B) *if A and B are statistically independent*

And, just as before, it can be extended:

P(A and B and C and …) = P(A)*P(B)*P(C)*… *if A, B, C, … are statistically independent*

Analyzing and improving the reliability of Systems

It is often useful to be able estimate the reliability of a whole system from the reliability of the individual components. (The system in question might be a satellite to be sent into space, or a computer, or a process for responding to accidents.) This enables designers to predict reliability levels without having to build and test the whole system. It also enables them to see weak spots in

the system, and to experiment with ways of improving these - perhaps by installing backups for critical components.

There are two established ways of analyzing the reliability of a system. The first uses *reliability block diagrams* to model the reliability of individual components and to predict the reliability of the whole system. The second uses *fault tree analysis* to model the probabilities of failure of individual components and so predict the probability of the whole system failing. Since the reliability of the system is 100% minus the probability of failure (as a percentage), the two approaches are equivalent so there is no point in doing both. Here we will look at fault tree analysis. (Reliability block diagrams do the same thing in a slightly different way but come to the same conclusions.)

Fault Tree Analysis

This is a commonly used technique in industry to evaluate reliability of a system. It was developed in the early 1960s to evaluate reliability of the Minuteman Launch Control System. Since then it has gained favour especially when analysing complex systems. Fault tree analysis begins by identifying the top event, known as the undesirable event of the system. The undesirable event of the system is caused by events generated and connected by logic gates such as AND, OR, etc (as you will see below). The following basic steps are involved in performing fault tree analysis:

i *Establishing system definition.*

ii *Constructing the fault tree.*

iii *Evaluating the fault tree qualitatively.*

iv *Collecting basic data such as components' failure rates, repair rates, and failure occurrence probability.*

v *Evaluating fault tree quantitatively.*

vi *Recommending corrective measures.*

This method is frequently used as a qualitative evaluation method in order to assist the designer, planner or operator in deciding how a system may fail and what remedies may be used to overcome the causes of failure. The method can also be used for quantitative evaluation, in which case the causes of system failure are gradually broken down into an increasing number of hierarchical levels until a level is reached at which reliability data is sufficient or precise enough for a quantitative assessment to be made. The appropriate data is then inserted into the tree at this hierarchical level and combined together using the logic of the tree to give the reliability assessment of the complete system being studied.

In order to illustrate the application of this method, consider the electric power requirements of the system in the following example. In this example, the failure event being considered is loss of the electric power. In practice the electric power requirements are both a.c. power, to supply energy for

prime movers, and d.c. power, to operate relays and contractors, both of which are required to ensure the successful operation of the electric power. Consequently the event 'loss of electric power' can be divided into two sub-events 'loss of a.c. power' and 'loss of d.c. power'. This is shown in the following figure with the events being joined by an OR gate as failure of *either, or both*, causes the system to fail.

A very simple fault tree

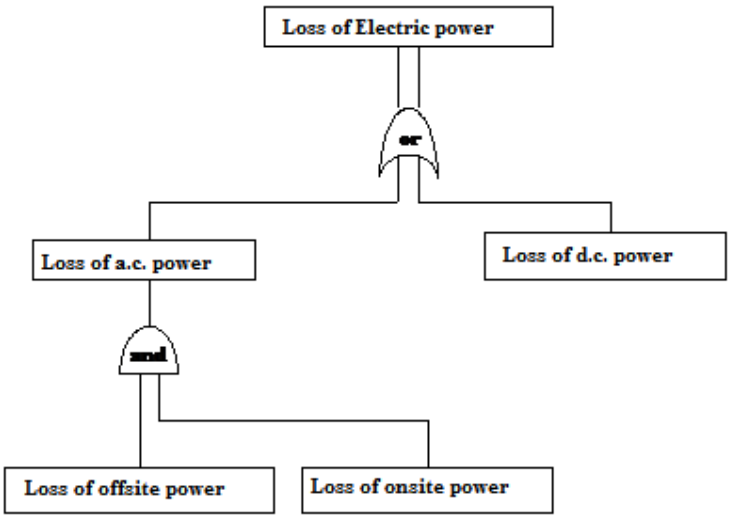

If this subdivision is insufficient, sub-events can be divided further. The event 'loss of a.c. power' may be caused by 'loss of offsite power' (the grid supply) and by 'loss of onsite power' (standby generators or similar devices). These are joined by an AND gate as they *both* have to fail for the a.c. power to fail. This process can be continued downwards to any required level of subdivision. After developing a fault tree, it is necessary to evaluate the probability of occurrence of the upper event by combining component probabilities using basic rules of probability and the logic defined in the fault tree.

In the present example, suppose the probabilities of the events at the bottom of the tree are:

Prob (Loss of offsite power) = 0.067

Prob (Loss of onsite power) = 0.075

Prob (Loss of dc power) = 0.005

These probabilities for the faults at the bottom of the tree can now be combined using the addition and multiplication rules of probability. For the AND gate at the bottom we multiply the probabilities to work out

Prob (Loss of a.c. power) = Prob(Loss of offsite power) x Prob (Loss of onsite power)

= 0.067 x 0.075 = 0.005025.

For the OR gate we add the probabilities to get the probability of the top event:

Prob (Loss of electric power) = Prob (Loss of a.c. power) + Prob (Loss of d..c power)

= 0.005025 + 0.005 = 0.010025.

This means that the probability of the top, undesirable, event – Loss of electric power – is about 1%. The calculation make two assumptions.

First, to multiply the two probabilities at the bottom of the tree we must assume they are *statistically independent*. This is reasonable if the onsite and offsite systems are separate so that the failure of one is independent of the other. Notice that the combined probability (Loss of a.c. power) is much less than the probability of either of the two probabilities in the AND gate. The principle here is that of reducing the risk of a fault by using a backup system: this will not work, of course, if the two systems are dependent so that if one fails there is an increased chance that other will fail too.

The second assumption is the assumption when we add the probabilities for the OR gate that the events are mutually exclusive (i.e. they don't overlap). With small probabilities, like we have here, this is reasonable because the probability of more than one fault occurring is small.

The case study and exercises below contain further examples of fault trees.

Case study: healthy laptop

This case study explores some of the concepts of reliability. It concerns a private provider of health care that runs a healthcare system with 8 acute care hospitals, 100 clinics, and home healthcare.

The home healthcare (HHC) division provides nurses (i.e. district nurses) to home care patients. It employs 125 nurses and each one is provided with a laptop that allows them to maintain patient records, point of care documentation, case loads, and payroll entry. The nurses maintain a remote connection with the central server to access new cases and upload any historical patient data gained throughout the day. The availability of the system, quality of the data and timeliness of transcribing information are critical to the nurses.

The nurses have been complaining that they are having a lot of problems with the laptops and the connection to the network and server. Reported laptop performance is very poor and 60% of the nurses have complained of problems of one sort or another. All the laptops are from the same manufacturer and have a similar configuration. Whenever a failure occurs the nurse has to return the laptop to the office, obtain a loan machine and then return that to the office when their own

laptop had been repaired. The company knows how many machines have been returned to the manufacturer for repair but did not itself keep any records of failure symptoms and causes – that is done by the vendor's field service group who made the repairs.

Discussions between HHC's IT Director and the manufacturer's Field Support Manager are getting progressively more heated with HHC threatening to switch suppliers completely – with serious consequences for the manufacturer who actually supplied all HHC's IT hardware. Losing the laptop contract is likely could result in the entire customer being lost.

In order to try and bring the relationship back onto a more business like footing HHC decided to bring in an independent reliability specialist to arbitrate.

The manufacturer did not, and would not, publish reliability data for its products, saying that

> "We do not quote MTBF (Mean Time Before Failure) numbers for our products or believe this type of information should be used as a meaningful description of the quality of our systems or component parts within those systems. MTBF is an older industry term that today has very little value and is mostly misinterpreted and misunderstood. MTBF is the point at which 63.2% of the population (all components of the given type built by a manufacturer), will fail. So for example, disk drives that claim 1,000,000 hour MTBF with 720 power on hours per month would take 115 years to reach the 63.2% mark. It does NOT mean that no failure should occur for 115 years. This is a total population statement from the manufacturer's point of view not from the customer point of view. As a Manufacturer, we know that some of that product will fail but we have no idea which ones, and we do not know the distribution of good or bad product shipped to any that may fail earlier than expected and therefore may achieve better or worse results and still meet the MTBF criteria."

The manufacturers were willing to show the defect data for HHC's laptops, but once again would not compare those results with a control sample of the same number of the same product. This laptop failure data (a Pareto diagram) was produced at the next meeting, as in the diagram below.

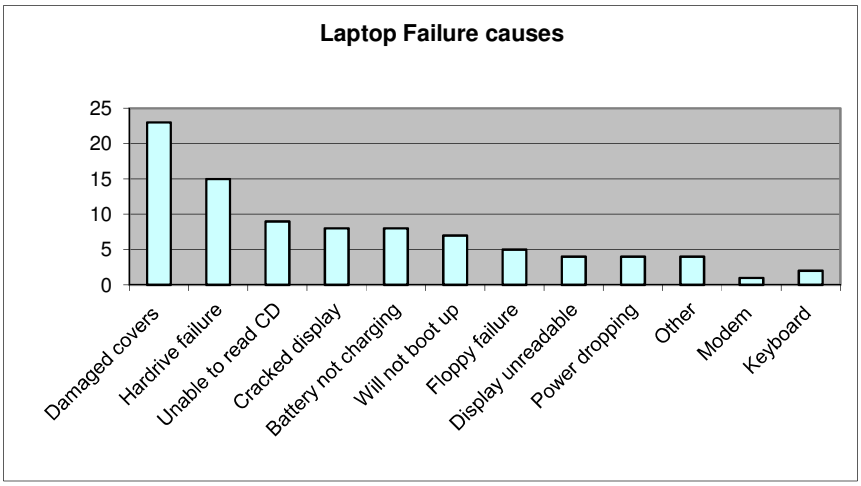

The average age of the machines was three years. This prompted some discussion that they were therefore fully depreciated and should be written off and replacements bought. The consultant thought that this was an accountant's view of the situation and was not a view that could be supported from a quality and reliability perspective. At three years the machines should be at the most reliable part of the reliability (bathtub) curve – well past early life failure and not yet approaching the wear-out part of the curve when failures start to increase. The high number of problems due to damaged covers, hard drives and CD failures made him suspect that part of the problem lay in the environment in which the laptops were used and the treatment that they received. If this were the case then replacing them with new machines would not solve the problem as the new ones would ultimately suffer the same fate.

The diagram below shows a fault tree for one of these problem categories – hard drive failures. (In order to make it legible it has not been fully developed—seek and logic failures have not been explored.)

Fault Tree - Hard Drive

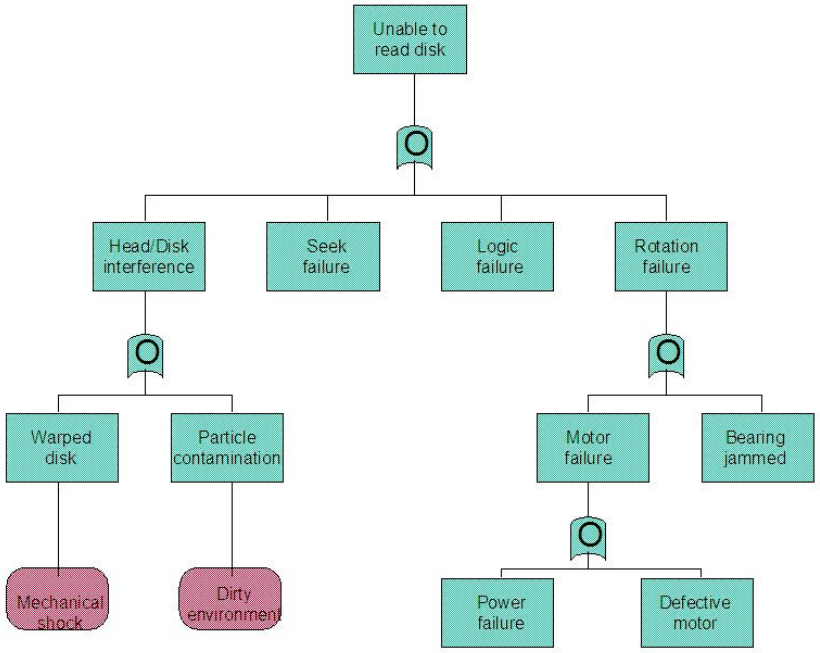

Exercises

8 In relation to the Healthy Laptop case study:

(a) Comment on the manufacturer's refusal to publish MTBF data – is their position reasonable?

(b) Where does the 63.2% come from?

(c) How can the fault tree be used to analyse and help prevent failures? Can it show how misuse can contribute to the failure of the component under review?

9 Two football teams are due to play each other three times in the next few weeks. The data from their matches in the past suggests that the probability of Team A winning is 50%, and the probability of Team B winning is 30%

a) What is the probability of a draw in the first match?

b) Assuming that the results of the three matches are statistically independent what is the probability of Team B winning all three matches?

c) Assuming that the results of the three matches are statistically independent what is the probability of Team B winning at least one of the matches?

d) Do you think the assumption of statistical independence is justified?

10 During World War II, the failure rate for bomber aircraft flying over Europe was believed to be a 4% chance of non-return from each mission, however experienced the pilot was. Calculate the probability that a crew member will survive 25 missions.

11 The diagram below is a cross-section diagram of a booster rocket outer shell at a joint between two stages. The system will fail only if both o-rings fail.

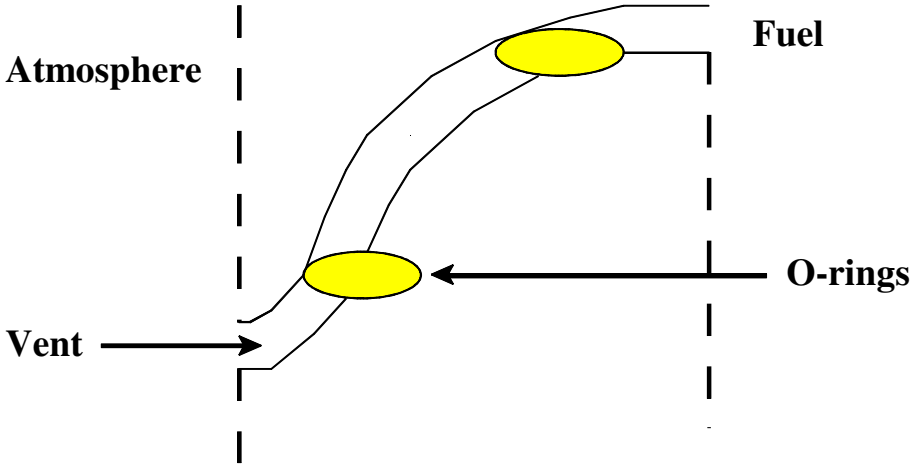

(a) Suppose the reliability of each o-ring is 0.95 during the most critical phase of flight. What is the system reliability? (You will need to work out the probability of failure for each O-ring.)

(b) Do you think the two reliabilities are likely to be statistically independent? Would this make any difference to your answer for the system reliability?

12 A certain electronic component has an exponential failure time with a mean of 50 hours.

(a) What is the rate of this component?

(b) What is the reliability of this component at 100 hours?

(c) What is the minimum number of these components that should be placed in parallel if we desire a reliability of 0.90 at 100 hours? (The idea of placing extra components in parallel is to provide a backup if the first component fails.)

13 A risk assessment of a proposed new railway included the following results:

Passenger train derailment would occur *if*

derailment occurred due to over-speeding (4.3×10^{-11}) *or*

derailment occurred due to rail faults (6.4×10^{-9}) *or*

derailment occurred due to rolling stock faults (3.9×10^{-9}) *or*

derailment occurred due to running into obstructions (7.1×10^{-9})

The figures in brackets represent the estimated probability of occurrence of each of these per train

kilometre. For example the first probability means that the probability for the event in question is 4.3×10^{-11} for each train for each kilometre. Clearly the probability for six trains, each travelling 100 km, would be 600 times as great.

The probabilities were estimated by breaking the events down into more detail and using historical data. For example, "derailment occurred due to rolling stock faults" was broken down into seven events which would lead to rolling stock faults - these included wheel failure and suspension system failure. Furthermore, suspension system failure would occur *if*

the suspension system deflated (1.5×10^{-6}), *and*

the emergency springs failed (1.0×10^{-4})

The assigned task is:

(a) Draw a fault tree from the above information.

(b) Estimate the probability (per train kilometre) of suspension system failure and explain the assumptions on which your estimate is based.

(c) Estimate the probability (per train kilometre) of passenger train derailment (the top event of the fault tree) and explain the assumptions on which your answer is based.

(d) Assuming that the railway is 100 km long and there are 20 trains each way every day, estimate the mean time between passenger train derailments. Do you think this represents a satisfactory level of risk?

(e) What do you think of the accuracy and usefulness of this type of analysis?

(The events and probabilities are taken from a presentation by C. Leighton & C. Dennis at a conference on Risk Analysis and Assessment, Edinburgh, 1994.)

Notes on answers to exercises

1 The bath tub pattern would apply to a human being. Failure rates, in the sense of death, are higher in early infancy and old age. And I think it might apply to a car. Old cars are probably more liable to failure than new ones, and there is also a possibility that when a car is new failure rates may be higher as faults left over from the manufacturing process are found and remedied. On the other hand, if the manufacturer has checked carefully for the early faults the infant mortality stage may not occur.

2 My answers are below.

Time	Survivors	Failures in next 100 h	Reliability	Failure rate (per hr)
0	1000	160	100.00%	0.16%
100	840	86	84.0%	0.10%
200	754	78	75.4%	0.10%
300	676	70	67.6%	0.10%
400	606	64	60.6%	0.11%

500	542	58	54.2%	0.11%
600	484	52	48.4%	0.11%
700	432	43	43.2%	0.10%
800	389	42	38.9%	0.11%
900	347	36	34.7%	0.10%
1000	311		31.1%	

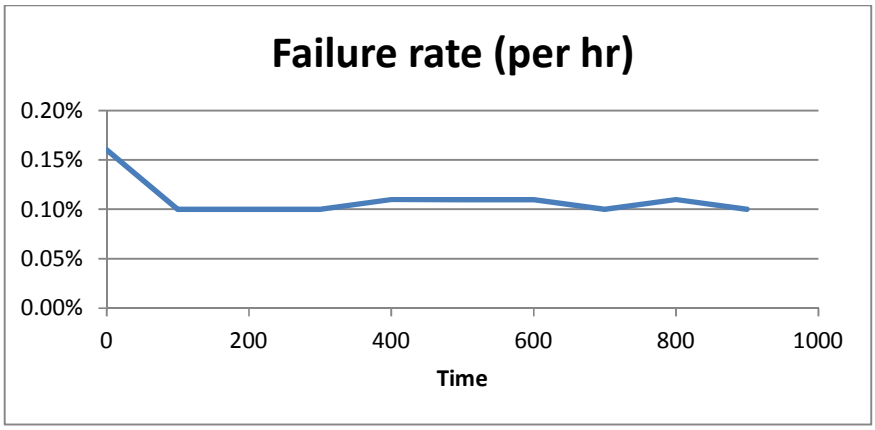

Note that the failure rate is *per hour.* (This is why the figures are so small.) A failure rate of 0.10% means that the probability of a failure in a given hour is 0.10%.

Note also that I am using data from the whole interval from 0-100 hours to calculate the failure rate at 0 hours. This is an approximation but it is all we can do. In the table I have assumed that the first failure rate (0.16%) applies to 0 hours – but I might equally well have said that it applies to 50 hours.

This result shows that the failure rate is initially 0.16% per hour, and then drops to around 0.10% to 0.11%. There is an infant mortality phase, a useful life phase with a fairly constant failure rate, but no wear out phase. However, there are still 311 components left after the 1000 hour test - the wear out phase may become apparent if the test were to be prolonged. Obviously, the best way to show this pattern would be to draw a graph of failure rate against time.

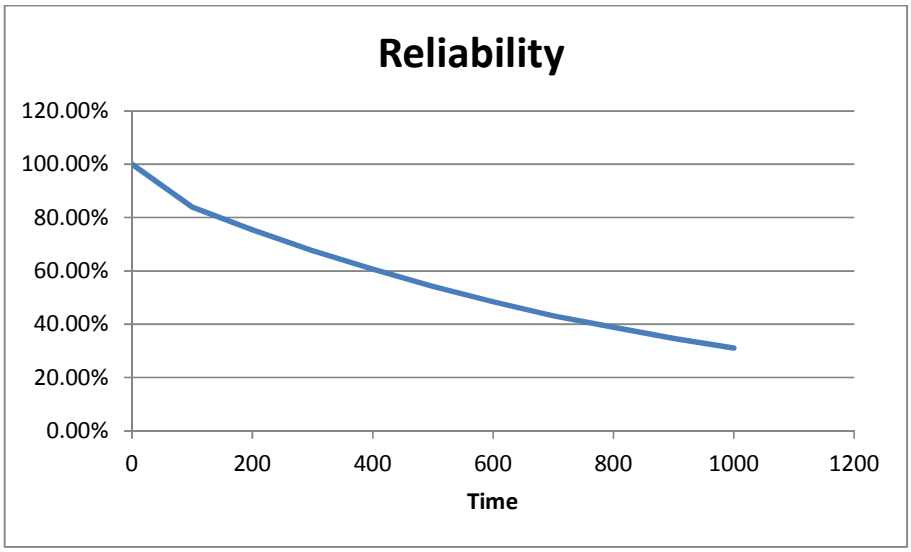

This graph is very similar to the exponential distribution except that it is a bit steeper for the first 100 hours because of the higher failure rate.

3 The compressor has a constant failure rate and therefore the reliability follows the exponential distribution:

$$R = e^{-\lambda t}$$

Failure rate λ = 0.000001 f/hr., operational time t = 2500 hours

Reliability = $e^{-(0.000001 * 2500)}$ = 0.9975 or 99.75%

4 The reliability of the product is given to be 0.98. The reliability of an exponential distribution is given by:

$$R = e^{-\lambda t}, \text{ so}$$

$$0.98 = e^{(-\lambda * 5000)}$$

Taking natural logarithms on both sides (see Appendix), we get,

$$-0.02020 = -\lambda * 5000$$

Therefore $\lambda = 4.04 * 10^{-6} = 0.00000404$ f/hr

(Alternatively, you could use a trial and error process with your calculator to find a negative number which gives you 0.98 when you press the e^x button. This number must be equal to λ*5000, so you can divide the number by 5000 to find λ.)

5 (a) Since the failure rate is constant, we will use the exponential distribution. Also, the MTBF = 25 hours. We know, for an exponential distribution, MTBF = 1/λ.

Therefore λ = 1/25 = 0.04

The reliability function is given by: $R(t) = e^{-\lambda t} = e^{-(0.04 * t)}$

(b) The reliability of the item at 30 hours = $e^{-0.04 * 30} = 0.3012$

6 Assuming the exponential model, the failure rate is

1/MTBF = 0.001

So

R(500) = $e^{-500*0.001} = e^{-0.5} = 0.61$, so the probability of the equipment operating for 500 hours without failure is 61%.

7 The MTTF is 2312 hours (simply the average). The evidence suggests that the ovens do *not* follow the exponential distribution, because the failure rate is obviously not constant: there are no failures between in the first 1000 hours, 2 failures between 1000 and 2000 hours, and 8 between 2000 and 3000 hours. This gives an estimated failure rate of 0 for the first 1000 hours, 20%/1000 for the next 1000 hours, and 100%/1000 for the next 1000 hours – this certainly increases over time)

8 In these situations it is always very frustrating for customers and user groups to be told that they will not be given Mean Time Between Failure (MTBF) and other failure rate data. From the manufacturer's perspective this is a sensible approach because there is always a danger that the MTBF data will be misinterpreted by the customer. MTBF is a measure applied to the population as a whole, not to a very small sample.

It is true that caution must exercised when applying averages to individual cases. However, provided that the customer understands that the MTBF is an average figure, the MTBF would still be useful, particularly for a customer such as HHC which buys large numbers of laptops. In this situation, a significant departure from the average figures would indicate either a problem with the laptops supplied, or that they are being used in unsuitable conditions.

The 63.2% comes from the exponential reliability distribution. If the MTBF is 1,000,000 hours, the failure rate is 0.000001 per hour and the exponential distribution gives

$R = e^{-0.000001 \times 1000000} = e^{-1} = 0.368 = 36.8\%$

which means that the other 63.2% must have failed. The quotation from the manufacturer is misleading in that it is not a *definition* of the MTBF, but rather a fact that follows from the exponential distribution.

The fault tree is a good illustration of how these can help in the analysis and prevention of failures. It shows how external factors such as mechanical shock and dirt can damage components. Other fault trees should be constructed for the other sub-systems contained in the laptop schematic diagram and then combined into a fault tree for the entire laptop system. The first stage of the analysis of the fault tree is a qualitative evaluation of the subsystem to understand how failure occurs. The next stage is to gather data on component failure rates and repair rates to allow a quantitative evaluation – where are the major sources of failure are, and the actions that can be taken to minimise or eliminate them.

9a The match must end in a win for Team A, *or* a win for Team B, *or* a draw. These three probabilities are mutually exclusive so they must add up to 100%. This means the probability of a

draw is 20% because 50%+30%+20%=100%.

9b P(B winning three matches)=P(B wins 1st match)*P(B wins 2nd match)*P(B wins 3rd)

$=30^{25}\% * 30\% * 30\% = 27/1000 = 0.027 = 2.7\%$

so it is unlikely!

9c This is a little more difficult. The easiest approach to "at least one" questions is to work out the probability of the event never happening, and then subtract it from 1. In this case, the probability of Team B failing to win each match is 70% (100%-30%), so the probability of failing to win all three matches is

$$70\% * 70\% * 70\% = 7/10 * 7/10 * 7/10 = 343/1000 = 34.3\%$$

In all other circumstances Team B will win at least one match so the probability is

$$P(\text{B wins at least one match}) = 100\% - 34.3\% = 65.7\%$$

Which is fairly likely!

9d This is a difficult question to which there is no easy answer. Perhaps if Team B wins the first match they would become more confident and so more likely to win the next match? Or perhaps Team A would get more determined.

10 The probability of surviving 10 missions is $(1-0.04)^{25} = 0.96^{25} = 0.36 = 36\%$.

11 a Assuming statistical independence we the probability of one failing is

$$0.05 \times 0.05 = 0.0025$$

So the system reliability is $1-0.0025 = 0.9975 = 99.75\%$.

b This answer assumes the two potential failures are statistically independent - which may be unlikely in practice. A single factor - such as temperature - may be responsible for failures in both. This means that, if we know one o-ring has failed, our estimate of the probability of the other failing will be higher. In the extreme case, one o-ring fails whenever the other does, and the reliability of the system is 0.95. In practice, the reliability of the whole system is likely to be between 0.95 and 0.9975 (the answer based on the assumption of independence).

11 a $1/50 = 2\%$ per hour

b $R(100) = e^{-0.02 \times 100} = 0.1353$, which is not very good!

c The parallel system will only fail if *all* components fail. The probability of each failing is 1-0.1353 = 0.8647.

If there are n in parallel we need

$1 - 0.8647^n = 0.9$, or

$0.8647^n = 0.1$

By trial and error (or see the Appendix), n = 16, so we need 16 components in parallel.

13 (a) The fault tree is on the next page.

(b) Assuming the probabilities of the two events (suspension system deflated and emergency springs failed) are *independent*, we can multiply the probabilities, so the estimate for the probability of suspension system failure is 1.5×10^{-10}. In practice, this assumption of independence may not be fully realistic if some causes of the first event also make the second more likely. If the events are not independent, the probabilities may be much higher.

(c) To work out the probability of passenger train derailment we add the four probabilities to give 1.7×10^{-8}. Note that the probability of the first event - over-speeding - is negligible compared with the others.

This assumes that these are the *only* faults that will lead to derailment. Sabotage, for example, does not seem to be included. (Strictly, adding probabilities presupposes the events are mutually exclusive. It is possible that more than one event can occur. However, the probability of this is so small - of the order of 10^{-18} - that it can reasonably be ignored.)

(d) The probability of derailment per day is $100 \times 20 \times 2 \times 1.7 \times 10^{-8} = 6.8 \times 10^{-5}$ This is, in effect, a failure rate, so the mean time between derailments (failures) is $1/(6.8 \times 10^{-5})$ or 1.5×10^4 or 15000 days or about 41 years. Derailments can be expected, on average, every 41 years.

(e) The accuracy and usefulness of the model are obviously dependent on the data on which it is built. In particular, it is obviously important that all input probabilities are accurate (e.g. of the emergency springs failing), that assumptions about statistical independence are carefully checked, and that all lists of events specifying the possible ways a fault can occur are as complete as possible. This obviously requires systematic research.

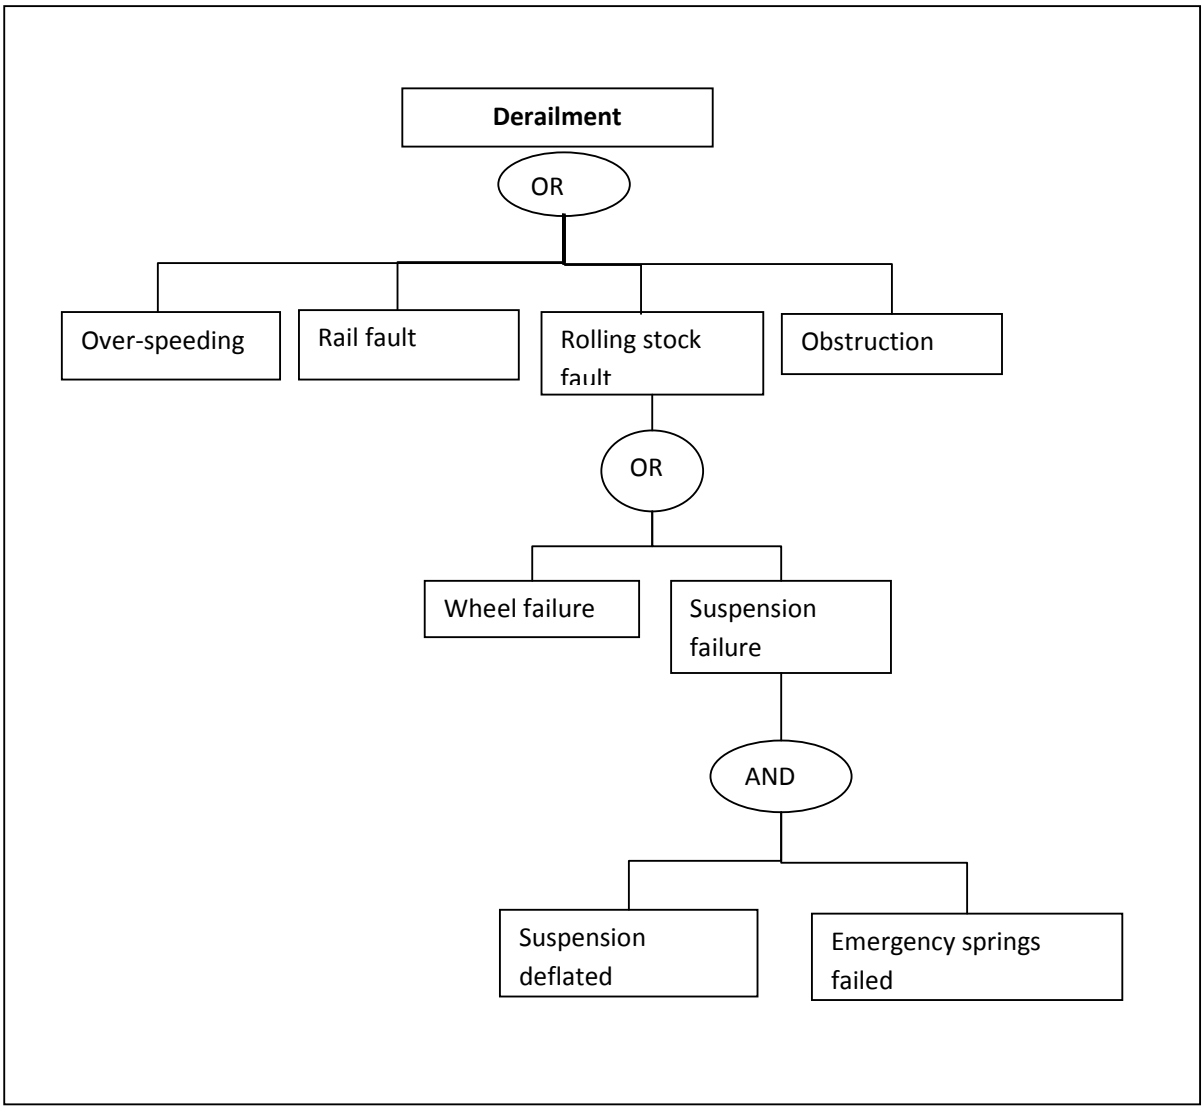

References

Besterfield, D. H. (2013). *Quality improvement.* Boston: Pearson.

O'Connor, P. D. (2002). *Practical reliability engineering, 4th edition.* Chichester: Wiley.

Mathematical Appendix

Powers and roots

As you will doubtless know

pppp (four p's multiplied together) can be written as p^4

You may also need to do this backwards. Suppose that

$p^4 = 0.7$ (p^4 is p^4 on a spreadsheet.)

and you want to know *p*. *p* is the fourth root of 0.7 or

$$p = 0.7^{1/4} = 0.7^{0.25} = 0.915$$

To check, try 0.915^4. This should come to 0.7 (approximately).

Powers are also defined for negative and any other fractional index. Experiment with your calculator or spreadsheet.

Exponential functions and natural logarithms

e^x is a function which arises from the mathematics of constant rates of growth. You should have a button for it on your calculator, and on a spreadsheet the function will be EXP(X) or something similar. Some examples (all rounded to two decimal places):

$$e^1 = 2.72$$

$$e^3 = 20.09$$

$$e^{-1.6} = 0.20$$

The inverse function to e^x is the natural logarithm of x, $log_e(x)$ or $ln(x)$. This is useful if you want to know what *x* is if, for example,

$$e^x = 0.5$$

You can find the answer by using natural logarithms:

$$x = log_e(0.5) = -0.69315$$

This can be checked by working out $e^{-0.69315}$. It should be very close to 0.5.